

DEPARTMENT OF COMPUTER SCIENCE AND
ENGINEERING
INDIAN INSTITUTE OF TECHNOLOGY MADRAS
CHENNAI – 600036

Advanced Encryption Schemes from Novel Lattice Assumptions

A Thesis

Submitted by

SIMRAN KUMARI

For the award of the degree

Of

DOCTOR OF PHILOSOPHY

May 2025



DEPARTMENT OF COMPUTER SCIENCE AND
ENGINEERING
INDIAN INSTITUTE OF TECHNOLOGY MADRAS
CHENNAI – 600036

Advanced Encryption Schemes from Novel Lattice Assumptions

A Thesis

Submitted by

SIMRAN KUMARI

For the award of the degree

Of

DOCTOR OF PHILOSOPHY

May 2025

For my brother.

THESIS CERTIFICATE

This is to undertake that the Thesis titled **ADVANCED ENCRYPTION SCHEMES FROM NOVEL LATTICE ASSUMPTIONS**, submitted by me to the Indian Institute of Technology Madras, for the award of **Doctor of Philosophy**, is a bonafide record of the research work done by me under the supervision of **Prof. Shweta Agrawal**. The contents of this Thesis, in full or in parts, have not been submitted to any other Institute or University for the award of any degree or diploma.

Chennai 600036

Date: May 11, 2025

Simran Kumari



Prof. Shweta Agrawal

Research advisor

Professor

Department of CSE

IIT Madras

LIST OF PUBLICATIONS

I. PUBLICATIONS FROM THE THESIS

1. Shweta Agrawal, Simran Kumari, Shota Yamada. Pseudorandom FE and iO with Applications. In TCC, 2025. (To appear).
2. Shweta Agrawal, Simran Kumari, Shota Yamada. Attribute Based Encryption for Turing Machines from Lattices. In CRYPTO, 2024. https://doi.org/10.1007/978-3-031-68382-4_11.
3. Shweta Agrawal, Simran Kumari, Anshu Yadav, Shota Yamada. Trace and Revoke with Optimal Parameters from Polynomial Hardness. In EUROCRYPT, 2023. https://doi.org/10.1007/978-3-031-30620-4_20.

II. OTHER MANUSCRIPTS

1. Shweta Agrawal, Simran Kumari, Ryo Nishimaki. Laconic Pre-Constrained Encryption. Cryptology ePrint Archive. <https://eprint.iacr.org/2024/1294>

ACKNOWLEDGEMENTS

This PhD journey has been one of the most defining chapters of my life – filled with moments of excitement, self-doubt, and perseverance. There were days of triumph and days of frustration, but through it all, I was fortunate to have people who believed in me even when I struggled to believe in myself. I take this opportunity to thank all who stood by me, encouraged me, and contributed to my growth during this PhD journey.

First and foremost, I express my gratitude to Prof. Shweta Agrawal, without whom this journey would not have been possible. From the very beginning, it was her immense efforts that set me on this path, and throughout, she has been a constant source of support, patience, and inspiration. Even through my many missteps, she was consistently patient and kind – never giving up on me, always encouraging me to grow, and guiding me with clarity through both professional and personal challenges. Her commitment to the growth of her students, willingness to help anyone who seeks it, dedication to strengthening the cryptography community in India, and relentless efforts to support women in science are truly admirable. If I can contribute even a fraction of what she has, it would be one of my greatest achievements.

I sincerely thank all my co-authors: Prof. Shweta Agrawal, Dr. Ryo Nishimaki from NTT Tokyo, Dr. Anshu Yadav from IST Austria, and Dr. Shota Yamada from AIST Tokyo. Shweta has introduced me to many exciting research problems – several of which now form the foundation of this thesis. Her sharp insights and engaging discussions have constantly pushed me to think deeper, and I know I will continue learning from her. I am greatly indebted to Anshu, Shota, and Ryo for enriching my research journey with many insightful technical discussions and guidance in writing papers. I would also like to express my gratitude to Dr. Venkata Koppula, Dr. Ryo Nishimaki, and Dr. Sikhar Patranabis for their valuable discussions during my internships at IIT Delhi, NTT Tokyo, and IBM Research Lab Bangalore, respectively.

I am grateful to my doctoral committee (DC) members: Prof. John Ebenezer Augustine, Dr. Venkata Koppula, Prof. Krishna M. Sivalingam, and Dr. Aishwarya Thiruvengadam – for their valuable guidance and feedback. A special thanks to Mrs. Jeyanthi Arumugam for her support with administrative matters, making numerous processes much smoother. I would also like to acknowledge Google Research for funding my travel to EUROCRYPT 2023 and CRYPTO 2024.

Beyond academics, IIT Madras has gifted me not just friendships, but bonds that feel like home. Among them, Anshu holds a special place. She has been there from my very first steps in cryptography – guiding me through my first proof, my first research paper, and many more firsts that followed. But beyond equations and theorems, she has been a constant – a friend whose presence brought comfort and strength. Akhil has been another anchor through this journey. From deep cryptographic discussions to impromptu conversations about music and drama, he has been a steady presence, always ready with unfiltered opinions and unwavering support. I cannot imagine my PhD life without the warmth and laughter that Anshu and Akhil brought into it. I also cherish the memories I have shared with Nibedita, Aditya and Ishika – late-night snack runs, discovering amazing food across Chennai and Bengaluru, shopping detours that were supposed to be *quick*, and conversations that effortlessly stretched into hours. Our shared meals and small getaways brought so much joy and lightness to an otherwise overwhelming journey. I would also like to thank Antonio and Rohin for various fun outings and coffee breaks.

I am grateful to the IIT Madras campus – my silent companion through this journey. From cycling through its quiet roads at night with friends – to long conversations under the open sky, the campus has been a space for reflection and connection. Some of my most cherished moments have been walking on its scenic paths – sometimes in deep discussions with Anshu, sometimes alone, lost in thought, seeking clarity and calmness.

Finally, my deepest gratitude goes to my family – my parents, my brother, and my

husband – for being my foundation throughout this journey. They have been my biggest cheerleaders, standing by me through moments of self-doubt, and giving me the strength to keep going. My brother has been one of the calmest, most thoughtful human in my life. With endless patience, he has listened to multiple instances of my overcooked jargon, pulled me out of more self-dug mental pits than I can count, and quietly shielded me from anything that might distract me, just so I could stay focused. He has always encouraged me to look at things more positively and shared books for my growth (many of which I *do* plan to finish!). My husband, too, has been a steady and loving presence – kind, understanding, and always there. His quiet support and care have meant more to me than I can put into words.

For every suggestion that shaped my thoughts, every correction that deepened my understanding, and every gesture of faith that lifted me – thank you. Your impact is woven into this journey.

ABSTRACT

KEYWORDS functional encryption; multi-input functional encryption; attribute-based encryption; predicate encryption; indistinguishability obfuscation; laconic obfuscation; pseudorandom functionality; evasive LWE; broadcast, trace and revoke ; Turing Machines

As digital systems grow more interconnected, there is an increasing need for encryption mechanisms that do more than simply protect data – they must also control *who* can access *what* and *under what conditions*. Traditional public-key encryption follows an “all-or-nothing” paradigm, where possession of a secret key grants full access to the message. This model is inadequate for settings such as cloud computing or data-sharing platforms, where different users may require access to different parts of the data. For instance, a manager may require access to financial records but not legal documents within the same encrypted dataset. To address this, advanced encryption primitives such as *Attribute-Based Encryption (ABE)* and *Functional Encryption (FE)* enable fine-grained access control, allowing decryption only when specific conditions are satisfied (ABE), or enabling recovery of only the output of a function on encrypted data (FE).

While these advanced primitives offer appealing forms of access control, realizing them from well-understood and quantum-safe assumptions – like the Learning With Errors (LWE) – has resisted satisfying solutions despite significant effort. This thesis explores a new pathway using the evasive LWE assumption (Wee, Eurocrypt 2022 and Tsabary, Crypto 2022), a recent strengthening of LWE that has enabled progress on long-standing open problems in the area. Our results include the following advances:

1. We construct an optimal broadcast, trace, and revoke (TR) system that allows content to be securely distributed to many users, supports revocation of users from future communications, and enables tracing of unauthorized redistribution. A desirable feature of this notion is that of embedded identities – where user identities are embedded in

their secret keys, eliminating the need for a public index-identity mapping and enhancing anonymity. We study the notion in both secret tracing (where the tracing key is private) and public tracing (where the tracing key is public) with embedded identities. In the public tracing setting, we provide a construction relying on standard polynomial-hardness assumptions, namely FE and special forms of ABE. In the secret tracing setting, we achieve a TR scheme with asymptotically optimal parameters for ciphertexts, public keys, and secret keys using Lockable Obfuscation (LO), a key-policy ABE from Boneh et al. [43] based on LWE, and a ciphertext-policy ABE for \mathcal{P} constructed by Wee [169] from evasive and tensor LWE.

2. We give the first ABE scheme for Turing machines supporting unbounded collusion from lattice assumptions. This allows policies expressed as general-purpose programs that run on unbounded-length inputs and support input-specific runtime – far beyond what circuit-based models can handle. Our construction uses LWE, evasive LWE, and a new assumption which we call the “circular tensor LWE” assumption. Towards our ABE for Turing machines, we obtain the first ciphertext-policy ABE for circuits of unbounded depth and size from the same assumptions.

3. We introduce and construct the first compact FE scheme for pseudorandom functionalities, under standard LWE and private-coin evasive LWE. Intuitively, a pseudorandom functionality means that the output of the circuit is indistinguishable from uniform for *every* input seen by the adversary. We demonstrate the power of our tool by achieving optimal key-policy and ciphertext-policy ABE for unbounded-depth circuits, improving upon prior work by: (i) replacing the circular evasive LWE used by Hsieh, Lin, and Luo [122] with plain evasive LWE, (ii) removing the need for circular tensor LWE from our prior work, and (iii) achieving asymptotically optimal parameters.

4. Finally, we compile our pseudorandom FE scheme into the first multi-input functional encryption (MIFE), where multiple users encrypt their inputs independently and a

functional key can compute a function jointly on them, and indistinguishability obfuscation (iO), which hides everything about a program except its input-output behavior, for pseudorandom functionalities. Our MIFE scheme relies on LWE and (private-coin) evasive LWE for constant arity functions, and a strengthening of evasive LWE for polynomial arity. Thus, we obtain the first MIFE and iO schemes for a nontrivial functionality from conjectured post-quantum assumptions.

These results demonstrate that evasive LWE offers a powerful ground for building expressive, efficient, and (conjectured) quantum-secure cryptographic primitives. Finally, we discuss recent attacks on evasive LWE and their repercussions. Following the first online posting of our prFE work, a series of counterexamples to evasive LWE emerged [19, 121, 85], showing that the intuition behind evasive LWE – specifically, that it avoids the zeroizing regime – is not universally valid. However, we show that by taking suitable precautions, it is possible to recover the security of our constructions. We view these attacks as a constructive step toward a deeper understanding of evasive LWE. In our opinion, constructions from disciplined new assumptions are important to make meaningful progress on long-standing problems that have resisted solutions from standard assumptions despite much effort.

CONTENTS

	Page
ACKNOWLEDGEMENTS	i
ABSTRACT	v
LIST OF TABLES	xv
LIST OF FIGURES	xvii
NOTATION	xix
CHAPTER 1 INTRODUCTION	1
1.1 Motivation	1
1.2 Overview of the Thesis	4
1.3 Organization of Thesis	8
CHAPTER 2 PRELIMINARIES	9
2.1 Lattices and Discrete Gaussians	9
2.1.1 Lattice Trapdoors	10
2.1.2 Hardness Assumptions	11
2.2 Attribute Based Encryption	14
2.3 Pseudorandom Function	18
CHAPTER 3 BROADCAST, TRACE AND REVOKE WITH OPTIMAL PARAMETERS FROM POLYNOMIAL HARDNESS	19
3.1 Introduction	19
3.1.1 Prior Work: Embedded Identity Trace and Revoke.	23
3.1.2 Our Results	23
3.1.3 Technical Overview	26
3.2 Preliminaries	44
3.2.1 Symmetric Key Encryption	44
3.2.2 Functional Encryption	45
3.3 Attribute Based Encryption	50
3.3.1 Key-Policy ABE by Boneh et al. [43]	51
3.3.2 Lockable Obfuscation	54
3.3.3 Laconic Oblivious Transfer	55
3.4 Revocable Predicate Encryption	57
3.5 Public-key RPE from FE and LWE	60
3.5.1 Construction	60
3.5.2 Security	64
3.5.3 Alternate Construction using LOT	85

3.6	Revocable Mixed Functional Encryption	89
3.6.1	Definition	89
3.6.2	Construction	91
3.6.3	Security	96
3.7	Secret Key RPE from Evasive and Tensor LWE	109
3.7.1	Construction	109
3.7.2	Security	112
3.8	Embedded Identity Trace and Revoke	118
3.8.1	Indexed Trace and Revoke with Embedded Identity	121
3.8.2	Bounded Trace and Revoke with Embedded Identity	122
3.8.3	Unbounded Trace and Revoke with Embedded Identity	124
3.9	Indexed Trace and Revoke with Embedded Identity	125
3.9.1	Construction	125
3.9.2	Security	129
3.10	Bounded Trace and Revoke with Embedded identity	141
3.10.1	Construction	141
3.10.2	Security	145
3.11	Unbounded Trace and Revoke with Embedded Identities	153
3.11.1	Construction	153
3.11.2	Security	157
3.12	Extension to Super-Polynomial Size Revocation List	162

CHAPTER 4 ATTRIBUTE BASED ENCRYPTION FOR TURING MACHINES FROM LATTICES 165

4.1	Introduction	165
4.1.1	Our Results	166
4.1.2	Technical Overview	168
4.2	Preliminaries	176
4.2.1	Garbled Circuits	176
4.2.2	Identity-Based Encryption	178
4.2.3	Turing Machines	180
4.2.4	Attribute Based Encryption	183
4.2.5	Tensors	184
4.2.6	Hardness Assumptions	185
4.2.7	GSW Homomorphic Encryption and Evaluation	185
4.2.8	BGG+ Homomorphic Evaluation Procedures	186
4.3	Bootstrapping Randomized Homomorphic Evaluation	188
4.3.1	Preparation	189
4.3.2	Noise Removal for Randomized Encoding	191
4.3.3	Randomized Bootstrapping a.k.a Structure Restoration	191
4.4	Ciphertext Policy ABE for Unbounded Depth Circuits	197
4.4.1	Construction	197
4.4.2	Our Assumption	201
4.4.3	Security Proof	202
4.5	Generic compiler: ABE for Turing Machines and NL	209
4.5.1	Generalized Bundling of Functionality	210

4.5.2	ABE for Turing Machines	220
4.5.3	ABE for NL	224
CHAPTER 5	COMPACT PSEUDORANDOM FUNCTIONAL ENCRYPTION FROM EVASIVE LWE	231
5.1	Introduction	231
5.1.1	Our Results	233
5.1.2	Additional Prior Work	235
5.1.3	Recent Attacks on Evasive LWE and Repercussions.	236
5.1.4	Technical Overview	240
5.2	Preliminaries	267
5.2.1	Blind Garbled Circuit	268
5.2.2	Blind Batch Encryption	270
5.2.3	Poly-Input Obfuscation for Pseudorandom Functionalities	272
5.3	Functional Encryption for Pseudorandom Functionalities	274
5.3.1	Definition	274
5.3.2	Construction	277
5.3.3	Security Proof for Pseudorandom Functionalities	282
5.3.4	Basing Security on Variant of Circular Evasive LWE ([122])	290
5.4	Laconic Pseudorandom Poly-Domain Obfuscation	299
5.4.1	Definition	299
5.4.2	Construction	301
5.4.3	Security Proof	305
5.5	Partial-Hiding prFE for Unbounded Depth with Optimal Parameters	309
5.5.1	Definition	309
5.5.2	Construction	313
5.5.3	Security	319
5.5.4	Reducing the Dependency on Private Input Length	325
5.5.5	Handling Longer Output	331
5.6	KP-ABE/PE for Unbounded Depth Circuits with Optimal Parameters	332
5.6.1	Construction of KP-ABE with Optimal Parameters	332
5.6.2	Security of KP-ABE	334
5.6.3	Predicate Encryption with Optimal Parameters	337
5.6.4	Extending the Message Space	338
5.7	Compiling KP-ABE to CP-ABE using prFE	339
5.7.1	Construction	339
5.7.2	Security	342
Appendix	346
5.A	Blind Batch Encryption from LWE	346
5.A.1	Basic Scheme from LWE	347
5.A.2	Bootstrapping the Basic Scheme	352
5.B	Bootstrapping AB-LFE to KP-ABE with unbounded depth using prFE	354
5.B.1	Attribute Based Laconic Functional Encryption	355
5.B.2	Construction of kpABE with Unbounded Depth	357
5.B.3	Using the ablFE from HLL	362

CHAPTER 6	PSEUDORANDOM MULTI-INPUT FUNCTIONAL ENCRYPTION AND APPLICATIONS	365
6.1	Introduction	365
6.1.1	Our Results	368
6.1.2	Technical Overview	370
6.2	Preliminaries	390
6.2.1	Hardness Assumptions	390
6.2.2	Puncturable Pseudorandom Functions	392
6.2.3	Pseudorandom Functional Encryption	393
6.2.4	Predicate Encryption	395
6.2.5	Multi-Input Predicate Encryption	396
6.2.6	ID-Based Non-Interactive Key Exchange	398
6.3	Multi-Input FE for Pseudorandom Functionalities	400
6.3.1	Definition	400
6.3.2	Construction for n -input prFE	402
6.3.3	Security Proof for General n	408
6.3.4	Security Proof for Constant n (with Weaker Assumption)	418
6.4	Multi-Input Predicate Encryption for Polynomial Arity for P	420
6.4.1	Construction	420
6.4.2	Security	423
6.5	Indistinguishability Obfuscation for Pseudorandom Functionalities	428
6.5.1	Definition	428
6.5.2	Construction	430
6.5.3	Security	431
6.6	Polynomial Domain IO for Pseudorandom Functionalities	434
6.6.1	Definition	434
6.6.2	Construction for Fixed Input Domain	437
6.6.3	Construction for Flexible Input Domain	441
6.6.4	Extending the Output Length.	447
6.7	Instantiating the Random Oracles Using prIO	448
6.7.1	Full-Domain Hash Signatures (Selectively Secure) from prIO	448
6.7.2	Discussion about Other Applications.	450
6.8	ID-Based Non-Interactive Key Exchange	451
6.8.1	Construction	451
6.8.2	Security Proof	452
Appendix	457
6.A	Pseudorandom FE with Stronger Security	457
6.A.1	Proof for Non-Uniform κ -prCT Security	457
CHAPTER 7	EVASIVE LWE: CLASSES, ATTACKS, AND REPERCUSSIONS	461
7.1	Public-Coin Evasive LWE	462
7.2	Private-coin Binding Evasive LWE.	463
7.2.1	Attacks	464

7.3	Private-coin Hiding Evasive LWE.	467
7.3.1	Attacks and Repercussions	468
7.4	Circular Small-Secret Evasive LWE	470
7.5	Contrived Functionality Attacks.	472
7.6	Our Perspective.	473
CHAPTER 8 CONCLUSIONS		475
BIBLIOGRAPHY		477
CURRICULUM VITAE		495
DOCTORAL COMMITTEE		497

LIST OF TABLES

Table	Caption	Page
3.1	State of the art with Public Traceability.	24
3.2	State of the Art with Secret Traceability.	26
5.1	State of the Art: Attribute Based Encryption	233

LIST OF FIGURES

Figure	Caption	Page
3.1	Overview of our constructions.	43
3.2	Function to compute kpABE ciphertexts.	63
3.3	Compute and Compare function CC	94
3.4	Security Experiment: Exp-Ind-TR	121
3.5	Security Experiment: Exp-BD-TR	123
3.6	Security Experiment: Exp-TR	124
3.7	Comparison Function $f[j, \ell, b]$	126
3.8	Index Tracing	128
3.9	Identity Tracing	128
3.10	Algorithm Bnd-isGoodDecoder	143
3.11	Algorithm Bnd-Subtrace	144
3.12	Oracle \tilde{D}	144
3.13	Algorithm isGoodDecoder for Unbounded EITR	155
3.14	Algorithm: SubTrace for Unbounded EITR	155
3.15	Oracle \tilde{D} for Unbounded EITR	156
4.1	Function F	194
4.2	Circuit $U_{i,x,t}$	221
4.3	Circuit $U_{i,M}$	223
4.4	Circuit $U_{i,x,t,s}$	227
4.5	Circuit $U_{i,M}$	228
5.1	Description of the Sampler for Evasive LWE	284
5.2	Description of the Sampler for circular small-secret evasive LWE	295
5.3	The Circuit $C[\text{dig}, E, \text{sd}](i)$	304
5.4	Circuit to compute garbled labels corresponding to \mathbf{x}_{priv}	316
5.5	The Circuit $F_1[\mathbf{r}]$	316
5.6	The Circuit $F_2[\mathbf{r}, \text{dig}_C]$	317
6.1	Function F_i	405
6.2	Full Domain Hash	449
6.3	Full Domain Hash*	449
6.4	Description of the circuit $F[\text{sd}, \widetilde{\text{sd}}, g^\beta, g^\gamma]$	454

NOTATION

$[m, n]$, for $m, n \in \mathbb{N}$	$\{i : m \leq i \leq n\}$
$[n]$, for $n \in \mathbb{N}$	$\{1, \dots, n\}$
M (bold capital letter)	Matrix M
v (bold small letter)	Vector v
$\mathcal{D}_{\Lambda, \sigma}$	Discrete Gaussian distribution over lattice Λ with std. deviation σ and mean 0
$\log x$	Base 2 logarithm of x
$\text{negl}(\lambda)$	A negligible function of λ
$\ \mathbf{v}\ $	ℓ_∞ norm of vector v
$\text{poly}(\lambda)$	Polynomial in λ
PPT	Probabilistic Polynomial Time
$\mathbf{1}_{\ell \times n}$ (resp. $\mathbf{0}_{\ell \times n}$)	A matrix of dimensions $\ell \times n$, having each entry as 1 (resp. 0)
$\mathbf{1}_\ell$ (resp. $\mathbf{0}_\ell$)	Vector $(1, \dots, 1) \in \mathbb{Z}^\ell$ (resp. $(0, \dots, 0) \in \mathbb{Z}^\ell$)
$\mathbf{x} \parallel \mathbf{y}$ (resp. $\mathbf{X} \parallel \mathbf{Y}$)	Horizontal concatenation of vectors x and y (resp. matrices X and Y)
$D_1 \approx_c D_2$	Distributions D_1 and D_2 are computationally indistinguishable
$D_1 \approx_s D_2$	Distributions D_1 and D_2 are statistically indistinguishable
$D_1 \equiv D_2$	Distributions D_1 and D_2 are perfectly indistinguishable

CHAPTER 1

INTRODUCTION

1.1 MOTIVATION

Encryption is a fundamental tool for securing data in an increasingly interconnected world, allowing users to share sensitive information over insecure networks or storage systems. However, traditional public-key encryption enforces an “all-or-nothing” decryption model, where possession of a decryption key grants full access to the plaintext. As cloud computing, distributed systems, and data-driven applications continue to expand, concerns about privacy and controlled access to sensitive information have become more pressing. Standard encryption mechanisms fall short in scenarios requiring selective access – such as an enterprise system where employees should access records based on their roles, departments, or clearance levels, without revealing the rest of the database, or a content provider who wants to restrict users to only the portions of a dataset relevant to their subscription level. To overcome these limitations, advanced cryptographic primitives such as *Attribute-Based Encryption (ABE)* [157, 116] and *Functional Encryption (FE)* [157, 46] provide fine-grained access control, enabling decryption only under specified conditions (ABE), or permitting recovery of specific functions of the encrypted data (FE).

While these primitives offer powerful capabilities, constructing them for expressive class of functionalities from well-understood and quantum-safe assumptions has remained a significant challenge. Ideally, cryptographic security is most compelling when it rests on standard computational assumptions. The ideal paradigm ensures that breaking a cryptographic scheme is no easier than solving a well-studied hard problem – such as factoring, discrete logarithms, or learning with errors (LWE). This not only instills confidence in security but also creates a meaningful feedback loop: any successful attack

on the scheme advances our understanding of the underlying computational hardness. However, not all cryptographic constructions fit neatly into this framework. Several advanced primitives, such as functional encryption for general circuits, have remained out of reach when relying solely on well-established lattice assumptions, despite significant efforts by the community. This highlights the need for principled new assumptions that can push the boundaries of theoretical cryptography.

Evasive LWE. A promising development in this direction is the *evasive* LWE assumption, introduced independently by Wee [169] and Tsabary [163]. This assumption was designed to support the security of cryptographic primitives – such as witness encryption and broadcast encryption – that had long resisted reductions to standard lattice-based hardness assumptions. While techniques from lattice-based cryptography have yielded candidate constructions for these primitives, formal security proofs under standard assumptions remained elusive. Evasive LWE was proposed to bridge this gap, offering a principled strengthening of standard LWE that enables security proofs for a broader class of functionalities.

To recall, the standard LWE assumption posits that for a random matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$, a secret vector $\mathbf{s} \in \mathbb{Z}_q^n$, and an error vector \mathbf{e} drawn from a discrete Gaussian distribution, the distribution $(\mathbf{A}, \mathbf{A}^\top \mathbf{s} + \mathbf{e})$ is computationally indistinguishable from (\mathbf{A}, \mathbf{u}) , where \mathbf{u} is uniform in \mathbb{Z}_q^m .

The evasive LWE assumption strengthens this by allowing the adversary additional auxiliary information beyond standard LWE samples. Roughly, it asserts that if the *precondition*

$$(\mathbf{B}, \mathbf{P}, \mathbf{s}^\top \mathbf{B} + \mathbf{e}_\mathbf{B}, \mathbf{s}^\top \mathbf{P} + \mathbf{e}_\mathbf{P}, \text{aux}) \approx_c (\mathbf{B}, \mathbf{P}, \$, \$, \text{aux})$$

holds (where $\$$ denotes uniform), then so does the corresponding *postcondition*

$$(\mathbf{B}, \mathbf{P}, \mathbf{s}^\top \mathbf{B} + \mathbf{e}_\mathbf{B}, \mathbf{B}^{-1}(\mathbf{P}), \text{aux}) \approx_c (\mathbf{B}, \mathbf{P}, \$, \mathbf{B}^{-1}(\mathbf{P}), \text{aux}).$$

Here, $\mathbf{B}^{-1}(\mathbf{P})$ refers to a low-norm matrix \mathbf{K} satisfying $\mathbf{BK} = \mathbf{P} \bmod q$.

This models a setting where the adversary learns LWE samples (e.g., $\mathbf{s}^\top \mathbf{B} + \mathbf{e}_\mathbf{B}$) and a trapdoor preimage \mathbf{K} , enabling computation of $\mathbf{s}^\top \mathbf{P} + \mathbf{e}_\mathbf{B} \mathbf{K}$. The core heuristic behind evasive LWE is that the term $\mathbf{s}^\top \mathbf{P} + \mathbf{e}_\mathbf{B} \mathbf{K}$ is *flooded* by the high-entropy component $\mathbf{s}^\top \mathbf{P}$ and hence appears pseudorandom. The structure in $\mathbf{e}_\mathbf{B} \mathbf{K}$ becomes irrelevant under this masking, rendering the trapdoor useless. This shifts the setting away from the “zeroizing” regime, where small-norm relations over the integers enable attacks.

The evasive LWE assumption is broadly categorized into two variants: the *public-coin* and *private-coin* formulations. In the public-coin variant, proposed by Wee [169], the sampler Samp outputs matrices \mathbf{B} , \mathbf{P} , and auxiliary information \mathbf{aux} , where \mathbf{aux} includes all the random coins used by Samp. This version suffices for constructing primitives such as optimal broadcast encryption [169] and multi-authority ABE [166]. Vaikuntanathan, Wee, and Wichs [164] introduced the *private-coin* variant, in which the sampler’s randomness is hidden from the adversary. This stronger formulation has enabled a broader range of powerful constructions, including witness encryption and null-iO [163, 164], unbounded-depth KP-ABE [122] and CP-ABE [13], adaptively sound SNARGs [146, 131], and functional encryption and iO for pseudorandom functionalities [12, 14, 65].

The hope behind evasive LWE was to facilitate progress on long-standing, presumably “intermediate” cryptographic goals – such as constructing broadcast encryption from lattices. Its simplicity and generality made it a compelling foundation for achieving concrete advances in primitives that had resisted significant progress for over a decade. This thesis explores the broadening of the cryptographic landscape enabled by evasive LWE, demonstrating how this assumption can be leveraged to construct expressive, efficient, and provably secure primitives in the (conjectured) post-quantum regime.

1.2 OVERVIEW OF THE THESIS

In this section, we briefly outline the contributions presented in this thesis, following the chronological order in which the works were developed.

Broadcast, Trace, and Revoke (TR) with Embedded Identities. In a broadcast, trace and revoke system an encryptor can specify a list $L \subseteq N$ of revoked users so that (i) users in L can no longer decrypt ciphertexts, (ii) ciphertext size is independent of L , (iii) a pirate decryption box supports tracing of compromised users. The “holy grail” of this line of work is a construction which resists unbounded collusions, achieves all parameters (including public and secret key) sizes independent of $|L|$ and $|N|$, and is based on polynomial hardness assumptions. A desirable feature in this notion is the ability to embed user identities directly into secret keys [151]. This eliminates the need to store an index-to-identity mapping and improves user anonymity, particularly in the public trace setting. Our contribution in this domain is as follows.

1. *Public Trace Setting:* We provide a construction which achieves optimal parameters, supports embedding identities (from an exponential space), relies on polynomial hardness assumptions, namely compact functional encryption (FE) and a key-policy attribute based encryption (ABE) with special efficiency properties, and enjoys adaptive security with respect to the revocation list. The prior work [151] achieved optimal parameters and embedded identities, relied on the stronger tool of indistinguishability obfuscation and achieved only selective security with respect to the revocation list.
2. *Secret Trace Setting:* We provide the first construction with optimal ciphertext, public and secret key sizes and embedded identities from any assumption outside Obfustopia. Our construction relies on Lockable Obfuscation which can be constructed using LWE [111, 172] and two ABE schemes: the kpABE scheme by [43] which relies on LWE and the cpABE scheme by [169] which was constructed using the new assumptions called (*public-coin*) *evasive and tensor* LWE.

Moreover, by relying on subexponential security of LWE, both our constructions can also support a *super-polynomial* sized revocation list, so long as it allows efficient representation and membership testing. Ours is the first work to achieve this, to the best of our knowledge.

Attribute-Based Encryption for Turing Machines. In an ABE scheme, a ciphertext encodes a message m and a public attribute \mathbf{x} , while a secret key embeds a function f ;

decryption reveals m if and only if $f(\mathbf{x}) = 1$. Two common variants are key-policy (kpABE), where f is in the key, and ciphertext-policy (cpABE), where f is in the ciphertext. Typically, f is represented as a circuit.

While there has been significant progress in constructing ABE for general circuits [116, 107, 43, 26, 169, 122], the circuit model imposes rigid input sizes and worst-case evaluation time. To overcome these issues, several works [165, 102, 32, 15, 136, 103, 104, 17, 18, 142] have explored ABE for uniform computation models. Without relying on iO or compact FE, the best known constructions in this regime support non-deterministic log-space Turing machines from pairings [142]. In the post-quantum setting, the strongest result [17] supports non-deterministic finite automata under LWE, but only in the symmetric-key setting. We give the first ABE scheme for Turing machines with unbounded collusion resistance from lattice assumptions. The encryptor encodes an attribute \mathbf{x} , a time bound t , and a message m into the ciphertext. The secret key embeds a Turing machine M , and decryption returns m if and only if $M(\mathbf{x}) = 1$.

We achieve the first (conjectured) post-quantum (i) ABE for NL from LWE, (public-coin) evasive LWE, and tensor LWE; and (ii) an ABE for all Turing machines from LWE, (private-coin) evasive LWE, and a new assumption we call circular tensor LWE, which incorporates circularity into tensor LWE. As a stepping stone for ABE for TM, we also obtain the first cpABE for circuits of unbounded depth and size from the same assumptions.

FE and family for pseudorandom functionality. Functional encryption (FE) allows decryption of $f(\mathbf{x})$ from a ciphertext for input \mathbf{x} and a secret key for function f , revealing nothing else. FE is a powerful theoretical tool and can be used to construct advanced primitives, most notably indistinguishability obfuscation (iO), which is considered “crypto-complete” [34, 90, 129]. To obtain iO from FE, the scheme must satisfy *compactness*, meaning ciphertexts should be sublinear in the size of the

supported circuits. There has been substantial research effort in the community for instantiating FE (or directly iO) from well-understood assumptions, leading to a sequence of exciting results [128, 129, 5, 20, 171, 96, 82, 129, 130, 154]. The breakthrough work of Jain, Lin and Sahai [129] finally obtained the first construction of compact FE for P from standard assumptions. This has been subsequently improved by [130, 154]. However, these constructions rely heavily on pairings, which are quantum insecure and limit the diversity of underlying assumptions.

In the lattice world, several candidates for compact FE and iO exist [5, 20, 171, 96, 82, 120, 126], but they either rely on heuristics or the underlying assumptions have been broken. This raises a central open question: *Can we construct compact functional encryption for any nontrivial functionality from simple lattice assumptions?*

We introduce and construct the first compact FE scheme for pseudorandom functionalities, under standard LWE and private-coin evasive LWE. Intuitively, a pseudorandom functionality means that the output of the circuit is indistinguishable from uniform for *every* input seen by the adversary. We demonstrate the power of our tool by achieving optimal key-policy and ciphertext-policy ABE for unbounded-depth circuits, improving upon prior work by: (i) replacing the circular evasive LWE used by Hsieh, Lin, and Luo [122] with plain evasive LWE, (ii) removing the need for circular tensor LWE from our prior work [13], and (iii) achieving asymptotically optimal parameters. Finally, we compile our pseudorandom FE scheme into the first multi-input functional encryption (MIFE), where multiple users encrypt their inputs independently and a functional key can compute a function jointly on them, and indistinguishability obfuscation (iO), which hides everything about a program except its input-output behavior, for pseudorandom functionalities. Our MIFE scheme relies on LWE and (private-coin) evasive LWE for constant arity functions, and a strengthening of evasive LWE for polynomial arity. Thus, we obtain the first MIFE and iO schemes for a nontrivial functionality from conjectured post-quantum assumptions. Our tools of MIFE and iO for pseudorandom functionalities

appear quite powerful and yield extremely simple constructions when used in applications. We believe they provide a new pathway for basing “extreme” cryptography, which has so far required full fledged iO , on the presumably weaker evasive LWE in the post quantum regime.

Attack on Evasive and Repercussions. Evasive LWE has been studied in two regimes: the *private-coin* setting, where random coins used by the sampler is hidden, and the *public-coin* setting, where it is revealed to the adversary. While early counterexamples in the private-coin case [164] relied on contrived auxiliary information, no known attacks existed against the public-coin variants, including circular evasive LWE [122]. This lent credibility to evasive LWE as a “middle ground” assumption between standard LWE and stronger assumptions needed for iO .

Recent works [19, 121, 85] present new attacks showing the vulnerabilities of evasive LWE (even in the public-coin setting), particularly when maliciously crafted samplers or contrived functionalities are allowed. Specifically, [19, 121] demonstrate that a non-black-box use of pseudorandom functions can violate the post-condition while maintaining the pre-condition of the initial-version of the assumption used by our prFE scheme. These attacks also apply to the circular, small-secret variant of evasive LWE [122]. The works [65, 19] show impossibility results for pseudorandom functionalities under contrived, self-referential function classes. However, these are analogous to known limitations in other idealized models like ROM [68] and VBB [34], and do not apply to natural functionalities used in our constructions. The attacks by [66] target settings where either \mathbf{B} or \mathbf{P} is hidden. In our case, \mathbf{B} is public, and \mathbf{P} is computable from auxiliary information which is public.

Countermeasures As discussed in [19], several refinements to the evasive LWE assumption can be made to avoid known counterexamples arising from malicious samplers. One natural approach leverages the fact that, in cryptographic constructions, the key generation

algorithm is an honest party that holds the master secret key. Since the sampler is invoked by the key generator, it is reasonable, and often realistic, to assume that the function representation used in key generation can be enforced to follow a canonical structure. Furthermore, [19] outlines refinements to the evasive LWE assumption that exclude such malicious samplers. We follow their approach and restrict the class of samplers used in our constructions to avoid these vulnerabilities.

Our Perspective. We view these attacks as an important step in clarifying the scope and limitations of evasive LWE. Rather than discarding the assumption outright, we believe it should be refined to disallow malicious samplers or pathological functionalities. Carefully designed “safe zones” can preserve its utility – especially for functionalities that remain out of reach from standard assumptions. Finally, we remark that proposing principled new assumptions, by its very nature, is highly non-trivial and it is unrealistic to expect the perfect formulation in the very first attempt. In our judgment, a balance of caution and risk is beneficial in this context.

1.3 ORGANIZATION OF THESIS

The rest of the thesis is organized as follows: we provide some preliminaries, used commonly in multiple chapters, in the thesis, in Chapter 2. In Chapter 3 we present our constructions of broadcast, trace, and revoke with embedded identities in both public and secret trace settings. In Chapter 4 we construct the first lattice-based ABE for Turing machines with unbounded collusions and also obtain the first ciphertext-policy ABE for unbounded-depth circuits. Chapter 5 introduces and constructs the notion of FE for pseudorandom functionalities, leading to optimal ABE and predicate encryption schemes. In Chapter 6 we define and construct the notion of multi-input FE and indistinguishability obfuscation (iO) for pseudorandom functionalities, and demonstrates several applications. We also discuss the recent attacks on evasive and possible countermeasures in Chapter 7. We conclude the thesis in Chapter 8.

CHAPTER 2

PRELIMINARIES

In this chapter, we provide some preliminaries commonly used in this thesis. Additional preliminaries are given in respective chapters.

2.1 LATTICES AND DISCRETE GAUSSIANS

In this section, we recall some facts on lattices and discrete Gaussian distributions.

Definition 2.1 (Lattice). An m -dimensional lattice Λ is a discrete additive subgroup of \mathbb{R}^m . For an integer $n < m$ and a rank n matrix $\mathbf{B} \in \mathbb{R}^{m \times n}$, $\Lambda(\mathbf{B}) = \{\mathbf{B}\mathbf{x} : \mathbf{x} \in \mathbb{Z}^n\}$ is the lattice generated by integer linear combinations of columns of matrix \mathbf{B} . The matrix \mathbf{B} is called a *basis* of the lattice.

Definition 2.2 (Integral lattice). An m -dimensional integral lattice Λ is a full-rank subgroup of \mathbb{Z}^m .

Among these lattices are the “ q -ary” lattices defined as follows: for any integer $q \geq 2$ and any $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$, we define

$$\Lambda_q^\perp(\mathbf{A}) := \{\mathbf{e} \in \mathbb{Z}^m : \mathbf{A} \cdot \mathbf{e} = \mathbf{0} \bmod q\}.$$

For a vector $\mathbf{u} \in \mathbb{Z}_q^n$, we define the following coset of $\Lambda_q^\perp(\mathbf{A})$:

$$\Lambda_q^{\mathbf{u}}(\mathbf{A}) := \{\mathbf{e} \in \mathbb{Z}^m : \mathbf{A} \cdot \mathbf{e} = \mathbf{u} \bmod q\}.$$

We have $\Lambda_q^{\mathbf{u}}(\mathbf{A}) = \Lambda_q^\perp(\mathbf{A}) + \mathbf{t}$ for any \mathbf{t} such that $\mathbf{A} \cdot \mathbf{t} = \mathbf{u} \bmod q$.

Definition 2.3 (Gaussian distribution). For any vector $\mathbf{c} \in \mathbb{R}^m$ and any real $s > 0$, the (spherical) Gaussian function with standard deviation parameter s and center \mathbf{c} is defined

as:

$$\forall \mathbf{x} \in \mathbb{R}^m, \rho_{s,\mathbf{c}}(\mathbf{x}) = \text{Exp} \left(-\frac{\pi \|\mathbf{x} - \mathbf{c}\|^2}{s^2} \right).$$

The Gaussian distribution is $\mathcal{D}_{s,\mathbf{c}}(\mathbf{x}) = \rho_{s,\mathbf{c}}(\mathbf{x})/s^m$.

The (spherical) *discrete Gaussian distribution* over a lattice $\Lambda \subseteq \mathbb{R}^m$, with standard deviation parameter $s > 0$ and center \mathbf{c} is defined as:

$$\forall \mathbf{x} \in \Lambda, \mathcal{D}_{\Lambda,s,\mathbf{c}} = \frac{\rho_{s,\mathbf{c}}(\mathbf{x})}{\rho_{s,\mathbf{c}}(\Lambda)},$$

where $\rho_{s,\mathbf{c}}(\Lambda) = \sum_{\mathbf{x} \in \Lambda} \rho_{s,\mathbf{c}}(\mathbf{x})$. When $\mathbf{c} = \mathbf{0}$, we omit the subscript \mathbf{c} .

2.1.1 Lattice Trapdoors

Let us consider a matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$. For all $\mathbf{V} \in \mathbb{Z}_q^{n \times m'}$, we let $\mathbf{A}^{-1}(\mathbf{V})$ be an output distribution of $\text{SampZ}(\gamma)^{m \times m'}$ conditioned on $\mathbf{A} \cdot \mathbf{A}^{-1}(\mathbf{V}, \gamma) = \mathbf{V}$; where $\text{SampZ}(\gamma)$ is a sampling algorithm for the truncated discrete Gaussian distribution over \mathbb{Z} with parameter $\gamma > 0$ whose support is restricted to $z \in \mathbb{Z}$ such that $|z| \leq \sqrt{n}\gamma$. A γ -trapdoor for \mathbf{A} is a trapdoor that enables one to sample from the distribution $\mathbf{A}^{-1}(\mathbf{V}, \gamma)$ in time $\text{poly}(n, m, m', \log q)$ for any \mathbf{V} . We slightly overload notation and denote a γ -trapdoor for \mathbf{A} by \mathbf{A}_γ^{-1} .

The following properties had been established in a long sequence of works [98, 71, 6, 7, 148, 57].

Lemma 2.1 (Properties of Trapdoors). *Lattice trapdoors exhibit the following properties.*

1. Given \mathbf{A}_τ^{-1} , one can obtain $\mathbf{A}_{\tau'}^{-1}$ for any $\tau' \geq \tau$.
2. Given \mathbf{A}_τ^{-1} , one can obtain $[\mathbf{A} \parallel \mathbf{B}]_\tau^{-1}$ and $[\mathbf{B} \parallel \mathbf{A}]_\tau^{-1}$ for any \mathbf{B} .
3. There exists an efficient procedure $\text{TrapGen}(1^n, 1^m, q)$ that outputs $(\mathbf{A}, \mathbf{A}_{\tau_0}^{-1})$ where $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ for some $m = O(n \log q)$ and is 2^{-n} -close to uniform, where $\tau_0 = \omega(\sqrt{n \log q \log m})$.

Useful Lemmata.

Lemma 2.2 (tail and truncation of $\mathcal{D}_{\mathbb{Z},\gamma}$). *There exists $B_0 \in \Theta(\sqrt{\lambda})$ such that*

$$\Pr[x \leftarrow \mathcal{D}_{\mathbb{Z},\gamma} : |x| > \gamma B_0(\lambda)] \leq 2^{-\lambda} \quad \text{for all } \gamma \geq 1 \text{ and } \lambda \in \mathbb{N}.$$

Lemma 2.3 (Smudging Lemma [166]). *Let λ be a security parameter. Take any $a \in \mathbb{Z}$ where $|a| \leq B$. Suppose $\gamma \geq B\lambda^{\omega(1)}$. Then the statistical distance between the distributions $\{z : z \leftarrow \mathcal{D}_{\mathbb{Z},\gamma}\}$ and $\{z + a : z \leftarrow \mathcal{D}_{\mathbb{Z},\gamma}\}$ is $\text{negl}(\lambda)$.*

Lemma 2.4 (Leftover Hash Lemma). *Fix some $n, m, q \in \mathbb{N}$. The leftover hash lemma states that if $m \geq 2n \log q$, then for $\mathbf{A} \leftarrow \mathbb{Z}_q^{n \times m}$, $\mathbf{x} \leftarrow \{0, 1\}^m$ and $\mathbf{y} \leftarrow \mathbb{Z}_q^n$ the statistical distance between $(\mathbf{A}, \mathbf{A} \cdot \mathbf{x})$ and (\mathbf{A}, \mathbf{y}) is negligible. More concretely, it is bounded by $q^n \sqrt{2^{1-m}}$.*

2.1.2 Hardness Assumptions

Assumption 2.5 (The LWE Assumption). Let $n = n(\lambda)$, $m = m(\lambda)$, and $q = q(\lambda) > 2$ be integers and $\chi = \chi(\lambda)$ be a distribution over \mathbb{Z}_q . We say that the $\text{LWE}(n, m, q, \chi)$ hardness assumption holds if for any PPT adversary \mathcal{A} we have

$$|\Pr[\mathcal{A}(\mathbf{A}, \mathbf{s}^\top \mathbf{A} + \mathbf{e}^\top) \rightarrow 1] - \Pr[\mathcal{A}(\mathbf{A}, \mathbf{v}^\top) \rightarrow 1]| \leq \text{negl}(\lambda)$$

where the probability is taken over the choice of the random coins by the adversary \mathcal{A} and $\mathbf{A} \leftarrow \mathbb{Z}_q^{n \times m}$, $\mathbf{s} \leftarrow \mathbb{Z}_q^n$, $\mathbf{e} \leftarrow \chi^m$, and $\mathbf{v} \leftarrow \mathbb{Z}_q^m$. We also say that $\text{LWE}(n, m, q, \chi)$ problem is subexponentially hard if the above probability is bounded by $2^{-n^\epsilon} \cdot \text{negl}(\lambda)$ for some constant $0 < \epsilon < 1$ for all PPT \mathcal{A} .

As shown by previous works [155, 57], if we set $\chi = \text{SampZ}(\gamma)$, the $\text{LWE}(n, m, q, \chi)$ problem is as hard as solving worst case lattice problems such as gapSVP and SIVP with approximation factor $\text{poly}(n) \cdot (q/\gamma)$ for some $\text{poly}(n)$. Since the best known algorithms for 2^k -approximation of gapSVP and SIVP run in time $2^{\tilde{O}(n/k)}$, it follows that the above $\text{LWE}(n, m, q, \chi)$ with noise-to-modulus ratio 2^{-n^ϵ} is likely to be (subexponentially) hard for some constant ϵ .

'p[];/-=]

Assumption 2.6 (Evasive LWE). [169, 22, 19] Let $n, m, t, m', q \in \mathbb{N}$ be parameters and λ be a security parameter. Let χ and χ' be parameters for Gaussian distributions. For Samp that outputs

$$\mathbf{S} \in \mathbb{Z}_q^{m' \times n}, \mathbf{P} \in \mathbb{Z}_q^{n \times t}, \text{aux} \in \{0, 1\}^*$$

on input 1^λ and for PPT adversaries \mathcal{A}_0 and \mathcal{A}_1 , we define the following advantage functions:

$$\text{Adv}_{\mathcal{A}_0}^{\text{PRE}}(\lambda) \stackrel{\text{def}}{=} \Pr[\mathcal{A}_0(\mathbf{B}, \mathbf{S}\mathbf{B} + \mathbf{E}, \mathbf{S}\mathbf{P} + \mathbf{E}', \text{aux}) = 1] - \Pr[\mathcal{A}_0(\mathbf{B}, \mathbf{C}_0, \mathbf{C}', \text{aux}) = 1] \quad (2.1)$$

$$\text{Adv}_{\mathcal{A}_1}^{\text{POST}}(\lambda) \stackrel{\text{def}}{=} \Pr[\mathcal{A}_1(\mathbf{B}, \mathbf{S}\mathbf{B} + \mathbf{E}, \mathbf{K}, \text{aux}) = 1] - \Pr[\mathcal{A}_1(\mathbf{B}, \mathbf{C}_0, \mathbf{K}, \text{aux}) = 1] \quad (2.2)$$

where

$$(\mathbf{S}, \mathbf{P}, \text{aux}) \leftarrow \text{Samp}(1^\lambda),$$

$$\mathbf{B} \leftarrow \mathbb{Z}_q^{n \times m},$$

$$\mathbf{C}_0 \leftarrow \mathbb{Z}_q^{m' \times m}, \mathbf{C}' \leftarrow \mathbb{Z}_q^{m' \times t},$$

$$\mathbf{E} \leftarrow \mathcal{D}_{\mathbb{Z}, \chi}^{m' \times m}, \mathbf{E}' \leftarrow \mathcal{D}_{\mathbb{Z}, \chi'}^{m' \times t}$$

$$\mathbf{K} \leftarrow \mathbf{B}^{-1}(\mathbf{P}) \text{ with standard deviation } O(\sqrt{m \log(q)}).$$

We say that the *evasive* LWE (EvLWE) assumption with respect to the sampler class \mathcal{SC} holds if for every PPT $\text{Samp} \in \mathcal{SC}$ and \mathcal{A}_1 , there exists another PPT \mathcal{A}_0 and a polynomial $Q(\cdot)$ such that

$$\text{Adv}_{\mathcal{A}_0}^{\text{PRE}}(\lambda) \geq \text{Adv}_{\mathcal{A}_1}^{\text{POST}}(\lambda)/Q(\lambda) - \text{negl}(\lambda) \quad \text{and} \quad \text{Time}(\mathcal{A}_0) \leq \text{Time}(\mathcal{A}_1) \cdot Q(\lambda). \quad (2.3)$$

We conjecture that for reasonable class of samplers, the evasive LWE assumption holds. In particular, we conjecture that our sampler $\text{Samp}_{\text{prFE}}(1^\lambda)$ used for the security proof of

our prFE for natural class of functions should be in the secure class of samplers \mathcal{SC} for which the evasive LWE holds.

Remark 1. In the above definition, all the LWE error terms are chosen from the same distribution $D_{\mathbb{Z},\chi}$. However, in our security proof, we often consider the case where some of LWE error terms are chosen from $D_{\mathbb{Z},\chi}$ and others from $D_{\mathbb{Z},\chi'}$ with different $\chi \gg \chi'$. The evasive LWE assumption with such a mixed noise distribution is implied by the evasive LWE assumption with all LWE error terms being chosen from $D_{\mathbb{Z},\chi}$ as above definition, since if the precondition is satisfied for the latter case, that for the former case is also satisfied. To see this, it suffices to observe that we can convert the distribution from $D_{\mathbb{Z},\chi'}$ into that from $D_{\mathbb{Z},\chi}$ by adding extra Gaussian noise.

In multiple of our security proofs, we may require the auxiliary information to include terms dependent on \mathbf{S} . Furthermore, we may want to prove the pseudorandomness of such auxiliary information. The following lemma from [22] enables this. In the lemma, we separate the auxiliary information into two parts \mathbf{aux}_1 and \mathbf{aux}_2 , where \mathbf{aux}_1 is typically the part dependent on \mathbf{S} . The lemma roughly says that \mathbf{aux}_1 is pseudorandom in the post condition distribution, if it is pseudorandom in the precondition distribution.

Lemma 2.7 (Lemma 3.4 in [22]). *Let $n, m, t, m', q \in \mathbb{N}$ be parameters and λ be a security parameter. Let χ and χ' be Gaussian parameters. Let Samp be a PPT algorithm that takes as input 1^λ and outputs*

$$\mathbf{S} \in \mathbb{Z}_q^{m' \times n}, \mathbf{aux} = (\mathbf{aux}_1, \mathbf{aux}_2) \in \mathcal{S} \times \{0, 1\}^* \text{ and } \mathbf{P} \in \mathbb{Z}_q^{n \times t}$$

for some set \mathcal{S} . Furthermore, we assume that there exists a public deterministic poly-time algorithm Reconstruct that allows to derive \mathbf{P} from \mathbf{aux}_2 , i.e. $\mathbf{P} = \text{Reconstruct}(\mathbf{aux}_2)$.

We introduce the following advantage functions:

$$\mathcal{A}_{\text{Adv}}^{\text{PRE}'}(\lambda) \stackrel{\text{def}}{=} \Pr[\text{Adv}(\mathbf{B}, \mathbf{SB} + \mathbf{E}, \mathbf{SP} + \mathbf{E}', \mathbf{aux}_1, \mathbf{aux}_2) = 1] - \Pr[\text{Adv}(\mathbf{B}, \mathbf{C}_0, \mathbf{C}', \mathbf{c}, \mathbf{aux}_2) = 1] \quad (2.4)$$

$$\mathcal{A}_{\text{Adv}}^{\text{POST}'}(\lambda) \stackrel{\text{def}}{=} \Pr[\text{Adv}(\mathbf{B}, \mathbf{SB} + \mathbf{E}, \mathbf{K}, \text{aux}_1, \text{aux}_2) = 1] - \Pr[\text{Adv}(\mathbf{B}, \mathbf{C}_0, \mathbf{K}, \mathbf{c}, \text{aux}_2) = 1] \quad (2.5)$$

where

$$(\mathbf{S}, \text{aux} = (\text{aux}_1, \text{aux}_2), \mathbf{P}) \leftarrow \text{Samp}(1^\lambda),$$

$$\mathbf{B} \leftarrow \mathbb{Z}_q^{n \times m}$$

$$\mathbf{C}_0 \leftarrow \mathbb{Z}_q^{m' \times m}, \mathbf{C}' \leftarrow \mathbb{Z}_q^{m' \times t}, \mathbf{c} \leftarrow \mathcal{S}$$

$$\mathbf{E} \leftarrow \mathcal{D}_{\mathbb{Z}, \chi}^{m' \times m}, \mathbf{E}' \leftarrow \mathcal{D}_{\mathbb{Z}, \chi}^{m' \times t}$$

$$\mathbf{K} \leftarrow \mathbf{B}^{-1}(\mathbf{P}) \text{ with standard deviation } O(\sqrt{m \log(q)}).$$

Then, under the Evasive-LWE (cited above in Assumption 2.6) with respect to a sampler $\text{Samp} \in \mathcal{SC}$, for a sampler class \mathcal{SC} , if $\mathcal{A}_{\text{Adv}}^{\text{PRE}'}(\lambda)$ is negligible for any PPT adversary Adv , so is $\mathcal{A}_{\text{Adv}}^{\text{POST}'}(\lambda)$ for any PPT adversary Adv .

2.2 ATTRIBUTE BASED ENCRYPTION

We define both ciphertext policy attribute-based encryption (cpABE) and key policy attribute-based encryption (kpABE) in a unified form below.

Let $R : \mathcal{X} \times \mathcal{Y} \rightarrow \{0, 1\}$ be a relation where \mathcal{X} and \mathcal{Y} denote “ciphertext attribute” and “key attribute” spaces, respectively. Ideally, we would like to have an ABE scheme that handles the relation R directly, where we can encrypt w.r.t any ciphertext attribute $x \in \mathcal{X}$ and can generate a secret key for any key attribute $y \in \mathcal{Y}$. However, in many cases, we are only able to construct a scheme that poses restrictions on the ciphertext attribute space and key attribute space. To capture such restrictions, we introduce a parameter prm and consider subsets of the domains $\mathcal{X}_{\text{prm}} \subseteq \mathcal{X}$ and $\mathcal{Y}_{\text{prm}} \subseteq \mathcal{Y}$ specified by it and the function R_{prm} defined by restricting the function R on $\mathcal{X}_{\text{prm}} \times \mathcal{Y}_{\text{prm}}$.

An attribute-based encryption (ABE) scheme for $R = \{R_{\text{prm}} : \mathcal{X}_{\text{prm}} \times \mathcal{Y}_{\text{prm}} \rightarrow \{0, 1\}\}_{\text{prm}}$

and a message space \mathcal{M} is defined by the following algorithms.

- **Setup**($1^\lambda, \text{prm}$) \rightarrow (mpk, msk). The setup algorithm takes as input the unary representation of the security parameter λ and a parameter prm and outputs a master public key mpk and a master secret key msk.
- **Enc**(mpk, X, μ) \rightarrow ct_X . The encryption algorithm takes as input a master public key mpk, a ciphertext attribute $X \in \mathcal{X}_{\text{prm}}$, and a message $\mu \in \mathcal{M}$. It outputs a ciphertext ct_X .
- **KeyGen**(msk, Y) \rightarrow sk_Y . The key generation algorithm takes as input the master secret key msk and a key attribute $Y \in \mathcal{Y}_{\text{prm}}$. It outputs a private key sk_Y .
- **Dec**(mpk, $\text{sk}_Y, Y, \text{ct}_X, X$) \rightarrow μ or \perp . The decryption algorithm takes as input the master public key mpk, a private key sk_Y , private key attribute $Y \in \mathcal{Y}_{\text{prm}}$, a ciphertext ct_X and ciphertext attribute $X \in \mathcal{X}_{\text{prm}}$. It outputs the message μ or \perp which represents that the ciphertext is not in a valid form.

Definition 2.4 (Correctness). An ABE scheme for relation family R is correct if for all prm, $X \in \mathcal{X}_{\text{prm}}, Y \in \mathcal{Y}_{\text{prm}}$ such that $R(X, Y) = 0$, and for all messages $\mu \in \mathcal{M}$,

$$\Pr \left[\begin{array}{l} (\text{mpk}, \text{msk}) \leftarrow \text{Setup}(1^\lambda, \text{prm}), \\ \text{sk}_Y \leftarrow \text{KeyGen}(\text{msk}, Y), \\ \text{ct}_X \leftarrow \text{Enc}(\text{mpk}, X, \mu) : \\ \text{Dec}(\text{mpk}, \text{sk}_Y, Y, \text{ct}_X, X) \neq \mu \end{array} \right] = \text{negl}(\lambda)$$

where the probability is taken over the coins of Setup, KeyGen, and Enc.

Definition 2.5 (Sel-IND security for ABE). For an ABE scheme $\text{ABE} = \{\text{Setup}, \text{Enc}, \text{KeyGen}, \text{Dec}\}$ for a relation family $R = \{R_{\text{prm}} : \mathcal{X}_{\text{prm}} \times \mathcal{Y}_{\text{prm}} \rightarrow \{0, 1\}\}_{\text{prm}}$ and a message space \mathcal{M} and an adversary \mathcal{A} , let us define Sel-IND security game $\text{Exp}_{\text{ABE}, \mathcal{A}}(1^\lambda)$ as follows.

1. \mathcal{A} outputs prm and the challenge ciphertext attribute $X^* \in \mathcal{X}_{\text{prm}}$.
2. **Setup phase:** On input $1^\lambda, \text{prm}$, the challenger samples $(\text{mpk}, \text{msk}) \leftarrow \text{Setup}(1^\lambda, \text{prm})$ and gives mpk to \mathcal{A} .
3. **Query phase:** During the game, \mathcal{A} adaptively makes the following queries, in an arbitrary order. \mathcal{A} can make unbounded many key queries, but can make only single challenge query.

a) **Key Queries:** \mathcal{A} chooses an input $Y \in \mathcal{Y}_{\text{prm}}$. For each such query, the challenger replies with $\text{sk}_Y \leftarrow \text{KeyGen}(\text{msk}, Y)$.

b) **Challenge Query:** At some point, \mathcal{A} submits a pair of equal length messages $(\mu_0, \mu_1) \in \mathcal{M}^2$ to the challenger. The challenger samples a random bit $b \leftarrow \{0, 1\}$ and replies to \mathcal{A} with $\text{ct}_{X^*} \leftarrow \text{Enc}(\text{mpk}, X^*, \mu_b)$.

We require that $R(X^*, Y) = 1$ holds for any Y such that \mathcal{A} makes a key query for Y in order to avoid trivial attacks.

4. **Output phase:** \mathcal{A} outputs a guess bit b' as the output of the experiment.

We define the advantage $\text{Adv}_{\text{ABE}, \mathcal{A}}^{\text{Sel-IND}}(1^\lambda)$ of \mathcal{A} in the above game as

$$\text{Adv}_{\text{ABE}, \mathcal{A}}^{\text{Sel-IND}}(1^\lambda) := |\Pr[\text{Exp}_{\text{ABE}, \mathcal{A}}(1^\lambda) = 1 | b = 0] - \Pr[\text{Exp}_{\text{ABE}, \mathcal{A}}(1^\lambda) = 1 | b = 1]|.$$

The ABE scheme ABE is said to satisfy **Sel-IND** security (or simply *selective security*) if for any stateful PPT adversary \mathcal{A} , there exists a negligible function $\text{negl}(\cdot)$ such that $\text{Adv}_{\text{ABE}, \mathcal{A}}^{\text{Sel-IND}}(1^\lambda) = \text{negl}(\lambda)$.

We can consider the following stronger version of the security where we require the ciphertext to be pseudorandom.

Definition 2.6 (Sel-INDr security for ABE). We define **Sel-INDr** security game similarly to **Sel-IND** security game except that the adversary \mathcal{A} chooses single message μ instead of (μ_0, μ_1) at the challenge phase and the challenger returns $\text{ct} \leftarrow \text{Enc}(\text{mpk}, X^*, \mu)$ if $b = 0$ and a random ciphertext $\text{ct} \leftarrow \mathcal{CT}$ from a ciphertext space \mathcal{CT} if $b = 1$. Here, we assume that uniform sampling from the ciphertext space \mathcal{CT} is possible without any parameter other than the security parameter λ . We define the advantage $\text{Adv}_{\text{ABE}, \mathcal{A}}^{\text{Sel-INDr}}(1^\lambda)$ of the adversary \mathcal{A} accordingly and say that the scheme satisfies **Sel-INDr** security if the quantity is negligible.

We also consider the very selective notion of security.

Definition 2.7 (VerSel-IND security for ABE). We define **VerSel-IND** security game similarly to **Sel-IND** security game except that the adversary \mathcal{A} outputs the key queries Y_1, \dots, Y_Q , where Q is the number of key queries made by \mathcal{A} , along with the challenge ciphertext attribute X^* in the beginning of the security game. We define the advantage $\text{Adv}_{\text{ABE}, \mathcal{A}}^{\text{VerSel-IND}}(1^\lambda)$ of the adversary \mathcal{A} accordingly and say that the scheme satisfies

VerSel-IND security if the quantity is negligible.

Definition 2.8 (VerSel-INDr security for ABE). We define VerSel-IND security game similarly to Sel-INDr security game except that the adversary \mathcal{A} outputs the key queries Y_1, \dots, Y_Q , where Q is the number of key queries made by \mathcal{A} , along with the challenge ciphertext attribute X^* in the beginning of the security game. We define the advantage $\text{Adv}_{\text{ABE}, \mathcal{A}}^{\text{VerSel-INDr}}(1^\lambda)$ of the adversary \mathcal{A} accordingly and say that the scheme satisfies VerSel-INDr security if the quantity is negligible.

Definition 2.9 (Ada-IND security for ABE). We define Ada-IND security game similarly to Sel-IND security game except that the adversary \mathcal{A} can choose the challenge ciphertext attribute X^* adaptively. We define the advantage $\text{Adv}_{\text{ABE}, \mathcal{A}}^{\text{Ada-IND}}(1^\lambda)$ of the adversary \mathcal{A} accordingly and say that the scheme satisfies Ada-IND security if the quantity is negligible.

In the following, we recall definitions of various ABEs by specifying the relation.

Ciphertext-policy Attribute Based encryption (cpABE). We define cpABE for circuit class $\{C_{\ell(\lambda), d(\lambda)}\}_\lambda$ by specifying the relation. Here, $C_{\ell(\lambda), d(\lambda)}$ is a set of circuits with binary output whose input length is $\ell(\lambda)$ and the depth is at most $d(\lambda)$. Note that we do not pose any restriction on the size of the circuits. We define $A_\lambda^{\text{cpABE}} = C_{\ell(\lambda), d(\lambda)}$ and $B_\lambda^{\text{cpABE}} = \{0, 1\}^\ell$. Furthermore, we define the relation R_λ^{cpABE} as

$$R_\lambda^{\text{cpABE}}(C, \mathbf{x}) = C(\mathbf{x}).$$

Key-policy Attribute Based encryption (kpABE). To define kpABE for circuits, we simply swap key and ciphertext attributes in cpABE for circuits. More formally, to define kpABE for circuits, we define $A_\lambda^{\text{kpABE}} = \{0, 1\}^\ell$ and $B_\lambda^{\text{kpABE}} = C_{\ell(\lambda), d(\lambda)}$. We also define $R_\lambda^{\text{kpABE}} : A_\lambda^{\text{kpABE}} \times B_\lambda^{\text{kpABE}} \rightarrow \{0, 1\}$ as

$$R_\lambda^{\text{kpABE}}(\mathbf{x}, C) = C(\mathbf{x}).$$

The above relations also holds for circuit class $\{C_{\ell(\lambda)}\}_\lambda$ which is the set of circuits with binary output whose input length is $\ell(\lambda)$ and the depth is unbounded.

2.3 PSEUDORANDOM FUNCTION

Definition 2.10 (Pseudorandom functions (PRF)). A pseudorandom function (PRF) family $\mathcal{F} = \{F^K\}_{K \in \mathcal{K}}$ with a key space \mathcal{K} , a domain \mathcal{X} , and a range \mathcal{Y} is a function family that consists of functions $F^K : \mathcal{X} \rightarrow \mathcal{Y}$. Let \mathcal{R} be a set of functions consisting of all functions whose domain and range are \mathcal{X} and \mathcal{Y} respectively. A PRF family \mathcal{F} is said to be secure if for any PPT adversary \mathcal{A} , the following condition holds,

$$|\Pr[\mathcal{A}^{F^K(\cdot)}(1^\lambda) = 1] - \Pr[\mathcal{A}^{R(\cdot)}(1^\lambda) = 1]| \leq \text{negl}(\lambda),$$

where $K \leftarrow \mathcal{K}$ and $R \leftarrow \mathcal{R}$.

CHAPTER 3

BROADCAST, TRACE AND REVOKE WITH OPTIMAL PARAMETERS FROM POLYNOMIAL HARDNESS

3.1 INTRODUCTION

Traitor Tracing. Traitor tracing (TT) schemes were first proposed by Chor, Fiat, and Naor [74] to enable content providers to trace malicious users who exploit their secret keys to construct illegal decryption boxes. More formally, a TT system is a public key encryption system comprising N users for some large polynomial N . Each user $i \in [N]$ is provided with a unique secret key sk_i for decryption, and there is a common public key pk which is used by the content distributor to encrypt content. If any collection of users attempts to create and sell a new decoding box that can be used to decrypt the content, then the tracing algorithm, given black-box access to any such pirate decoder, is guaranteed to output an index $i \in [N]$ of one of the corrupt users, which in turn allows to hold them accountable. The literature has considered both public and secret tracing, where the former requires knowledge of a secret key to run the trace procedure and the latter does not suffer from this restriction.

Broadcast Encryption. Broadcast Encryption [86] (BE) introduced by Fiat and Naor, is also an N user system which supports an encrypted broadcast functionality. In BE, a content provider can transmit a single ciphertext over a broadcast channel so that only an authorized subset $S \subseteq N$ of users can decrypt and recover the message. More formally, each user $i \in [N]$ is provided with a unique decryption key sk_i and a ciphertext ct_m for a message m also encodes an authorized list S so that sk_i decrypts ct_m if and only if $i \in S$. Evidently, public key encryption provides a trivial construction of BE with ciphertext of size $O(N)$ – thus, the focus in such schemes is to obtain short ciphertext,

ideally logarithmic in N .

Broadcast, Trace and Revoke. Naor and Pinkas [150] suggested a meaningful interleaving of these two functionalities so that traitors that are identified by the TT scheme can be removed from the set of authorized users in a BE scheme. To capture this, they defined the notion of “Broadcast, Trace and Revoke” (or simply “Trace and Revoke”, which we denote by TR) where the content provider in a broadcast encryption scheme includes a list L of revoked users in the ciphertext, and sk_i works to decrypt ct_L if $i \notin L$. Moreover, it is required that revocation remain compatible with tracing, so that if an adversary builds a pirate decoder that can decrypt ciphertexts encrypted with respect to L , then the tracing algorithm should be able to output a corrupt non-revoked user who participated in building the illegal decoder. Trace and revoke systems provide a functionality which is richer than a union of BE and TT, since the traitor traced by the latter must belong to the set of non-revoked users for the guarantee to be meaningful. As such, TT schemes have been challenging to construct even given TT and BE schemes.

The Quest for Optimal Parameters. All the above primitives have been researched extensively over decades, resulting in a long sequence of beautiful constructions, non-exhaustively [44, 50, 48, 49, 151, 135, 112, 113, 27, 169]. A central theme in this line of work is to achieve optimal parameters, namely optimal sizes for the ciphertext, public key and secret key (and understanding tradeoffs thereof), while still supporting unbounded collusion resistance. Towards this, the powerful hammer of indistinguishability obfuscation (iO) [34] yielded the first feasibility results for traitor tracing [49] as well as trace and revoke [151] while multilinear maps [89, 80] led to the first construction for broadcast encryption [48]. Though there has been remarkable progress in the construction of iO from standard assumptions, with the breakthrough work of Jain, Lin and Sahai [129, 130] finally reaching this goal, iO is an inherently subexponential assumption [93] because the challenger is required to check whether two circuits are functionally equivalent, which can take exponential time in general. Indeed,

all known constructions of iO assume subexponential hardness of the underlying algebraic assumptions. To address this limitation, a sequence of works [93, 15, 56, 92, 139] has sought to replace iO by polynomially hard assumptions such as functional encryption in different applications.

Optimal TT, BE and TR from Polynomial Assumptions: For traitor tracing, the first construction from standard assumptions was finally achieved by the seminal work of Goyal, Koppula and Waters [112] in the secret trace setting, from the Learning With Errors (LWE) assumption. For broadcast encryption, this goal was achieved by Agrawal and Yamada [27] from LWE and the bilinear GGM. In the standard model, Agrawal, Wichs and Yamada [24] provided a construction from a non-standard knowledge assumption on pairings, while Wee [169] provided a construction from a new assumption on lattices, called Evasive and Tensor LWE. For trace and revoke, the only construction without iO that achieves collusion resistance and optimal parameters is by Goyal, Vusirikala and Waters (GVW) [115] from *positional witness encryption* (PWE) which is a polynomial hardness assumption. However, their construction incurs an exponential loss in the security proof, requiring the underlying PWE to satisfy subexponential security. Moreover, although PWE is not an inherently subexponential assumption as are iO and witness encryption (WE), we do not currently know of any *constructions* of PWE that rely on standard polynomial hardness assumptions. In particular, [129, 130] do *not* imply PWE from polynomial hardness.

Pathway via Secret Tracing. Both the iO and PWE based constructions of TR [151, 115] achieve public tracing. Taking a lesson from TT, where optimal parameters were achieved from standard assumptions only in the secret trace setting [112], a natural approach towards optimal TR from better assumptions is to weaken the tracing algorithm to be secret key. This approach has been explored in a number of works – the current best parameters are achieved by Zhandry [174] who obtains the best known tradeoff in ciphertext, public key and secret key size. In particular, Zhandry [174] showed that all

parameters can be of size $O(N^{1/3})$ by relying on the bilinear generic group model (GGM). Note that the generic group model is a strong assumption, and indeed a construction secure in this model cannot be considered as relying on standard assumptions, since several non-standard assumptions on pairings are secure in the GGM. Prior to [174], Goyal et al. [114] provided a construction from LWE and Pairings, but their overall parameters are significantly worse – while their ciphertext can be arbitrarily small, $O(N^\epsilon)$, their public key is $O(N)$ and secret key is $O(N^c)$ for some large constant c ¹.

Thus, a central open question in TR is:

Can we construct collusion resistant Trace and Revoke with optimal parameters from concrete polynomial assumptions?

Embedding Identities. Traditionally, it was assumed that tracing the index $i \in [N]$ of a corrupt user is enough, and there is an external mapping, maintained by the content distributor or some other party which associates the number i to the identity of the user, i.e. name, national identity number and such, which is then used to ensure accountability. The work of Nishimaki, Wichs and Zhandry (henceforth NWZ) [151] argued that this assumption is problematic since it implies that a user must trust the content provider with her confidential information. Storing such a map is particularly worrisome in the setting of public tracing since the user either cannot map the recovered index to an actual person, or the index-identity map must be stored publicly.

NWZ provided an appealing solution to the above conundrum – they suggest that identifying information be embedded in the key of the user, so that if a coalition of traitors constructs a pirate decoder, the tracing algorithm can directly retrieve the identifying

¹Zhandry [174] states that the secret size in [114] is $O(N^2)$ but in fact the exponent is much larger due to the usage of arithmetic computations in NC^1 , which blows up the circuit size associated with the ABE secret keys.

information from one of the keys that was used to construct the decoder and no one needs to keep any records associating users to indices. Notably, the identities can live in an *exponential* sized space, which introduces significant challenges in the tracing procedure. Indeed, handling an exponential space in the tracing procedure is the key contribution of NWZ. They also provided constructions of traitor tracing as well as trace and revoke with embedded identities, denoted by EITT and EITR respectively, from various assumptions.

3.1.1 Prior Work: Embedded Identity Trace and Revoke.

In the *public* trace setting, the only work that achieves embedded identity trace and revoke (EITR) with full collusion resistance is that of NWZ. However, while it takes an important first step, the construction by NWZ suffers from the following drawbacks:

1. *Reliance on Subexponential Hardness Assumption.* The construction relies on indistinguishability obfuscation [34], which appears to be an inherently subexponential assumption as discussed above.
2. *Selective Security in Revocation List:* Despite relying on adaptive security of functional encryption, the notion of security achieved by their construction is selective – the adversary must announce the revocation list before making any key requests or seeing the challenge ciphertext.

In the *secret* trace setting, the work of Kim and Wu [135] achieves EITR from the subexponential Learning With Errors (LWE) assumption. However, their construction incurs a ciphertext size that grows with the size of the revocation list. Additionally, while they can achieve adaptive security with respect to the revocation list, this is either by incurring an exponential loss in the security proof, or by assuming sub-exponential security for an ingredient scheme.

3.1.2 Our Results

In our work, we provide the first constructions with optimal parameters from polynomial assumptions, which additionally support embedded identities from an exponential space. We detail our contributions below.

Work	CT	SK	PK	Trace Space	Sel/Adp	Asspn	Identities
[151]	1	1	1	Exp	Selective	Subexp (iO)	Yes
[115]	1	1	1	Poly	Adaptive	Subexp (subexp PWE)	No
This	1	1	1	Exp	Adaptive	Poly (FE and Special ABE)	Yes

Table 3.1: State of the art with Public Traceability.

Public Trace Setting. We provide a construction of Trace and Revoke with public tracing which overcomes the limitations of NWZ – (i) it relies on polynomial hardness assumptions, namely functional encryption and “special” attribute based encryption, both of which can be constructed using standard polynomial hardness assumptions [43, 129, 130] (ii) it enjoys adaptive security in the revocation list.

A detailed comparison with prior work is provided in Table 3.1.

Our Assumptions. Functional Encryption (FE) and Attribute Based Encryption (ABE) are generalizations of Public Key Encryption. In FE, a secret key corresponds to a circuit C and a ciphertext corresponds to an input x from the domain of C . Given a function key sk_C and a ciphertext ct_x , the decryptor can learn $C(x)$ and nothing else. It has been shown that FE implies iO [29, 40] albeit with exponential loss. The aforementioned work of Jain, Lin and Sahai [129, 130] provides a construction of compact FE from polynomial hardness assumptions, namely LPN, PRG in NC_0 and pairings. ABE is a special case of FE in which the input can be divided into a public and private part (x, m) and the circuit C in the secret key sk_C is only evaluated on the public part x in the ciphertext $\text{ct}_{x,m}$. The private message m is revealed by decryption if and only if $C(x) = 1$. While FE implies ABE in general, we require our underlying ABE to satisfy special

efficiency properties, which is not generically implied by FE. However, the desired ABE can be instantiated using the construction of Boneh et al. [43] which is based on LWE.

Secret Trace Setting. In the secret trace setting, we achieve the optimal size of $O(\log N)$ for ciphertext, public and secret key by relying on Lockable Obfuscation (LO) [111, 172] and two special ABE schemes – one, the key-policy scheme with special efficiency properties by Boneh et al. [43] which is based on LWE, and two, a ciphertext-policy ABE for \mathbf{P} which was recently constructed by Wee [169] using the new evasive and tensor LWE assumptions. Along the way, we show that a small modification to the TR construction by Goyal et al. [114] yields a ciphertext of size $O(\log N)$ as against their original $O(N^\epsilon)$, from LWE and pairings. However, this construction retains the large public and secret keys of their construction, which depend at least linearly on N . Our results are summarized in Table 3.2.

Our Assumptions. We remark that while FE has now been constructed from standard assumptions [129, 130], the reliance of these constructions on pairings makes it insecure in the post-quantum regime. From lattices, constructions of FE rely on strong, non-standard assumptions which are often subject to attack [5, 21, 171, 96, 82, 126]. Hence, there is an active effort in the community [169, 163, 164] to construct advanced primitives from the hardness of weaker assumptions in the lattice regime. The new assumptions by Wee, also independently discovered by Tsabary [163], are formulated for designing ciphertext-policy ABE which is much weaker than FE since ABE is an all or nothing primitive in contrast to FE. As such, these are believed to be much weaker than lattice based assumptions that have been introduced in the context of FE or iO . In particular, based on the current state of art, evasive LWE is required even for broadcast encryption in the lattice regime, and is therefore necessary for the generalization of broadcast encryption studied in this work.

Super-polynomial Revoke List. Lastly, by relying on subexponential security of LWE, both our constructions can support a *super-polynomial* sized revocation list, so long as it

Work	CT	SK	PK	Trace Space	Asspn	Identities
[114]	N^ϵ	N^c	N	Poly	LWE and Pairings	No
[174]	N^a	N^{1-a}	N^{1-a}	Poly	GGM Pairings	No
[135]	L	1	1	Exp	Subexp LWE	Yes
This work[11]	1	1	1	Exp	Evasive Tensor LWE	Yes
Modified [114]	1	N^c	N	Poly	LWE and Pairings	No

Table 3.2: State of the Art with Secret Traceability.

allows efficient representation and membership testing. Ours is the first work to achieve this, to the best of our knowledge.

3.1.3 Technical Overview

We proceed to give an overview of our techniques. We begin by defining the notion of revocable predicate encryption (RPE) in both the public and secret setting, then describe the ideas used to instantiate this primitive. Finally we outline how to upgrade public/secret RPE to build trace and revoke with embedded identities with public/secret tracing.

Revocable Predicate Encryption. NWZ introduced the notion of revocable functional encryption (RFE) and used it to construct EITR with public tracing. Subsequently, Kim and Wu [135] adapted this notion to the secret key setting, under the name of revocable predicate encryption (RPE) and used it to construct EITR with secret tracing. In this work, we extend Kim and Wu’s notion of RPE to the public key setting and use it to construct EITR with public tracing. Our notion of RPE in the public setting is similar to but weaker than RFE² – it only supports “all or nothing” decryption in contrast to

²Syntactically, RPE is “ciphertext-policy” while RFE is “key-policy”, i.e. the function is emdedded in the ciphertext in RPE as against the key in RFE.

RFE. This weaker notion nevertheless suffices to construct EITR and moreover admits constructions from weaker assumptions.

In RPE, the key generation algorithm takes as input the master secret key msk , a label $\text{lb} \in \mathcal{L}$ and an attribute $x \in \mathcal{X}$. It outputs a secret key $\text{sk}_{\text{lb},x}$. The encryption algorithm takes as input the encryption key ek , a function f , a message $m \in \mathcal{M}$, and a revocation list $L \subseteq \mathcal{L}$. It outputs a ciphertext ct . Decryption recovers m if $f(x) = 1$ and $\text{lb} \notin L$. In the public variant of RPE, ek is a public key, while in the secret variant, ek is a secret key. In the secret variant, the scheme is also required to support a public “broadcast” functionality, i.e. there exists a public encryption algorithm that allows anyone to encrypt a message with respect to the “always-accept” policy, i.e. a policy that evaluates to true for all inputs. This is analogous to the primitive of “mixed FE” introduced by [112].

In terms of security, we require RPE to satisfy message hiding and function hiding. At a high level, message hiding stipulates that an adversary cannot distinguish between encryptions of (f, m_0) and (f, m_1) as long as every key query for (lb, x) satisfies $f(x) = 0$ or $\text{lb} \in L$. Function hiding stipulates that an adversary cannot distinguish between encryptions of (f_0, m) and (f_1, m) as long as every key query for (lb, x) satisfies $f_0(x) = f_1(x)$ or $\text{lb} \in L$.

Before we describe our constructions, we highlight the chief difficulties that are inherent to designing RPE:

1. *Independence of parameter sizes from $|L|$.* A key requirement in TR schemes is that the ciphertext size should be independent of the length of the revocation list L – this constraint must also be satisfied by the underlying RPE, in both the secret and public setting. In our work, we insist that even the public and secret keys satisfy $|L|$ independence. This constraint is inherited from broadcast encryption, and is challenging to satisfy. Further, note that L must be unbounded – its length cannot be fixed during setup, which introduces additional difficulties.
2. *Encrypted Computation.* While the revocation list L need not be hidden by the ciphertext, the function f^3 in the ciphertext is required to be hidden, as formalized

³For the informed reader, this function encodes the “index” and function hiding corresponds to “index

by our function hiding requirement. Yet, this hidden function must participate in computing $f(x)$ where x is provided in the key. This requirement makes TR schemes worryingly close to collusion resistant functional encryption, an “obfustopia” primitive which we want to avoid in the secret trace setting.

Constructing Public Revocable Predicate Encryption. We proceed to describe the main ideas in constructing public RPE.

Overview of NWZ. The work of NWZ addresses the challenge of making the ciphertext size independent of $|L|$ by using a somewhere statistically binding (SSB) hash and hides the function f by using a functional encryption scheme, where f is encrypted in the ciphertext. However, they must additionally rely on iO – at a high level, this is because they require the *decryptor* to compute the SSB opening π and then run SSB verification on it (details of how SSB algorithms work are not relevant for this overview). In turn, the reason they need the decryptor to compute the opening π is because this needs both the set L and the label lb , which are available only to the decryptor – note that the encryptor has only L and the key generator has only lb . Now, since the decryptor has to compute π and run SSB verification, and since the program that computes SSB verification has some secrets, the decryptor is allowed to obtain obfuscation of this program. To implement this idea, they nest iO inside a compact FE scheme so that FE decryption outputs an iO which is then run by the decryptor on openings that it computes.

Trading iO for ABE. Above, note that the usage of iO is caused by the usage of SSB, which in turn is used to compress L . However, compression of a list has been achieved by much weaker primitives than iO in the literature of broadcast encryption – in particular, the construction of optimal broadcast encryption by Agrawal and Yamada uses the much weaker primitive of ABE (with special efficiency properties) to achieve this. However, ABE does not permit hiding anything other than a message, in particular, an ABE

hiding” in the literature.

ciphertext cannot encrypt our function f since we desire f to participate in computation. ABE only permits computation on public values, and using ABE to encode f would force f to be public which we cannot allow.

In order to get around this difficulty, we leverage the power of functional encryption (FE), which permits encrypted computation and exactly fills the gap over ABE that we require. A natural candidate for RPE would be to simply use FE to encrypt f , L and m , and encode x and lb in the secret key for a functionality which tests that $\text{lb} \notin L$, that $f(x) = 1$ and outputs m if so. Indeed, this approach using FE is folklore, and was explicitly discussed by NWZ. Yet, they end up with a construction that additionally uses SSB, iO , a puncturable PRF and secret key encryption scheme because of the requirement of size independence from $|L|$ – we do not have candidates for FE with ciphertext size independent of the public attributes. In short, ABE gives us L compactness (in some cases by encoding L in the secret key [43] and in some cases by encoding L in the ciphertext [33]) but does not hide f , whereas FE gives the opposite.

Synthesis of ABE and FE. We address this conundrum by combining the two primitives in a way that lets us get the best of both. In particular, we use ABE to check that $\text{lb} \notin L$ and use FE to compute $f(x)$. Evidently, the two steps cannot be performed independently in order to resist mix and match attacks so we use nesting, i.e. we use FE to generate ABE ciphertexts. Here, care is required, because ABE encryption takes L as input and done naively, this strategy will again induce a size dependence on L . We address this challenge by using the special ABE by Boneh et al. [43] which enjoys succinct secret keys and encoding L in the ABE secret key. In more detail, we let the RPE encryption generate $\text{ABE.sk}(C_L)$ for a circuit C_L which takes as input lb and checks that $\text{lb} \notin L$. Additionally, it generates an FE ciphertext for the function f and message m . The RPE key generator computes an FE key for a function which has (lb, x) hardwired and takes as input a function f , checks whether $f(x) = 1$ and if so, generates a fresh ABE ciphertext with attribute lb and message m . Thus the decryptor can first

compute FE decryption to recover the ABE ciphertext $\text{ABE.ct}(\text{lb}, m)$ and then use ABE decryption with $\text{ABE.sk}(C_L)$ to output m if and only if $\text{lb} \notin L$. It is easy to verify that this construction achieves optimal parameters – this is because ABE has optimal parameters and we used FE only for a simple functionality that does not involve L .

Putting it all Together. The above description is over-simplified and ignores technical challenges such as how to leverage indistinguishability based security of FE, how to generate the randomness used for ABE encryption and such others – we refer the reader to Section 3.5 for details. However, even having filled in these details, we get only a selectively secure RPE. Substantial work and several new ideas are required for adaptive security, as we discuss next.

Adaptive Security. Next, we outline our ideas to achieve adaptive security, namely where the revocation list L is chosen adaptively by the adversary. Note that to avoid complexity leveraging, we are required to rely only on the selective security of the underlying ABE – this creates multiple technical difficulties which are resolved by very carefully using specific algebraic properties of our ingredients.

Leveraging Late Generation of ABE. Our first observation is that full adaptive security of ABE may be unnecessary, since in our construction of RPE, the generation of the ABE instance is deferred until the generation of the challenge ciphertext, at which time the set of revoked users is known. This intuition turns out to be true, but via a complicated security proof as we outline next. Below, we consider the case of function hiding in the RPE ciphertext, the case of message hiding is similar.

Recall that function hiding says that two ciphertexts encoding (f_b, m, L) , where $b \in \{0, 1\}$ should be indistinguishable so long as for any requested key $\text{sk}_{\text{lb},x}$ it holds that $f_0(x) = f_1(x)$ or $\text{lb} \in L$. Note that the adversary is permitted to query for keys that allow decryption of the ciphertexts, i.e. $f_0(x) = f_1(x) = 1$.

Embedding ABE CTs in FE keys. In order to use ABE security to prove RPE security, a first (by now standard) step is to use the “trapdoor technique” [70, 28, 59], which allows us to hardwire ABE ciphertexts into FE secret keys. In the security game with the ABE challenger, the reduction submits the label lb associated with each RPE secret key as its challenge attribute and embeds the returned ABE ciphertext into the FE key. Here we immediately run into a difficulty, since in the RPE setting some ABE ciphertexts are decryptable by the adversary and we cannot leverage ABE security. Moreover, we cannot even hope to guess which keys will correspond to decryptable ABE ciphertexts since there are an unbounded polynomial number of key queries in the RPE security game. The same difficulty is faced by NWZ and is the main reason why their construction does not achieve adaptive security in the revocation list.

Polynomial Function Space Suffices for TR. To overcome this hurdle, we leverage the serendipitous fact that for the purpose of constructing TR, it suffices to construct RPE whose function space (recall that functions are encoded in the ciphertext) is only of polynomial size. This observation, which was implicitly present in [113], is abstracted and used explicitly in our proof. In particular, we can assume that the reduction algorithm knows the challenge functions (f_0, f_1) at the beginning of the game, since it can simply guess them. Now, given the secret key query (lb, x) , the reduction checks whether $f_0(x) = f_1(x)$. If yes, then there is no need to use ABE security, for the ABE ciphertexts in this case will encode the same message, and will hence be independent of the challenge bit. On the other hand, if $f_0(x) \neq f_1(x)$, then we have by the admissibility condition that $\text{lb} \in L$, even when L is not known. In this case, the reduction can use the security of the ABE without any difficulty.

Additional Hurdles Stemming from ABE Selective Security. We now highlight another challenge in the proof. For concreteness, let us consider the second key query $(\text{lb}^{(2)}, x^{(2)})$, which we assume is a pre-challenge query, and assume that $f_0(x^{(2)}) \neq f_1(x^{(2)})$. Hence, by the above discussion, we are required to use ABE security for the ciphertexts with

attribute $\text{lb}^{(2)}$. However, according to the selective definition, the reduction is required to choose the challenge attribute at the very start of the game, without even seeing the public parameters. At the same time, the reduction is required to simulate the ABE ciphertext for the first key query, before receiving the second key query from the adversary, that is, without seeing the ABE parameters, leading to an apparent impasse.

We address this issue by considering the following two cases separately: for the first query $(\text{lb}^{(1)}, x^{(1)})$, we have (1) $f_0(x^{(1)}) \neq f_1(x^{(1)})$ or (2) $f_0(x^{(1)}) = f_1(x^{(1)})$. In first case, it is tempting to think that one can simply use a hybrid argument to change the ABE ciphertext associated with each key query satisfying $f_0(x^{(i)}) \neq f_1(x^{(i)})$ for $i \in [2]$. However, this does not work as is, since the ABE ciphertext may leak information about the ABE public key. To address this, we rely on the pseudorandomness of ciphertexts in our ABE [43] due to which we are guaranteed that the ciphertext does not reveal any information about the public parameters, enabling the hybrid strategy above. To handle the second case, we change the way in which the ABE ciphertext for the first key is generated. In more detail, we stop hardwiring the the ABE ciphertext into the first key and instead generate it directly using ABE parameters. This removes the aforementioned problem since we no longer need to embed the ABE ciphertext or public key into the first FE key. To enable this idea, we introduce additional branch of trapdoor mode for the construction to separate the paths of computation for the cases $f_0(x) = f_1(x)$ and $f_0(x) \neq f_1(x)$. To handle post-challenge queries, we need to address additional challenges, which we do not describe here. We refer the reader to Section 3.5 for details.

Handling Super-polynomial Revocation List. Our construction (also the secret version, described next) organically supports super-polynomially large revocation list, something that was not known before, to the best of our knowledge. In more detail, let L be a list of super-polynomial size, such that L can be represented as a string of polynomial length and there exists a circuit C_L of polynomial size which takes as input some string lb and checks whether $\text{lb} \in L$ or not. Note that any super-polynomially large list must have

efficient representation in order to even allow various algorithms to read it. Then, the key generation of [43] can naturally encode the circuit C_L as before and the construction works as before. A subtlety that arises with super-polynomial L is that when we deal with post challenge key queries in the proof, we have to deal with the ABE queries in the order of key first and ciphertext later. With polynomial size L , this does not pose a problem because when the adversary chooses L , all the labels for which we use ABE security are in L and we can perform a hybrid argument over these labels. However, this is not possible for super polynomial L , which requires to rely on subexponentially secure LWE. Please see Section 3.12 for details.

Instantiating Public RPE. Overall, armed with the above ideas, we get a public RPE from compact FE and efficient ABE supporting exponential sized identity space and adaptive security in the revocation list L . Currently, we only know how to instantiate our desired ABE from LWE [43], whereas FE can be instantiated in multiple different ways. A natural candidate would be the FE from standard assumptions [129, 130] which relies on pairings, LPN and low depth PRG – in this case, our RPE will require the extra assumption of LWE. Another option is to instantiate FE with a post-quantum candidate [96, 144, 171, 5, 82] from non-standard strengthenings of LWE – this has the advantage that the ABE does not incur any extra assumption in the final construction. For super-polynomial L , we need subexponential hardness of LWE in either pathway to instantiation, as discussed above.

Alternative Construction Based on Laconic OT. Here, we sketch an alternative construction of RPE based on laconic OT (LOT) [73] that works when the number N of possible labels is polynomially bounded (i.e., the identity space is of polynomial size). Since LOT is known to be possible from various assumptions, this diversifies the assumptions that we need to rely on. The basic idea is to replace ABE with LOT. In more detail, the encryptor chooses LOT parameters instead of ABE parameters and computes the digest of the list of recipients (or equivalently, the list of revoked users),

which is represented as a binary string of length N with 1 for non-revoked identities. The digest, whose size is independent of N , is then embedded into the FE ciphertext. Then, FE decryption yields LOT encryption of the message for the label $lb \in [N]$, which is the label associated with the secret key, instead of ABE ciphertext. The LOT ciphertext is encrypted so that it can be decrypted only when the lb -th bit of the binary string representing the list of recipients is 1. We note that this idea does not extend for identities from exponentially large space and cannot therefore support embedded identities any more.

Revocable Predicate Encryption in Private Setting. For private revocable predicate encryption, our starting point is the work of Goyal et al. [114], who show how to combine “broadcast mixed FE” (called BMFE) together with ABE to achieve RPE (via a different abstraction which they call AugBE). They construct BMFE by adding the broadcast functionality to the primitive of mixed FE defined by [112]. They embed BMFE ciphertext into an ABE ciphertext to achieve RPE, where BMFE is constructed from LWE and ABE is instantiated using pairings.

Supporting Exponential Identity Space. To begin, we upgrade their notion of BMFE to support an exponential space of identities (which we refer to as labels) towards the goal of embedded identity trace and revoke. We refer to our notion as Revocable Mixed FE (denoted by RMFE) and construct it from LWE. Both [114] and our work start with a mixed FE scheme and add broadcast to it, but their construction builds upon the scheme based on constrained PRFs [72] while ours begins with the scheme based on Lockable Obfuscation (LO), also from [72]. Our construction of RMFE deviates significantly from theirs, and achieves significantly better secret key size – $O(\log N)$ as against $O(N)$ – in addition to supporting exponential instead of polynomial space. We describe this construction next.

Mixed FE. The notion of mixed FE was introduced by Goyal, Koppula and Waters in the context of traitor tracing [112]. Identifying and constructing this clever primitive is the key insight that enables [112] to construct traitor tracing with optimal parameters from LWE. Mixed FE is, as the name suggests, a mix of public and secret key FE. Thus, it has a secret as well as a public encryption procedure. The secret encryption procedure takes as input a function f and computes ct_f . This is decryptable by a key SK_x to recover $f(x)$. The adversary can make one query to the encryption oracle in addition to getting the challenge ciphertext for challenge (f_0, f_1) . It can also make an unbounded number of key requests so long as $f_0(x) = f_1(x)$. The public encryption algorithm computes a ciphertext for the “always accept” function, i.e. a function which evaluates to 1 for any input x . It is required that the public ciphertext be indistinguishable from the secret ciphertext.

One of the constructions of mixed FE suggested by [72] uses a secret key FE scheme (SKFE) to construct the secret encryption algorithm and leverages the power of lockable obfuscation (LO) to construct the public encryption procedure. Recall that in a lockable obfuscation scheme [111, 172] there exists an obfuscation algorithm Obf that takes as input a program C , a message m and a (random) “lock value” α and outputs an obfuscated program \tilde{P} . One can evaluate the obfuscated program on any input x to obtain as output m if $P(x) = \alpha$ and \perp otherwise. Intuitively, the idea of [72] is to wrap the FE ciphertext using LO and to define the public key encryption algorithm as outputting a simulated version of the LO obfuscated circuit, which is publicly sampleable.

In more detail, the construction works as follows. The secret key for a user with input x is an SKFE secret key $\text{SKFE.sk}(x)$. The secret ciphertext of MFE for function f is constructed as follows.

1. First, SKFE ciphertext $\text{SKFE.ct}(H_{f,\alpha})$ is generated, where α is a freshly chosen random value and $H_{f,\alpha}$ is a circuit that takes as input x and outputs α if $f(x) = 0$ and 0 otherwise.

2. Then, LO with lock value α and any message $m \neq \perp$ is used to obfuscate the circuit $\text{SKFE.Dec}(\text{SKFE.ct}(H_{f,\alpha}), \cdot)$, namely the circuit that takes as input an SKFE secret key and decrypts the hardwired ciphertext using this.

The decryption result of MFE is defined as 1 if the evaluation result of the LO circuit on the given input SKFE secret key is \perp and 0 otherwise. Correctness follows from correctness of SKFE and LO. In particular, if $f(x) = 0$, then SKFE decryption outputs α , which unlocks the LO to give m , otherwise \perp . By definition, MFE decryption will output 1 if LO outputs \perp which happens when $f(x) = 1$, and 0 otherwise.

Revocable Mixed FE. RMFE augments MFE so that the encryption algorithms (both secret and public) now include a revocation list L and the secret key additionally includes a label lb . A secret key $\text{sk}_{\text{lb},x}$ decrypts a secret ciphertext $\text{ct}_{f,L}$ to recover $f(x)$ if $\text{lb} \notin L$ and 1 otherwise. For a public ciphertext ct_L , the output of decryption is always 1 regardless of which secret key is being used. For security, we need two properties: function hiding and mode hiding. For function hiding, we require that a secret ciphertext $\text{ct}_{f_0,L}$ is indistinguishable from $\text{ct}_{f_1,L}$ if for all queries, either $f_0(x) = f_1(x)$ or $\text{lb} \in L$. For mode hiding, we require that a secret ciphertext $\text{ct}_{f,L}$ is indistinguishable from a public ciphertext ct_L . Recall that L is not required to be hidden, but we require that the parameters do not depend on $|L|$.

To extend MFE to RMFE, we retain the idea of letting the secret ciphertext be an LO obfuscated circuit and public ciphertext be the simulated LO. To incorporate the list L , we must ensure that the LO lock value α is recovered only when $f(x) = 0$ and $\text{lb} \notin L$. To do so, we consider two subsystems such that one system outputs partial decryption result α_1 only when $f(x) = 0$ and the second system outputs partial decryption result α_2 only when $\text{lb} \notin L$ such that $\alpha = \alpha_1 + \alpha_2$. We must ensure that α_1 and α_2 are user specific decryption results to avoid collusion attacks.

Note that the second subsystem, which entails L , should be constructed so that the hardwired values inside the circuit do not depend on $|L|$, but still control access to the

value α_2 depending on L . To satisfy these apparently conflicting requirements, we make use of the unique algebraic properties of the ABE construction by Boneh et al [43], as described below. For the first subsystem, we use SKFE.

In more detail, our candidate scheme is as follows.

1. **Secret Key:** The RMFE secret key consists of $\text{ABE.ct}(\text{lb}, K)$ and $\text{SKFE.ct}((x, K, R))$ where K and R are user specific random strings, lb is used as an attribute and K is the plaintext for ABE encryption.
2. **Ciphertext:** To generate RMFE ciphertext, the secret key encryption procedure is as follows:
 - It first generates $\text{ABE.sk}(C_L)$, where C_L is a circuit that takes as input a label lb and outputs 1 only when $\text{lb} \notin L$.
 - It also generates $\text{SKFE.sk}(H_{f,\alpha})$, where $H_{f,\alpha}$ takes as input (x, K, R) and outputs $K \oplus \alpha$ if $f(x) = 0$ and R if $f(x) = 1$.
 - Now, consider the circuit $CC[\text{ABE.sk}(C_L), \text{SKFE.sk}(H_{f,\alpha})]$, which takes as an input the pair $(\text{ABE.ct}, \text{SKFE.ct})$, decrypts both ABE and SKFE ciphertexts using their respective keys, and then outputs the XOR between the decryption results.
 - The final ciphertext is an LO of $CC[\text{ABE.sk}(C_L), \text{SKFE.Enc}(H_{f,K,\alpha})]$ with lock value α and any arbitrary message $m \neq \perp$.

By key compactness of [43], the size of $\text{ABE.sk}(C_L)$ is independent of $|L|$. A subtle point here is that ABE decryption is happening inside the LO and this depends on L . If the LO must process L , then the size of the LO and hence ciphertext blows up with L ! Fortunately, the algebraic structure of the ABE scheme we use [43] again comes to our rescue. At a high level, ABE decryption can be divided into an “ L -dependent” step which results in a short processed ciphertext, followed by an “ L -independent” step. Importantly, the L -dependent step does not depend on the ABE secret key which is hardwired in the LO and hence inaccessible, and can hence be performed *outside* the LO by the decryptor! The resultant short processed ciphertext can then be provided as input to the LO preventing the problematic size blowup.

RMFE Proof Overview. Next we outline some of the ideas developed for the security proof. For ease of understanding, we limit ourselves to the simpler setting where the adversary does not have access to the encryption oracle. This restriction can be removed using combinatorial tricks, similar to [72]. For security, we must argue two properties – mode indistinguishability and function hiding. The former can be established by relying on security of SKFE and LO analogously to the MFE proof in [72]. Hence, we focus on function hiding for the rest of the overview, which is subtle and requires several new ideas.

For function hiding, we must make use of the security of ABE and SKFE. Intuitively, security of SKFE guarantees that the values encoded in SKFE ciphertexts and secret keys are hidden, beyond what is revealed by decryption.⁴ First note that given a key for (lb, x) such that $f_0(x) = f_1(x)$, no information about the challenge bit is revealed by decryption, since the decryption results of SKFE are the same for both cases. The case with $f_0(x) \neq f_1(x)$ is more challenging. Let us assume $f_0(x) = 0$ and $f_1(x) = 1$. In this case, the decryption result of the challenge ciphertext is R or $K \oplus \alpha$ depending on the value of the challenge bit. Since both are random strings, it is tempting to conclude that they do not reveal any information of the challenge bit.

However, in reality, information about K is encoded in the ciphertext $\text{ABE.ct}(\text{lb}, K)$ and creates a correlation which must be handled. Indeed, a computationally unbounded attacker can learn the challenge bit by breaking open the ABE ciphertext, recovering K and then correlating it with the decryption result of SKFE. Hence, security of ABE must play a role and fortunately, we show that security of ABE suffices to overcome this difficulty. Recall that our security definition of RMFE requires that if $f_0(x) \neq f_1(x)$, then it should hold that $\text{lb} \in L$. This means that the ciphertext $\text{ABE.ct}(\text{lb}, K)$ is computationally indistinguishable from $\text{ABE.ct}(\text{lb}, 0)$, since the only ABE secret key available to the adversary is $\text{ABE.sk}(C_L)$. Now, in the adversary's view, both $K \oplus \alpha$

⁴We note that we need message *and* function hiding security for the underlying SKFE, while [72] only needs message hiding security.

and R are random strings that are independent from other parameters. Therefore, the adversary cannot obtain any information of the challenge bit from the decryption result in this case as well. For more details, please see Section 3.6.

Comparison with the BMFE by Goyal et al. [114]. We observe that both our RMFE as well as the BMFE by [114] rely solely on LWE. However, our secret key is $\text{ABE.ct}(\text{lb}, K)$ and $\text{SKFE.ct}((x, K, R))$, which has optimal size, being clearly independent of N and L . In contrast their secret key depends linearly on N . We also observe that our RMFE can support an exponentially large space of identities, while their BMFE does not.

Combining RMFE and ABE to get RPE. Finally, we nest our RMFE inside an outer ABE scheme to obtain RPE. This step is very similar to [114], but we need to use a different ABE scheme. In particular, in the construction of RPE in [114], a key policy ABE (kpABE) is used to encrypt the message m with attributes as the RMFE ciphertext along with the list L . The RPE secret key for (x, lb) is a kpABE secret key for a the RMFE decryption circuit $\text{RMFE.Dec}(\text{RMFE.sk}, \cdot, \cdot)$.

An obvious difficulty here is that encoding the attribute $(L, \text{RMFE.ct})$ in the ABE ciphertext can cause the ciphertext size to depend on the size of L . To avoid this blowup, [114] use a special kpABE which has the property that the ciphertext size is independent of the size of the attribute. They instantiate this kpABE with the scheme [33] which uses pairings⁵. However, we cannot use [33, 160] because of the following two reasons:

1. First, the ABE scheme by [33] only supports NC_1 . However, our circuit $\text{RMFE.Dec}(\text{RMFE.sk}, \cdot, \cdot)$ does not fit into NC_1 ⁶.
2. Furthermore, even if the above problem could be resolved, using [33] is problematic since their ABE has secret and public keys at least as large as $O(|L|)$. While the scheme of [114] also suffers from this blow-up, our goal is to obtain short keys, independent of $|L|$.

⁵In fact, one could instead use the kpABE constructed by [160]. This enjoys the same efficiency properties and is based on the standard DLIN assumption as against the q -type assumption of [33].

⁶The informed reader may wonder whether we can solve this issue by using preprocessing as in [114] but this does not work due to technical reasons.

The first problem cannot be resolved even if we use the ABE schemes for circuits [107, 43], since their ciphertext size also depends on $|L|$. To instantiate our ABE, we use recent construction of compact cpABE from evasive and tensor LWE [169], whose parameter sizes depend only on the input length of the circuit and are independent of its size. Armed with the above ideas, we suggest the following RPE:

1. The encryption algorithm of RPE, given m, f, L computes RMFE ciphertext encoding (f, L) and then computes $\text{cpABE.Enc}(\text{RMFE.ct}, \cdot, L, m)$.
2. The key generation algorithm RPE given (lb, x) , computes RMFE secret key for (lb, x) and outputs $\text{cpABE.sk}(\text{RMFE.sk})$.

Correctness of RPE follows from correctness of cpABE and RMFE while optimality of parameters follows from the efficiency of the underlying schemes. In particular, observe that all parameters are independent of $|L|$. Also note that evasive and tensor LWE are required only to instantiate cpABE with the desired efficiency. If future work standardizes the assumptions underlying the cpABE, our construction will inherit these assumptions. For more details, we refer the reader to Section 3.7.

Instantiating Secret RPE. Currently, the only two suitable ABE schemes that we know to instantiate our compiler are the LWE based kpABE by Boneh et al. [43] and the evasive and tensor LWE based cpABE by Wee [169]. These two ABEs give us a secret RPE scheme supporting exponential identities and with optimal parameters, from evasive and tensor LWE. Note that this construction does *not* achieve adaptive security in the revoke list. Nevertheless, it is the first construction of optimal RPE, even without embedded identities, from any assumption outside Obfustopia. Note that the usage of a non-standard assumption outside of obfustopia (in particular, only from lattice techniques) is somewhat inherent given that even broadcast encryption *without* tracing requires non-standard assumption if we instantiate it only from lattices. We are hopeful that future improvements in cpABE will yield a construction from completely standard assumptions.

Trace and Revoke with Optimal Ciphertext from LWE and Pairings. Along the way, we observe that the broadcast and trace construction provided by Goyal et al. [114], without embedded identities, can be easily modified to achieve at least optimal ciphertext size, from the same assumptions. At a high level, they construct a broadcast mixed FE from LWE with optimal ciphertext size and then nest this inside the kpABE by [33], which enjoys ciphertext size independent of the attribute length, and can support computation in NC_1 . Since their BMFE decryption does not fit into NC_1 , they preprocess the ciphertext so that part of the decryption is performed “outside”, namely, they group $\log N$ matrix tuples into c groups of $(\log N)/c$ tuples each. Then they precompute all possible $2^{(\log N)/c} = N^{1/c}$ subset-products within each group. Due to this, BMFE decryption only needs to multiply together c of the preprocessed matrices, which can be done in NC_1 so long as c is constant. Unfortunately, this step increases their ciphertext size to $O(N^\epsilon)$ for any $\epsilon > 0$ though the BMFE ciphertext size was optimal.

We observe that they are “under-using” the ciphertext size independence of [33] – in particular, while the attribute length has indeed been blown up to $O(N^\epsilon)$, this does not affect the ciphertext size of [33]. Moreover, while the attribute must also be provided outside in the clear, this part can be compressed, i.e. the preprocessing which expands the attribute to size N^ϵ can be performed by the decryptor directly by grouping and multiplying matrices as described above, and there is no need for the encryptor to provide this expanded form. Thus, their scheme tweaked with this simple modification already achieves ciphertext of optimal size, though with large secret key $O(N^c)$ for some large constant c .

Trace and Revoke from Revocable Predicate Encryption. It remains to show how to construct the final goal of trace and revoke with embedded identities. As discussed earlier, we follow [151, 135] and use the abstraction of RPE to build trace and revoke. However, to embed identities in our trace and revoke schemes, we deviate from these

works and instead build upon ideas developed by [113] (henceforth GKW) in the context of traitor tracing.

Embedded Identity Traitor Tracing (EITT) by GKW. The work of Goyal, Koppula and Waters [113] provided an alternative approach for embedding identities in traitor tracing schemes. A well known approach for constructing Traitor Tracing systems suggested by Boneh, Sahai and Waters [38] is via the intermediate primitive of *Private Linear Broadcast Encryption* (PLBE), which allows to construct a tracing algorithm that performs a linear search over the space of users to recover the traitor. Since the number of users was polynomial, this algorithm could be efficient. However, if we allow arbitrary identities to correspond to user indices then the space over which this search must be performed becomes exponential even if the number of users is polynomial, and the trace algorithm is no longer efficient. The main new idea in NWZ that enables them to handle exponentially large identity spaces is to replace a linear search over indices by a clever generalization of binary search, which efficiently solves an “oracle jump problem” which in turn suffices for tracing.

Goyal, Koppula and Waters (GKW) provided an alternate route to the problem of embedding identities. Instead of using PLBE and generalizing the search procedure, they instead extend the definition of PLBE to support embedded identities, denoted by EI-PLBE, and then used this to get a full fledged EITT scheme. This approach has the notable advantage that even if the space of identities is exponential, it can use the fact that the number of users is only *polynomial* and hence rely on only *selective* security of the underlying primitives. In particular, they demonstrate a “nested” tracing approach, where the tracing algorithm works in two steps: first, it outputs a set of indices that correspond to the users that are traitors, and then it uses each index within this set to recover the corresponding identity. Additionally, GKW provide a sequence of (increasingly stronger) TT primitives with embedded identities, namely, indexed EITT, bounded EITT and finally unbounded EITT where unbounded EITT satisfies the most general notion of

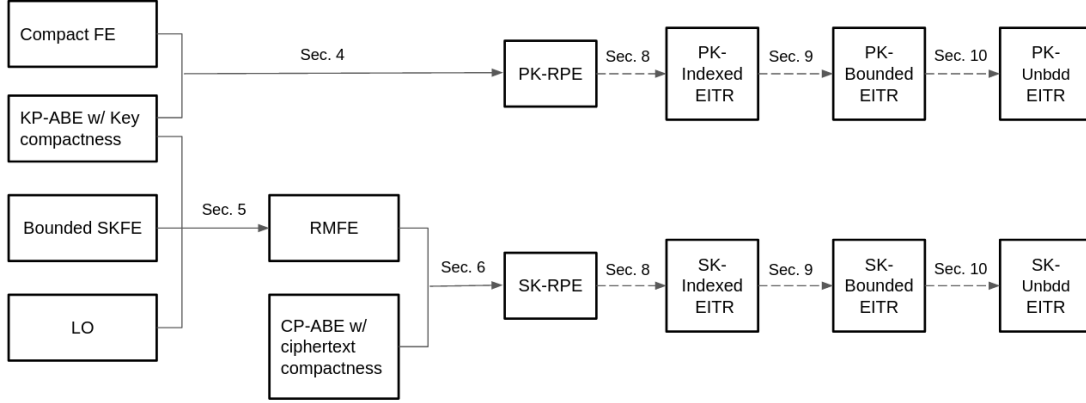


Figure 3.1: Overview of our constructions.

embedded identity traitor tracing. They also provide generic transformations between these notions, which allows to focus on the weakest notion for any new instantiation.

Embedded Identity Trace and Revoke (EITR). We adapt the approach of **GKW** and show how to use their nested approach to trace embedded identities even in the more challenging setting of trace and revoke. As in their case, this lets us use polynomial hardness assumptions in obtaining EITR, in contrast to **NWZ**. We also define indexed, bounded and unbounded EITR and provide transformations between them. Our definitions as well as transformations are analogous to **GKW** albeit care is required to incorporate the revoke list L in each step and adapt the definitions and proofs of security accordingly. We then construct indexed EITR using RPE, and obtain unbounded EITR via our generic conversions.

We note that our framework unifies the approaches of Kim and Wu [135] who used the framework of RPE in the context of TR and that of **GKW** who used the framework of EI-PLBE in the context of TT, to obtain EITR. This unification yields a clean abstraction which can be used for both public and secret key settings. We believe this framework is of independent interest. We refer the reader to Sections 3.8, 3.9, 3.10 and 3.11 for details. An overview of our constructions is provided in Figure 3.1.

Organization of the Chapter. We provide some additional preliminaries in Section 3.2. We define RPE in Section 3.4 and provide a construction of public-key RPE in Section 3.5. We give our construction of RMFE in Section 3.6. Then we construct secret-key RPE using RMFE in Section 3.7. We define different versions of trace and revoke with embedded identities in Section 3.8 and construct indexed-EITR in Section 3.9, bounded-EITR in Section 3.10 and unbounded-EITR in Section 3.11. Our goal is unbounded-EITR and other variants are introduced as intermediate goals. Secret and public tracing unbounded-EITR can be obtained by applying the conversions in Section 3.9, 3.10, and 3.11 to the secret-key and public key RPE, respectively. We show how to extend our constructions for super-polynomial sized revocation list in Section 3.12.

3.2 PRELIMINARIES

3.2.1 Symmetric Key Encryption

Definition 3.1 (Syntax). A symmetric key encryption scheme for message space $\mathcal{M} = \{\mathcal{M}_\lambda\}_{\lambda \in [\mathbb{N}]}$ and key space $\mathcal{K} = \{\mathcal{K}_\lambda\}_{\lambda \in [\mathbb{N}]}$ and ciphertext space $\mathcal{CT}_{\text{SKE}}$ has the following syntax:

$\text{Setup}(1^\lambda) \rightarrow \text{sk}$. The setup algorithm takes as input the security parameter λ and outputs a secret key sk .

$\text{Enc}(\text{sk}, m) \rightarrow \text{ct}$. The encryption algorithm takes as input the secret key sk and a message $m \in \mathcal{M}_\lambda$ and outputs a ciphertext ct .

$\text{Dec}(\text{sk}, \text{ct}) \rightarrow m'$. The decryption algorithm takes as input a secret key sk and a ciphertext ct and outputs a message $m' \in \mathcal{M}_\lambda$.

Correctness: A SKE scheme is said to be correct if there exists a negligible function $\text{negl}(\cdot)$ such that for all $\lambda \in \mathbb{N}$, for every message $m \in \mathcal{M}_\lambda$, we have

$$\Pr \left[\begin{array}{l} \text{sk} \leftarrow \text{Setup}(1^\lambda); \\ m' = m : \text{ct} \leftarrow \text{Enc}(\text{sk}, m); \\ m' = \text{Dec}(\text{sk}, \text{ct}). \end{array} \right] \geq 1 - \text{negl}(\lambda),$$

Security: A SKE scheme is said to have pseudorandom ciphertext if there exists a negligible function $\text{negl}(\cdot)$ such that for all $\lambda \in \mathbb{N}$, for every message $m \in \mathcal{M}_\lambda$, we have

$$\Pr \left[\begin{array}{l} \beta' = \beta : \\ \text{sk} \leftarrow \text{Setup}(1^\lambda); \\ \beta' \leftarrow \mathcal{A}^{\text{Enc}(\text{sk}, \cdot), \text{Enc}^\beta(\text{sk}, \cdot)}. \end{array} \right] \leq 1/2 + \text{negl}(\lambda),$$

where the $\text{Enc}(\text{sk}, \cdot)$ oracle, on input a message m , returns $\text{Enc}(\text{sk}, m)$ and $\text{Enc}(\text{sk}, \cdot)$ oracle, on input a message m , returns ct_β , where $\text{ct}_0 \leftarrow \text{Enc}(\text{sk}, m)$ and $\text{ct}_1 \leftarrow \mathcal{CT}_{\text{SKE}}$.

3.2.2 Functional Encryption

Here, we recall the definition of public-key and secret-key functional encryption.

Public-Key Functional Encryption

Consider a function family $\mathcal{F} = \{\mathcal{F}_\lambda\}_{\lambda \in \mathbb{N}}$, with input space $\mathcal{M} = \{\mathcal{M}_\lambda\}_{\lambda \in \mathbb{N}}$ and output space $\mathcal{Y} = \{\mathcal{Y}_\lambda\}_{\lambda \in \mathbb{N}}$, i.e, $\mathcal{F}_\lambda = \{f : \mathcal{M}_\lambda \rightarrow \mathcal{Y}_\lambda\}$. A public-key functional encryption scheme FE for \mathcal{F} consists of four polynomial time algorithms (Setup, KeyGen, Enc, Dec):

$\text{Setup}(1^\lambda) \rightarrow (\text{mpk}, \text{msk})$. The setup algorithm takes as input the security parameter λ and outputs a public key mpk and a master secret key msk .

$\text{KeyGen}(\text{msk}, f) \rightarrow \text{sk}_f$. The key generation algorithm takes as input the master secret key msk and a function $f \in \mathcal{F}_\lambda$ and it outputs a functional secret key sk_f .

$\text{Enc}(\text{mpk}, \mathbf{m}) \rightarrow \text{ct}$. The encryption algorithm takes as input the public key mpk and a message $\mathbf{m} \in \mathcal{M}_\lambda$ and outputs a ciphertext ct .

$\text{Dec}(\text{sk}_f, \text{ct}) \rightarrow y$. The decryption algorithm takes as input a functional secret key sk_f and a ciphertext ct and outputs $y \in \mathcal{Y}_\lambda$.

Definition 3.2 (Correctness). A functional encryption scheme is said to be correct if there exists a negligible function $\text{negl}(\cdot)$ such that for all $\lambda \in \mathbb{N}$, for every message $\mathbf{m} \in \mathcal{M}_\lambda$, and for every function $f \in \mathcal{F}_\lambda$, we have

$$\Pr \left[f(\mathbf{m}) \leftarrow \text{Dec}(\text{KeyGen}(\text{msk}, f), \text{Enc}(\text{mpk}, \mathbf{m})) \right] \geq 1 - \text{negl}(\lambda)$$

where $(\text{mpk}, \text{msk}) \leftarrow \text{Setup}(1^\lambda)$ and the probability is taken over the random coins of all algorithms.

Definition 3.3 (Selective Message Privacy). A functional encryption scheme over a function space $\mathcal{F} = \{\mathcal{F}_\lambda\}_{\lambda \in \mathbb{N}}$ is said to have selective message privacy (or simply is selectively secure) if for any PPT adversary \mathcal{A} , there exists a negligible function $\text{negl}(\cdot)$ such that for all $\lambda \in \mathbb{N}$, we have

$$\Pr \left[\begin{array}{l} (\mathbf{m}_0, \mathbf{m}_1) \leftarrow \mathcal{A}; \\ \mathcal{A}^{\text{KeyGen}(\text{msk}, \cdot)}(\text{mpk}, \text{ct}) = b : \quad (\text{mpk}, \text{msk}) \leftarrow \text{Setup}(1^\lambda); \\ \quad \quad \quad b \leftarrow \{0, 1\}; \text{ct} \leftarrow \text{Enc}(\text{mpk}, \mathbf{m}_b) \end{array} \right] \leq 1/2 + \text{negl}(\lambda),$$

where each key query for a function $f \in \mathcal{F}_\lambda$, queried by \mathcal{A} to the KeyGen oracle must satisfy the condition that $f(\mathbf{m}_0) = f(\mathbf{m}_1)$.

Compactness : We say that a FE scheme is compact if the running time of the encryption algorithm only depends on the security parameter and the input message length. In particular, it is independent of the complexity of the function class supported by the scheme.

Definition 3.4 (Fully Compact FE). A functional encryption scheme, $\text{FE} = (\text{FE.Setup}, \text{FE.KeyGen}, \text{FE.Enc}, \text{FE.Dec})$ for input space $\mathcal{M} = \{\mathcal{M}_\lambda\}_{\lambda \in \mathbb{N}}$ and

function class $\mathcal{F} = \{\mathcal{F}_\lambda\}_{\lambda \in \mathbb{N}}$, where each \mathcal{F}_λ is a finite collection of functions, is said to be fully compact if the running time of the encryption algorithm FE.Enc , on input FE.mpk and a message $\mathbf{m} \in \mathcal{M}_\lambda$, is $\text{poly}(\lambda, |\mathbf{m}|)$.

Jain, Lin and Sahai [129, 130] provided the first construction of FE with sublinear compactness, from standard assumptions. In more detail:

Lemma 3.1 ([130, 94, 30]). *There exists a public-key functional encryption scheme for polynomial sized circuits having selective security (as per Definition 3.3) and full compactness (as per Definition 3.4) with encryption time $\text{poly}(\lambda, |\mathbf{m}|)$, where $\lambda \in \mathbb{N}$ is the security parameter, \mathbf{m} is the input message, assuming LPN, DLIN and existence of boolean PRGs in NC^0 .*

Secret Key Functional Encryption

A secret key functional encryption scheme is the same as the public key functional encryption scheme, except that the setup algorithm only outputs msk and the encryption algorithm takes the master secret key msk as input, instead of the master public key mpk .

Definition 3.5 (Ada-IND Function-Message Privacy ([59], adapted)). A SKFE = (Setup, KeyGen, Enc, Dec) scheme over an input space $\mathcal{M} = \{\mathcal{M}_\lambda\}_{\lambda \in \mathbb{N}}$ and a function space $\mathcal{F} = \{\mathcal{F}_\lambda\}_{\lambda \in \mathbb{N}}$ is said to be Ada-IND function and message private if for any PPT algorithm \mathcal{A} , there exists a negligible function $\text{negl}(\cdot)$ such that

$$|\Pr[\text{Exp}_{\mathcal{A}}(\lambda, 0) = 1] - \Pr[\text{Exp}_{\mathcal{A}}(\lambda, 1) = 1]| \leq \text{negl}(\lambda)$$

where for each $b \in \{0, 1\}$ and $\lambda \in \mathbb{N}$, the $\text{Exp}_{\mathcal{A}}(\lambda, b)$ is defined as follows

1. $\text{msk} \leftarrow \text{Setup}(1^\lambda)$.
2. $b' \leftarrow \mathcal{A}^{\text{KeyGen}_b(\text{msk}, \cdot, \cdot), \text{Enc}_b(\text{msk}, \cdot, \cdot)}$, where the oracle $\text{KeyGen}_b(\text{msk}, \cdot, \cdot)$ on input (f_0, f_1) outputs $\text{KeyGen}(\text{msk}, f_b)$ and $\text{Enc}_b(\text{msk}, \cdot, \cdot)$ on input $(\mathbf{m}_0, \mathbf{m}_1)$ outputs $\text{Enc}(\text{msk}, \mathbf{m}_b)$.
3. Output b' .

and \mathcal{A} is admissible only if for all the key queries $(f_0, f_1) \in \mathcal{F} \times \mathcal{F}$ and encryption queries $(\mathbf{m}_0, \mathbf{m}_1) \in \mathcal{M} \times \mathcal{M}$, it must hold that $f_0(\mathbf{m}_0) = f_1(\mathbf{m}_1)$.

Note: We refer to the t -bounded Ada-IND function and message private SKFE scheme where Def. 3.5 holds only if at most t queries are made to the $\text{KeyGen}_b(\text{msk}, \cdot, \cdot)$ oracle.

Remark 2. We can construct a SKFE scheme satisfying t -bounded Ada-IND message privacy (without function privacy, i.e., $f_0 = f_1$ should hold in the security game) from one-way functions [31]. This implies a SKFE scheme satisfying t -bounded Ada-IND function-message privacy, using the conversion results from [59].

We also need a 2-bounded semi-adaptive simulation based function-message private SKFE scheme, defined next.

Definition 3.6 (Simulation Based Function-Message Privacy). A secret-key functional encryption scheme is said to satisfy t -bounded semi-adaptive simulation based function and message privacy, if for all PPT stateful algorithm \mathcal{A} , there exists PPT stateful algorithms $\text{Sim}^{\text{SK}}, \text{Sim}^{\text{CT}}$ such that:

$$\{\text{Exp}_{\mathcal{A}}^{\text{real}}(1^\lambda)\}_{\lambda \in \mathbb{N}} \approx_c \{\text{Exp}_{\mathcal{A}, \text{Sim}^{\text{SK}}, \text{Sim}^{\text{CT}}}^{\text{ideal}}(1^\lambda)\}_{\lambda \in \mathbb{N}}$$

where the real and ideal experiments of stateful algorithms $\mathcal{A}, \text{Sim}^{\text{SK}}, \text{Sim}^{\text{CT}}$ are as follows:

$\text{Exp}_{\mathcal{A}}^{\text{real}}(1^\lambda)$	$\text{Exp}_{\mathcal{A}, \text{Sim}}^{\text{ideal}}(1^\lambda)$
$\text{msk} \leftarrow \text{Setup}(1^\lambda)$	$\text{msk} \leftarrow \text{Setup}(1^\lambda)$
For $i \in [t]$:	For $i \in [t]$:
$\mathcal{A} \rightarrow f_i \in \mathcal{F};$	$\mathcal{A} \rightarrow f_i \in \mathcal{F};$
$\mathcal{A} \leftarrow \text{sk}_{f_i} = \text{KeyGen}(\text{msk}, f_i)$	$\mathcal{A} \leftarrow \text{sk}_{f_i} = \text{Sim}^{\text{SK}}(\text{msk}, 1^{ f_i })$
Repeat polynomially many times:	Repeat polynomially many times:
$\mathcal{A} \rightarrow \mathbf{m} \in \mathcal{M};$	$\mathcal{A} \rightarrow \mathbf{m} \in \mathcal{M};$
$\mathcal{A} \leftarrow \text{Enc}(\text{msk}, \mathbf{m})$	$\mathcal{A} \leftarrow \text{Sim}^{\text{CT}}(\text{msk}, \{f_i(\mathbf{m})\}_{i \in [t]})$
$\mathcal{A} \rightarrow b; \text{ Output } b$	$\mathcal{A} \rightarrow b; \text{ Output } b$

Lemma 3.2. *There exists a 2-bounded semi-adaptive simulation based function and message private SKFE scheme, assuming the existence of one-way functions.*

Proof. Let $\Pi = (\Pi.\text{Setup}, \Pi.\text{KeyGen}, \Pi.\text{Enc}, \Pi.\text{Dec})$ be a 2-bounded Ada-IND function-message private SKFE scheme. We construct a 2-bounded semi-adaptive simulation based function-message private SKFE $= (\text{Setup}, \text{KeyGen}, \text{Enc}, \text{Dec})$ scheme as follows:

1. $\text{Setup}(1^\lambda)$: Sample $\text{msk} \leftarrow \Pi.\text{Setup}(1^\lambda)$.
2. $\text{KeyGen}(\text{msk}, f)$: $\text{sk} \leftarrow \Pi.\text{KeyGen}(\text{msk}, E[f, 0])$. Here, $E[f, \text{mode}]$ is defined as $E[f, \text{mode}](\mathbf{m}, y_1, y_2) = y_{\text{mode}}$ for $\text{mode} \in \{0, 1, 2\}$, where $y_0 := f(\mathbf{m})$.
3. $\text{Enc}(\text{msk}, \mathbf{m})$: $\text{ct} \leftarrow \Pi.\text{Enc}(\text{msk}, (\mathbf{m}, \perp, \perp))$.
4. $\text{Dec}(\text{sk}, \text{ct})$: $\mathbf{m}' \leftarrow \Pi.\text{Dec}(\text{sk}, \text{ct})$.

Simulators for encryption and keygen are described as follows:

$\text{Sim}^{\text{CT}}(\text{msk}, \{f_1(\mathbf{m}), f_2(\mathbf{m})\})$ outputs $\Pi.\text{Enc}(\text{msk}, (\perp, f_1(\mathbf{m}), f_2(\mathbf{m})))$.

$\text{Sim}^{\text{SK}}(\text{msk}, 1^{|f|})$ outputs $\Pi.\text{KeyGen}(\text{msk}, E[1^{|f|}, b])$ for the b -th key query, for $b \in \{1, 2\}$.

Security:

Firstly, observe that the simulator thus constructed is valid since its output does not depend on the function f or input \mathbf{m} . Then proof of security follows directly from Ada-IND security of Π , because for any set of queries consisting of two functions f_1 and f_2 and polynomial many inputs $\{\mathbf{m}_i\}$, $f_b(\mathbf{m}_i) = g_b(\mathbf{z}_i)$ for $b \in \{1, 2\}$, where g_b is the circuit defined as $(E[1^{|f_b|}, b])$ and $\mathbf{z}_i = (\perp, f_1(\mathbf{m}_i), f_2(\mathbf{m}_i))$ used by the simulators to generate the simulated keys and ciphertexts, respectively. ■

3.3 ATTRIBUTE BASED ENCRYPTION

For the definition of key-policy (KPABE) and ciphertext-policy ABE, we refer to Section 2.2.

Boneh et al. [43] provided a construction of kpABE which satisfies key compactness and ciphertext succinctness. The following theorem summarizes the efficiency properties of their construction.

Theorem 3.3 (Properties of [43]). *There exists a key-policy ABE scheme $\text{kpABE} = (\text{kpABE.Setup}, \text{kpABE.KeyGen}, \text{kpABE.Enc}, \text{kpABE.Dec})$ for function class $\mathcal{C}_{\ell,d}$ which is selectively secure under the LWE assumption and has the following properties.*

In particular:

1. *Key Compactness.* We have $|\text{ABE.sk}_C| \leq \text{poly}(\lambda, d)$ for any $C \in \mathcal{C}_{\ell,d}$, where $(\text{ABE.mpk}, \text{ABE.msk}) \leftarrow \text{ABE.Setup}(1^\lambda)$ and $\text{ABE.sk}_C \leftarrow \text{ABE.KeyGen}(\text{ABE.msk}, C)$. In particular, the length of the secret key is independent of the attribute length ℓ and the size of the circuit C .
2. *Parameters Succinctness.* We have $|\text{ABE.mpk}|, |\text{ABE.msk}| \leq \text{poly}(\lambda, d, \ell)$ and $|\text{ABE.ct}| \leq \text{poly}(\lambda, d, \ell) + |\mu|$ for any $\mathbf{x} \in \mathcal{X}_\lambda$ and $\mu \in \mathcal{M}_\lambda$, where $(\text{ABE.mpk}, \text{ABE.msk}) \leftarrow \text{ABE.Setup}(1^\lambda)$ and $\text{ABE.ct} \leftarrow \text{ABE.Enc}(\text{ABE.mpk}, \mathbf{x}, \mu)$.
3. *Online-Offline Decryption.* The decryption algorithm $\text{Dec}(\text{mpk}, \text{sk}_C, C, \text{ct}_\mathbf{x}, \mathbf{x})$ can be divided into two parts, which we call
 - $\text{Dec}^{\text{off}}(\text{mpk}, C, \mathbf{x}) \rightarrow \text{off}$, which performs the heavier computation that involves the circuit C and attribute \mathbf{x} offline without knowing the ciphertext $\text{ct}_\mathbf{x}$ or the secret key sk_C to get a "help" off . We have that the size of off is $\text{poly}(\lambda, \ell, d)$ and the depth of the circuit $\text{Dec}^{\text{off}}(\cdot, \cdot, \cdot)$ is bounded by $\text{poly}(\lambda, \ell, \text{depth}(C))$, where $\text{depth}(C)$ is the depth of the circuit C .
 - $\text{Dec}^{\text{on}}(\text{sk}_C, \text{ct}_\mathbf{x}, \text{off})$ takes the help off generated offline along with the secret key sk_C and the ciphertext $\text{ct}_\mathbf{x}$ and outputs the underlying message μ . We note that this part does not take C as input and in particular, the size of the circuit $\text{Dec}^{\text{on}}(\cdot, \cdot, \cdot)$ is $\text{poly}(\lambda, \ell, d)$, which is independent from the size of the circuit C .

We will provide the detail of the kpABE scheme given by [43] in Sec. 3.3.1. There, we

will show that the scheme satisfies the online-offline decryption property defined above. We will also use the **cpABE** scheme given by [169]. The following theorem summarizes the properties of the scheme.

Theorem 3.4 (Properties of [169]). *There exists a ciphertext policy ABE scheme $\text{cpABE} = (\text{cpABE.Setup}, \text{cpABE.KeyGen}, \text{cpABE.Enc}, \text{cpABE.Dec})$ for function class $\mathcal{C}_{\ell,d}$, which is very selectively secure under the evasive and tensor LWE assumption and has the following properties. In particular:*

1. *Ciphertext Compactness.* We have $|\text{cpABE.ct}| \leq \text{poly}(\lambda, d) + |\mu|$ for any $C \in \mathcal{C}_{\ell,d}$ and $\mu \in \mathcal{M}_\lambda$, where $(\text{cpABE.mpk}, \text{cpABE.msk}) \leftarrow \text{cpABE.Setup}(1^\lambda)$ and $\text{cpABE.ct} \leftarrow \text{cpABE.Enc}(\text{cpABE.mpk}, C, \mu)$. In particular, the size of the ciphertext is independent from the size of the circuit C and its input length.
2. *Parameters Succinctness.* We have $|\text{cpABE.mpk}|, |\text{cpABE.msk}|, |\text{cpABE.sk}_x| \leq \text{poly}(\lambda, \ell, d)$ for any $x \in \{0, 1\}^\ell$, where $(\text{cpABE.mpk}, \text{cpABE.msk}) \leftarrow \text{cpABE.Setup}(1^\lambda)$ and $\text{cpABE.sk} \leftarrow \text{cpABE.KeyGen}(\text{cpABE.msk}, x)$.

3.3.1 Key-Policy ABE by Boneh et al. [43]

We will use the **kpABE** scheme proposed by Boneh et al. [43]. We provide the description of the scheme in the following while showing that the decryption algorithm can be divided into two phases as stated in Theorem 3.3. We note that the presentation here is largely based on [27].

Lattice Evaluation. The following is an abstraction of the evaluation procedure in previous LWE based FHE and ABE schemes. We follow the presentation by Tsabary [162].

Lemma 3.5 (Fully Homomorphic Computation [43]). *There exists a pair of deterministic algorithms $(\text{EvalF}, \text{EvalFX})$ with the following properties.*

- $\text{EvalF}(\mathbf{B}, F) \rightarrow \mathbf{H}_F$. Here, $\mathbf{B} \in \mathbb{Z}_q^{n \times m\ell}$ and $F : \{0, 1\}^\ell \rightarrow \{0, 1\}$ is a circuit.
- $\text{EvalFX}(F, \mathbf{x}, \mathbf{B}) \rightarrow \hat{\mathbf{H}}_{F, \mathbf{x}}$. Here, $\mathbf{x} \in \{0, 1\}^\ell$ and $F : \{0, 1\}^\ell \rightarrow \{0, 1\}$ is a circuit with depth d . We have $[\mathbf{B} - \mathbf{x} \otimes \mathbf{G}] \hat{\mathbf{H}}_{F, \mathbf{x}} = \mathbf{B} \mathbf{H}_F - F(\mathbf{x}) \mathbf{G} \bmod q$, where we

denote $[x_1 \mathbf{G} \parallel \dots \parallel x_k \mathbf{G}]$ by $\mathbf{x} \otimes \mathbf{G}$. Furthermore, we have $\|\mathbf{H}_F\|_\infty \leq m^{O(d)}$ and $\|\widehat{\mathbf{H}}_{F,\mathbf{x}}\|_\infty \leq m^{O(d)}$.

- The running time of $(\text{EvalF}, \text{EvalFX})$ is bounded by $\text{poly}(n, m, \log q, d)$.

Note that the last item implies that the circuits computing EvalF and EvalFX can be implemented with depth $\text{poly}(n, m, \log q, d)$.

Key-Policy ABE by Boneh et al. [43]

The scheme supports the circuit class $\mathcal{C}_{\ell(\lambda), d(\lambda)}$, which is the set of all circuits with input length $\ell(\lambda)$ and depth at most $d(\lambda)$ with arbitrary $\ell(\lambda) = \text{poly}(\lambda)$ and $d(\lambda) = \text{poly}(\lambda)$. In our case, we will set $d(\lambda) = \omega(\log \lambda)$. In the description below, we focus on the case where the message space is $\{0, 1\}$ for simplicity. To encrypt a long message, we run the construction in parallel to encrypt an SKE key $K \in \{0, 1\}^\lambda$ and then use it to encrypt the message.

Setup(1^λ): On input 1^λ , the setup algorithm defines the parameters $n = n(\lambda)$, $m = m(\lambda)$, noise distributions χ over \mathbb{Z} , $\tau_0 = \tau_0(\lambda)$, $\tau = \tau(\lambda)$, and $B = B(\lambda)$ as specified later. It then proceeds as follows.

1. Sample $(\mathbf{A}, \mathbf{A}_{\tau_0}^{-1}) \leftarrow \text{TrapGen}(1^n, 1^m, q)$ such that $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$.
2. Sample random matrix $\mathbf{B} = (\mathbf{B}_1, \dots, \mathbf{B}_\ell) \leftarrow (\mathbb{Z}_q^{n \times m})^\ell$ and a random vector $\mathbf{u} \leftarrow \mathbb{Z}_q^n$.
3. Output the master public key $\text{mpk} = (\mathbf{A}, \mathbf{B}, \mathbf{u})$ and the master secret key $\text{msk} = \mathbf{A}_{\tau_0}^{-1}$.

KeyGen(msk, C): The key generation algorithm takes as input the master public key mpk , the master secret key msk , and a circuit $C \in \mathcal{C}_{\ell, d}$ and proceeds as follows.

1. Set $F := \neg C$ to be the same circuit as C except that the output bit is flipped.⁷
2. Compute $\mathbf{H}_F = \text{EvalF}(\mathbf{B}, F)$ and $\mathbf{B}_F = \mathbf{B}\mathbf{H}_F$.

⁷While we follow the standard definition of ABE and require the decryption to be possible when $C(x) = 1$, it is when $C(\mathbf{x}) = 0$ in [43]. To fill the gap, we flip the output bit.

3. Compute $[\mathbf{A} \parallel \mathbf{B}_F]_{\tau}^{-1}$ from $\mathbf{A}_{\tau_0}^{-1}$ and sample $\mathbf{r} \in \mathbb{Z}^{2m}$ as $\mathbf{r}^{\top} \leftarrow [\mathbf{A} \parallel \mathbf{B}_F]_{\tau}^{-1}(\mathbf{u}^{\top})$.
4. Output the secret key $\mathbf{sk}_C := \mathbf{r}$.

$\text{Enc}(\text{mpk}, \mathbf{x}, \mu)$: The encryption algorithm takes as input the master public key mpk , an attribute $\mathbf{x} \in \{0, 1\}^{\ell}$, and a message $\mu \in \{0, 1\}$ and proceeds as follows.

1. Sample $\mathbf{s} \leftarrow \mathbb{Z}_q^n$, $e_0 \leftarrow \chi$, $\mathbf{e} \leftarrow \chi^m$, and $\mathbf{e}_{i,b} \leftarrow \widetilde{\chi}^m$ for $i \in [\ell]$ and $b \in \{0, 1\}$, where $\widetilde{\chi}^m$ is the distribution obtained by first sampling $\mathbf{x} \leftarrow \chi^m$ and $\mathbf{S} \leftarrow \{-1, 1\}^{m \times m}$ and then outputting $\mathbf{S}\mathbf{x}$.
2. Compute

$$\begin{aligned} \text{For all } i \in [\ell], b \in \{0, 1\}, \psi_{i,b} &:= \mathbf{s}(\mathbf{B}_i - b\mathbf{G}) + \mathbf{e}_{i,b} \in \mathbb{Z}_q^m \\ \psi_{2\ell+1} &:= \mathbf{s}\mathbf{A} + \mathbf{e} \in \mathbb{Z}_q^m, \psi_{2\ell+2} := \mathbf{s}\mathbf{u}^{\top} + e_0 + \mu \lceil q/2 \rceil \in \mathbb{Z}_q, \end{aligned}$$

3. Output the ciphertext $\text{ct}_{\mathbf{x}} := (\{\psi_{i,x_i}\}_{i \in [\ell]}, \psi_{2\ell+1}, \psi_{2\ell+2})$, where x_i is the i -th bit of \mathbf{x} .

$\text{Dec}(\text{mpk}, \mathbf{sk}_C, C, \text{ct}_{\mathbf{x}}, \mathbf{x})$: The decryption algorithm takes as input the master public key mpk , a secret key \mathbf{sk}_C for a circuit C , and a ciphertext $\text{ct}_{\mathbf{x}}$ for an attribute \mathbf{x} and proceeds as follows. The decryption algorithm can be divided into offline and online phase Dec^{off} and Dec^{on} , respectively, as defined below.

$\text{Dec}^{\text{off}}(\text{mpk}, C, \mathbf{x})$:

1. Compute $C(\mathbf{x})$ and $\widehat{\mathbf{H}}_{F,\mathbf{x}} = \text{EvalF}(F, \mathbf{x}, \mathbf{B})$, where $F = \neg C$.
2. Output $\text{off} = (C(\mathbf{x}), \widehat{\mathbf{H}}_{F,\mathbf{x}})$.

$\text{Dec}^{\text{on}}(\mathbf{sk}_C, \text{ct}_{\mathbf{x}}, \text{off})$:

1. Parse off as $(C(\mathbf{x}), \widehat{\mathbf{H}}_{F,\mathbf{x}})$ and $\text{ct}_{\mathbf{x}} \rightarrow (\{\psi_{i,x_i} \in \mathbb{Z}_q^m\}_{i \in [\ell]}, \psi_{2\ell+1} \in \mathbb{Z}_q^m, \psi_{2\ell+2} \in \mathbb{Z}_q)$, and $\mathbf{sk}_C = \mathbf{r} \in \mathbb{Z}^{2m}$. If any of the component is not in the corresponding domain or $C(\mathbf{x}) = 0$, output \perp .
2. Concatenate $\{\psi_{i,x_i}\}_{i \in [\ell]}$ to form $\psi_{\mathbf{x}} = (\psi_{1,x_1}, \dots, \psi_{\ell,x_{\ell}})$.

3. Compute

$$\psi' := \psi_{2\ell+2} - [\psi_{2\ell+1} \|\psi_{\mathbf{x}} \widehat{\mathbf{H}}_{F,\mathbf{x}}] \mathbf{r}^\top.$$

4. Output 0 if $\psi' \in [-B, B]$ and 1 if $[-B + \lceil q/2 \rceil, B + \lceil q/2 \rceil]$.

Parameters and Security. We choose the parameters for the scheme as follows for concreteness:

$$\begin{aligned} d &= \log n \log \log n, & n &= \tilde{\Theta}(\lambda^c), & m &= n^{1.1} \log q, & q &= n^{O(\log^3 n)}, \\ \chi &= \text{SampZ}(3\sqrt{n}), & \tau_0 &= n \log q \log m, & \tau &= n^{O(\log^2 n)} & B &= n^{O(\log^2 n)}, \end{aligned}$$

where c is some constant (e.g., $c = 1$).

It is easy to see that the efficiency requirement stated in Theorem 3.3 directly follows from the above parameter settings and Lemma 3.5.

Theorem 3.6 (Adapted from [43]). *The above scheme satisfies Sel-INDr security (Definition 2.6) if we assume the LWE assumption with $n^{O(\log^3 n)}$ approximation factor.*

3.3.2 Lockable Obfuscation

We define lockable obfuscation [111, 172] below. Consider a function family $\mathcal{F} = \{\mathcal{F}_\lambda\}_{\lambda \in \mathbb{N}}$, with input space $\mathcal{X} = \{\mathcal{X}_\lambda\}_{\lambda \in \mathbb{N}}$ and output space $\mathcal{Y} = \{\mathcal{Y}_\lambda\}_{\lambda \in \mathbb{N}}$, i.e., $\mathcal{F}_\lambda = \{f : \mathcal{X}_\lambda \rightarrow \mathcal{Y}_\lambda\}$. A lockable obfuscation scheme for \mathcal{F} consists of algorithms **Obf** and **Eval** with the following syntax:

Obf($1^\lambda, f, \alpha$) $\rightarrow \tilde{f}$. The obfuscation algorithm takes as input the security parameter λ , a function $f \in \mathcal{F}_\lambda$ and a lock value $\alpha \in \mathcal{Y}_\lambda$. It outputs an obfuscated program \tilde{f} .

Eval(\tilde{f}, x) $\rightarrow 1 \cup \{\perp\}$. The evaluation algorithm takes as input the obfuscated program \tilde{f} and an input $x \in \mathcal{X}_\lambda$. It outputs 1 or \perp .

Definition 3.7 (Correctness). A lockable obfuscation scheme is said to be correct if it satisfies the following properties:

1. For all $\lambda \in \mathbb{N}$, $f \in \mathcal{F}_\lambda$, $x \in \mathcal{X}_\lambda$ and $\alpha \in \mathcal{Y}_\lambda$ such that $f(x) = \alpha$, we have

$$\text{Eval}\left(\text{Obf}(1^\lambda, f, \alpha), x\right) = 1$$

2. There exists a negligible function $\text{negl}(\cdot)$ such that for all $\lambda \in \mathbb{N}$, $f \in \mathcal{F}_\lambda$, $x \in \mathcal{X}_\lambda$ and $\alpha \in \mathcal{Y}_\lambda$ such that $f(x) \neq \alpha$, we have

$$\Pr[\text{Eval}\left(\text{Obf}(1^\lambda, f, \alpha), x\right) = \perp] \geq 1 - \text{negl}(\lambda)$$

where the probability is taken over the random coins used during obfuscation.

Definition 3.8 (Security). A lockable obfuscation scheme is said to be secure if there is a PPT simulator Sim such that for all $f \in \mathcal{F}_\lambda$, we have

$$\text{Obf}(1^\lambda, f, \alpha) \approx_c \text{Sim}(1^\lambda, 1^{|f|})$$

where $\alpha \leftarrow \mathcal{Y}_\lambda$ and the probability is taken over the randomness of the obfuscator and simulator Sim .

Theorem 3.7 ([111, 172]). *There exists lockable obfuscation for all circuits with lock space $\{0, 1\}^\lambda$ from the LWE assumption.*

3.3.3 Laconic Oblivious Transfer

We define laconic oblivious transfer (LOT) [73] below. A LOT scheme consists of four algorithms crsGen , Hash , Send and Receive with the following syntax:

$\text{crsGen}(1^\lambda) \rightarrow \text{crs}$. It takes the security parameter λ as input and outputs a common reference string crs .

$\text{Hash}(\text{crs}, D) \rightarrow (\text{digest}, \hat{D})$. It takes as input a common reference string crs and a database $D \in \{0, 1\}^*$ and outputs a digest digest of the database and a state \hat{D} .

$\text{Send}(\text{crs}, \text{digest}, L, m_0, m_1) \rightarrow e$. It takes as input a common reference string crs , a digest digest , a database location $L \in \mathbb{N}$ and two messages m_0 and m_1 of length λ , and outputs a ciphertext e .

$\text{Receive}^{\hat{D}}(\text{crs}, e, L) \rightarrow m$. This is a RAM algorithm with random read access to \hat{D} .

It takes as input a common reference string crs , a ciphertext e , and a database location $L \in \mathbb{N}$. It outputs a message m .

Definition 3.9 (Correctness). A LOT scheme is said to be correct if for any database D of size at most $M = \text{poly}(\lambda)$ for any polynomial function $\text{poly}(\cdot)$, any memory location $L \in [M]$, and any pair of messages $(m_0, m_1) \in \{0, 1\}^\lambda \times \{0, 1\}^\lambda$, we have

$$\Pr \left[\begin{array}{l} \text{crs} \leftarrow \text{crsGen}(1^\lambda) \\ (\text{digest}, \hat{D}) \leftarrow \text{Hash}(\text{crs}, D) \\ e \leftarrow \text{Send}(\text{crs}, \text{digest}, L, m_0, m_1) \\ m \leftarrow \text{Receive}^{\hat{D}}(\text{crs}, e, L) \end{array} : m = m_{D[L]} \right] = 1,$$

where the probability is taken over the random choices made by crsGen and Send .

Definition 3.10 (Sender Privacy Against Semi-Honest Receivers). There exists a PPT simulator LOTSim such that the following holds. For any database D of size at most $M = \text{poly}(\lambda)$ for any polynomial function $\text{poly}(\cdot)$, any memory location $L \in [M]$, and any pair of messages $(m_0, m_1) \in \{0, 1\}^\lambda \times \{0, 1\}^\lambda$, let $\text{crs} \leftarrow \text{crsGen}(1^\lambda)$ and $\text{digest} \leftarrow \text{Hash}(\text{crs}, D)$. Then it holds that

$$(\text{crs}, \text{Send}(\text{crs}, \text{digest}, L, m_0, m_1)) \approx_c (\text{crs}, \text{LOTSim}(D, L, m_{D[L]})).$$

Theorem 3.8 ([73, 83, 58, 84]). *There exists laconic oblivious transfer where the length of digest is a fixed polynomial in security parameter λ , independent of the size of the database assuming either the Computational Diffie-Hellman assumption or the Factoring assumption or the Learning with Errors assumption. Moreover, the algorithm Hash runs in time $|D| \cdot \text{poly}(\log |D|, \lambda)$, Send and Receive run in time $\text{poly}(\log |D|, \lambda)$ for any database D of size at most $\text{poly}(\lambda)$ for any polynomial function $\text{poly}(\cdot)$.*

3.4 REVOCABLE PREDICATE ENCRYPTION

We define revocable predicate encryption (RPE), in both public and secret key setting. Since the two notions differ only in the encryption algorithm, we present them here in a unified way.

Definition 3.11. A RPE scheme for an attribute space $\mathcal{X} = \{\mathcal{X}_\lambda\}_{\lambda \in [\mathbb{N}]}$, a function family $\mathcal{F} = \{\mathcal{F}_\lambda\}_{\lambda \in [\mathbb{N}]}$ where $\mathcal{F}_\lambda = \{f : \mathcal{X}_\lambda \rightarrow \{0, 1\}\}$, a label space $\mathcal{L} = \{\mathcal{L}_\lambda\}_{\lambda \in [\mathbb{N}]}$ and a message space $\mathcal{M} = \{\mathcal{M}_\lambda\}_{\lambda \in [\mathbb{N}]}$ has the following probabilistic polynomial time algorithms:

$\text{Setup}(1^\lambda) \rightarrow (\text{mpk}, \text{msk})$. The setup algorithm takes the security parameter λ as input and it outputs a master public key mpk and a master secret key msk .

$\text{KeyGen}(\text{msk}, \text{lb}, x) \rightarrow \text{sk}_{\text{lb}, x}$ ⁸. The key generation algorithm takes as input the master secret key msk , a label $\text{lb} \in \mathcal{L}_\lambda$ and an attribute $x \in \mathcal{X}_\lambda$. It outputs a secret key $\text{sk}_{\text{lb}, x}$.

$\text{Enc}(\text{ek}, f, m, L) \rightarrow \text{ct}$. The encryption algorithm takes as input the encryption key ek , a function f , a message $m \in \mathcal{M}_\lambda$, and a revocation list $L \subseteq \mathcal{L}_\lambda$. It outputs a ciphertext ct .

$\text{Dec}(\text{sk}_{\text{lb}, x}, \text{ct}, L) \rightarrow m^{\text{Prime}}$. The decryption algorithm takes the secret key $\text{sk}_{\text{lb}, x}$, a ciphertext ct , and a revocation list L and it outputs $m^{\text{Prime}} \in \mathcal{M}_\lambda \cup \{\perp\}$.

In *public-key RPE*, we take $\text{ek} = \text{mpk}$ in the Enc algorithm, and in *secret-key RPE*, we take $\text{ek} = \text{msk}$. Furthermore, there is an additional algorithm in the secret key setting defined as follows:

$\text{Broadcast}(\text{mpk}, m, L) \rightarrow \text{ct}$. On input the master public key, a message m , and a revocation list $L \subseteq \mathcal{L}_\lambda$, the broadcast algorithm outputs a ciphertext ct .

⁸We want to point out that the secret key $\text{sk}_{\text{lb}, x}$ does not hide the corresponding label lb and attribute x and we assume these to be included in the secret key.

Definition 3.12 (Correctness). A revocable predicate encryption scheme is said to be correct if there exists a negligible function $\text{negl}(\cdot)$ such that for all $\lambda \in \mathbb{N}$, label $\text{lb} \in \mathcal{L}_\lambda$, attributes $x \in \mathcal{X}_\lambda$, predicates $f \in \mathcal{F}_\lambda$ such that $f(x) = 1$, all messages $m \in \mathcal{M}_\lambda$ and any set of revoked users $L \subseteq \mathcal{L}_\lambda$ such that $\text{lb} \notin L$, if we set $(\text{mpk}, \text{msk}) \leftarrow \text{Setup}(1^\lambda)$ and $\text{sk}_{\text{lb},x} \leftarrow \text{KeyGen}(\text{msk}, \text{lb}, x)$, then the following holds

$$\Pr [\text{Dec}(\text{sk}_{\text{lb},x}, \text{ct}, L) = m] \geq 1 - \text{negl}(\lambda),$$

for $\text{ct} \leftarrow \text{Enc}(\text{ek}, f, m, L)$ (*Encryption correctness*) and $\text{ct} \leftarrow \text{Broadcast}(\text{mpk}, m, L)$ (*Broadcast correctness*).

Security. In the following security definitions, we assume for simplicity that the adversary does not make key queries for same input (lb, x) more than once.

Definition 3.13 (q -query Message Hiding). Let $q(\cdot)$ be any fixed polynomial. A RPE scheme satisfies q -query message hiding property if for every PPT adversary \mathcal{A} , there exists a negligible function $\text{negl}(\cdot)$ such that for every $\lambda \in \mathbb{N}$, all messages $m \in \mathcal{M}_\lambda$ and any subset of revoked users $L \subseteq \mathcal{L}_\lambda$, the following holds

$$\Pr \left[\begin{array}{l} (\text{mpk}, \text{msk}) \leftarrow \text{Setup}(1^\lambda); \\ (f, m_0, m_1, L) \leftarrow \mathcal{A}^{\text{KeyGen}(\text{msk}, \cdot, \cdot), \text{Enc}(\text{ek}, \cdot, \cdot, \cdot)}(\text{mpk}); \\ \beta \leftarrow \{0, 1\}; \text{ct}_\beta \leftarrow \text{Enc}(\text{ek}, f, m_\beta, L); \\ \beta^{\text{Prime}} \leftarrow \mathcal{A}^{\text{KeyGen}(\text{msk}, \cdot, \cdot), \text{Enc}(\text{ek}, \cdot, \cdot, \cdot)}(\text{ct}_\beta) \end{array} \right] \leq \frac{1}{2} + \text{negl}(\lambda)$$

where \mathcal{A} can make at most $q(\lambda)$ queries to the encryption oracle $\text{Enc}(\text{ek}, \cdot, \cdot, \cdot)$, and \mathcal{A} is admissible if and only if for all the key queries (lb, x) to the $\text{KeyGen}(\text{msk}, \cdot, \cdot)$ oracle, either $f(x) = 0$ or $\text{lb} \in L$.

Definition 3.14 (q -query Selective Message Hiding). This is the same as the Def 3.13 except that \mathcal{A} outputs the revocation list L in the beginning of the game, before the Setup algorithm is run.

Definition 3.15 (q -query Very Selective Message Hiding). This is the same as the Def 3.14 except that \mathcal{A} outputs all the key queries (lb, x) to the $\text{KeyGen}(\text{msk}, \cdot, \cdot)$ oracle in the beginning of the game, before the Setup algorithm is run.

Definition 3.16 (q -query Function Hiding). Let $q(\cdot)$ be any fixed polynomial. A RPE scheme satisfies q -query function hiding property if for every PPT adversary \mathcal{A} , there exists a negligible function $\text{negl}(\cdot)$ such that for every $\lambda \in \mathbb{N}$, all messages $m \in \mathcal{M}_\lambda$ and any subset of revoked users $L \subseteq \mathcal{L}_\lambda$, the following holds

$$\Pr \left[\begin{array}{l} (\text{mpk}, \text{msk}) \leftarrow \text{Setup}(1^\lambda); \\ (f_0, f_1, m, L) \leftarrow \mathcal{A}^{\text{KeyGen}(\text{msk}, \cdot, \cdot), \text{Enc}(\text{ek}, \cdot, \cdot, \cdot)}(\text{mpk}); \\ \beta \leftarrow \{0, 1\}; \text{ct}_\beta \leftarrow \text{Enc}(\text{ek}, f_\beta, m, L); \\ \beta^{\text{Prime}} \leftarrow \mathcal{A}^{\text{KeyGen}(\text{msk}, \cdot, \cdot), \text{Enc}(\text{ek}, \cdot, \cdot, \cdot)}(\text{ct}_\beta) \end{array} \right] \leq \frac{1}{2} + \text{negl}(\lambda)$$

where \mathcal{A} can make at most $q(\lambda)$ queries to the encryption oracle $\text{Enc}(\text{ek}, \cdot, \cdot, \cdot)$, and \mathcal{A} is admissible if and only if for all the key queries (lb, x) to the $\text{KeyGen}(\text{msk}, \cdot, \cdot)$ oracle, either $f_0(x) = f_1(x)$ or $\text{lb} \in L$.

Definition 3.17 (q -query Selective Function Hiding). This is the same as the Def 3.16 except that \mathcal{A} outputs the revocation list L in the beginning of the game, before the Setup algorithm is run.

The following security notion is defined only for secret-key RPE scheme.

Definition 3.18 (q -query Selective Broadcast Security). Let $q(\cdot)$ be any fixed polynomial. A RPE scheme satisfies q -query selective broadcast security if there exists a negligible function $\text{negl}(\cdot)$ such that for every PPT adversary \mathcal{A} , for every $\lambda \in \mathbb{N}$, all messages $m \in \mathcal{M}_\lambda$ and any subset of revoked users $L \subseteq \mathcal{L}_\lambda$, the following holds

$$\Pr \left[\begin{array}{l} L \leftarrow \mathcal{A}(1^\lambda); \\ (\text{mpk}, \text{msk}) \leftarrow \text{Setup}(1^\lambda); \\ f, m \leftarrow \mathcal{A}^{\text{KeyGen}(\text{msk}, \cdot, \cdot), \text{Enc}(\text{msk}, \cdot, \cdot, \cdot)}(\text{mpk}); \\ \beta \leftarrow \{0, 1\}; \text{ct}_0 \leftarrow \text{Enc}(\text{msk}, f, m, L); \\ \text{ct}_1 \leftarrow \text{Broadcast}(\text{mpk}, m, L); \\ \beta^{\text{Prime}} \leftarrow \mathcal{A}^{\text{KeyGen}(\text{msk}, \cdot, \cdot), \text{Enc}(\text{msk}, \cdot, \cdot, \cdot)}(\text{ct}_\beta) \end{array} \right] \leq \frac{1}{2} + \text{negl}(\lambda)$$

where \mathcal{A} can make at most $q(\lambda)$ queries to the encryption $\text{Enc}(\text{msk}, \cdot, \cdot, \cdot)$ oracle and \mathcal{A} is admissible if and only if $f(x) = 1, \forall x \in \mathcal{X}_\lambda$.

Remark 3. In the *public-key* RPE scheme, the adversary \mathcal{A} can itself simulate the encryption oracle $\text{Enc}(\text{ek}, \cdot, \cdot, \cdot)$, as $\text{ek} = \text{mpk}$ in this setting. Therefore, in public-key setting, we refer to the security definitions without imposing the q -query bound on the encryption oracle.

Remark 4. We note that when the message space is binary, function space \mathcal{F}_λ is polynomially small and q is a constant, the weaker security definitions where adversary outputs the challenge function f , the challenge message m and the SK-Enc query functions $\{\tilde{f}_i\}_{i \in [q]}$ at the beginning of the game, before the $\text{Setup}(1^\lambda)$ algorithm is run, is equivalent to the definitions where the adversary outputs $f, m, \{\tilde{f}_i\}_{i \in [q]}$ adaptively. First, the functions can be guessed with polynomial loss. Furthermore, if we restrict the message space to be binary, we can guess the challenge message as well. To extend the message space, we can encrypt each bit by parallel systems.

3.5 PUBLIC-KEY RPE FROM FE AND LWE

3.5.1 Construction

In this section we provide our construction of a public key RPE scheme $\text{RPE} = (\text{RPE.Setup}, \text{RPE.KeyGen}, \text{RPE.Enc}, \text{RPE.Dec})$ for an attribute space $\mathcal{X} = \{\mathcal{X}_\lambda\}_\lambda$, a function family $\mathcal{F} = \{\mathcal{F}_\lambda\}_\lambda$ where $\mathcal{F}_\lambda = \{f : \mathcal{X}_\lambda \rightarrow \{0, 1\}\}$, a label space $\mathcal{L} = \{\mathcal{L}_\lambda\}_\lambda$ and a message space $\mathcal{M} = \{\mathcal{M}_\lambda\}_\lambda$ from polynomial hardness assumptions. We assume that $|\mathcal{F}_\lambda|$ and $|\mathcal{M}_\lambda|$ are bounded by some polynomial in λ . The restriction on $|\mathcal{F}_\lambda|$ is sufficient for our purpose and the restriction on $|\mathcal{M}_\lambda|$ can be removed by running the scheme in parallel.

Our construction uses the following building blocks:

1. A Sel-INDr secure key-policy ABE scheme $\text{kpABE} = (\text{kpABE.Setup}, \text{kpABE.Enc}, \text{kpABE.KeyGen}, \text{kpABE.Dec})$ for circuit class $\mathcal{C}_{\ell(\lambda), d(\lambda)}$ with parameter succinctness and key compactness (Theorem 3.3). Here $\ell(\lambda)$ is the input length and is the length of labels in our setting and the depth of the circuit is $d(\lambda) \in \omega(\log \lambda)$ to support unbounded revocation list. The message space of the scheme kpABE is $\mathcal{M} = \{\mathcal{M}_\lambda\}_\lambda$ and $\mathcal{CT}_{\text{kpABE}}$ denotes the ciphertext space. We assume that uniform sampling from $\mathcal{CT}_{\text{kpABE}}$ is efficiently possible without any

parameter.

2. A (fully) compact, selectively secure, public-key functional encryption scheme $\text{FE} = (\text{FE.Setup}, \text{FE.Enc}, \text{FE.KeyGen}, \text{FE.Dec})$ that supports polynomial sized circuits. We assume that the message space is sufficiently large so that it can encrypt an ABE master public key, a (description of) function $f \in \mathcal{F}_\lambda$, a PRF key, two secret keys of SKE, and a trit $\text{mode} \in \{0, 1, 2\}$.
3. A PRF $F : \{0, 1\}^\lambda \times \mathcal{X} \rightarrow \{0, 1\}^t$ where t is the length of the randomness used in kpABE encryption (Def. 2.10)
4. A symmetric key encryption schemes $\text{SKE} = (\text{SKE.KeyGen}, \text{SKE.Enc}, \text{SKE.Dec})$ with pseudorandom ciphertexts (Def. 3.1) We let $\mathcal{CT}_{\text{SKE}}$ denote the ciphertext space of SKE.⁹ We assume that uniform sampling from $\mathcal{CT}_{\text{SKE}}$ is efficiently possible without any parameter.

We describe our construction below.

$\text{RPE.Setup}(1^\lambda) \rightarrow (\text{RPE.mpk}, \text{RPE.msk})$. The setup algorithm does the following:

- Generate $(\text{FE.mpk}, \text{FE.msk}) \leftarrow \text{FE.Setup}(1^\lambda)$.
- Output $\text{RPE.mpk} = \text{FE.mpk}$ and $\text{RPE.msk} = \text{FE.msk}$.

$\text{RPE.KeyGen}(\text{RPE.msk}, \text{lb}, x) \rightarrow \text{RPE.sk}_{\text{lb},x}$. The key generation algorithm does the following:

- Sample random values $\gamma_1, \gamma_2, \delta \leftarrow \mathcal{CT}_{\text{SKE}}$.
- Construct a circuit $\text{Re-Enc}[\text{lb}, x, \gamma_1, \gamma_2, \delta]$ which has the label lb , attribute x , γ_1, γ_2 and δ hardwired, as defined in Figure 3.2.
- Compute $\text{FE.sk}_{\text{lb},x} \leftarrow \text{FE.KeyGen}(\text{FE.msk}, \text{Re-Enc}[\text{lb}, x, \gamma_1, \gamma_2, \delta])$.
- Output $\text{RPE.sk}_{\text{lb},x} = \text{FE.sk}_{\text{lb},x}$.

$\text{RPE.Enc}(\text{RPE.mpk}, f, m, L) \rightarrow \text{RPE.ct}$. The encryption algorithm does the following:

⁹ We note that we use the same ciphertext space for simplicity even though messages with different lengths are going to be encrypted. To have the same ciphertext space, we can, for example, pad short messages to be some fixed length, which is possible when the message length is bounded by some polynomial.

- Parse $\text{RPE.mpk} = \text{FE.mpk}$.
- Sample a F key $K \leftarrow \{0, 1\}^\lambda$.
- Generate $(\text{kpABE.mpk}, \text{kpABE.msk}) \leftarrow \text{kpABE.Setup}(1^\lambda)$.
- Compute $\text{FE.ct} \leftarrow \text{FE.Enc}(\text{FE.mpk}, (\text{kpABE.mpk}, f, m, K, 0, \perp, \perp))$.
- Construct a circuit C_L , with revocation list L hardwired defined as follows:
On input a label $\text{lb} \in \mathcal{L}_\lambda$,

$$C_L(\text{lb}) = 1 \text{ if and only if } \text{lb} \notin L. \quad (3.1)$$

Compute $\text{kpABE.sk}_L \leftarrow \text{kpABE.KeyGen}(\text{kpABE.msk}, C_L)$.

- Output $\text{RPE.ct} = (\text{kpABE.mpk}, \text{kpABE.sk}_L, \text{FE.ct})$.

$\text{RPE.Dec}(\text{RPE.sk}_{\text{lb},x}, \text{RPE.ct}, L) \rightarrow m'$. The decryption algorithm does the following:

- Parse $\text{RPE.ct} = (\text{kpABE.mpk}, \text{kpABE.sk}_L, \text{FE.ct})$ and $\text{RPE.sk}_{\text{lb},x} = \text{FE.sk}_{\text{lb},x}$.
- Compute $\text{ct}' = \text{FE.Dec}(\text{FE.sk}_{\text{lb},x}, \text{FE.ct})$.
- Construct circuit C_L from L and compute $m' = \text{kpABE.Dec}(\text{kpABE.mpk}, \text{kpABE.sk}_L, C_L, \text{ct}', \text{lb})$.
- Output m' .

Correctness. We now show that the above construction is correct via the following theorem.

Theorem 3.9. *Suppose FE and kpABE schemes are correct. Then the above construction satisfies the encryption correctness (Def. 3.12).*

Proof. Firstly, for any label lb , attribute x , and function f such that $f(x) = 1$, we have $\text{FE.Dec}(\text{FE.sk}_{\text{lb},x}, \text{FE.ct}) = \text{kpABE.ct}_{\text{lb}}$, where $\text{kpABE.ct}_{\text{lb}} = \text{kpABE.Enc}(\text{kpABE.mpk}, \text{lb}, m; F(K, (\text{lb}, x)))$, by the correctness of FE and the definition of Re-Enc. We then observe that $C_L(\cdot)$ can be implemented with

Function Re-Enc[lb, x , γ_1 , γ_2 , δ]

Hardwired values: A label lb, an attribute x , and SKE ciphertexts γ_1 , γ_2 , and δ .

Inputs: A kpABE master public key kpABE.mpk, a function $f \in \mathcal{F}_\lambda$, a message $m \in \mathcal{M}_\lambda$, a F key K , a trapdoor mode $\text{mode} \in \{0, 1, 2\}$ and ske keys ske.key₁ and ske.key₂.

Output : A kpABE ciphertext.

1. Parse the input as (ABE.mpk, f , m , K , mode, SKE.key₁, SKE.key₂).
2. Set $\tilde{m} = \begin{cases} m & \text{if } f(x) = 1 \\ 0 & \text{if } f(x) = 0. \end{cases}$
3. Compute kpABE.ct_{lb} = kpABE.Enc(kpABE.mpk, lb, \tilde{m} ; $F(K, (\text{lb}, x))$).
4. Compute flag = ske.Dec(ske.key₂, δ).
5. Compute out _{i} = ske.Dec(ske.key _{i} , γ_i) for $i \in \{1, 2\}$.
6. If mode = 0, output kpABE.ct_{lb}.
7. If mode = 1, output out₁.
8. If mode = 2, output $\begin{cases} \text{out}_2 & \text{if flag} = 1 \\ \text{kpABE.ct}_{\text{lb}} & \text{if flag} = 0. \end{cases}$

Figure 3.2: Function to compute kpABE ciphertexts.

depth $O(\log(|L| \cdot |\text{lb}|)) = O(\log \text{poly}(\lambda)) = O(\log \lambda) \leq d$ and thus $C_L(\cdot) \in \mathcal{C}_{\ell, d}$. Then if $\text{lb} \notin L$ we have $C_L(\text{lb}) = 1$ and hence from the correctness of the kpABE scheme it follows that $\text{kpABE.Dec}(\text{kpABE.mpk}, \text{kpABE.sk}_L, C_L, \text{kpABE.ct}_{\text{lb}}, \text{lb}) = m$. So the decryption correctly recovers the message when $f(x) = 1$ and $\text{lb} \notin L$. ■

Efficiency. Here we argue that our construction achieves optimal parameters. Namely, we show that the size of each parameter is independent from $|L|$. We note that $|f|$ refers to the description size of the function, not the circuit size that implements the function. When $|x|$ is very long and f has succinct description, the former can be much shorter

than the latter.

1. Public key size $|\text{RPE.mpk}|$: We have $|\text{RPE.mpk}| = |\text{FE.mpk}|$. Since we assumed that FE is fully compact (Def. 3.4), the length of FE.mpk only depends on the input length of Re-Enc. We have that the input length is $|\text{ABE.mpk}| + |f| + |m| + |K| + |\text{mode}| + 2|\text{SKE.key}| = |\text{ABE.mpk}| + |f| + O(\lambda)$. We have $|\text{ABE.mpk}| \leq \text{poly}(\lambda, |\mathbf{b}|, d) = \text{poly}(\lambda, |\mathbf{b}|)$ by Theorem 3.3. The total length is therefore $\text{poly}(\lambda, |f|, |\mathbf{b}|)$.
2. Secret key size $|\text{RPE.sk}_{|\mathbf{b},x}|$: We have $|\text{RPE.sk}_{|\mathbf{b},x}| = |\text{FE.sk}_{|\mathbf{b},x}|$. Since the size of the latter is polynomially dependent on the size of Re-Enc, we evaluate its size. We can see that the size of Re-Enc is polynomial in the total length of the input and the hardwired values. The length of the input is bounded by $\text{poly}(\lambda, |f|, |\mathbf{b}|)$ as analyzed in the above item. The length of the hardwired values are $|\mathbf{b}| + |x| + |\gamma_1| + |\gamma_2| + |\delta|$. We have $|\gamma_1| + |\gamma_2| + |\delta| = 3|\gamma_2|$ ¹⁰ and $|\gamma_2| = \text{poly}(\lambda, \text{kpABE.ct}_{|\mathbf{b}|}) = \text{poly}(\lambda, |\mathbf{b}|, d) = \text{poly}(\lambda, |\mathbf{b}|)$. Therefore, the size of Re-Enc is $\text{poly}(\lambda, |x|, |f|, |\mathbf{b}|)$ and so is the size of the secret key.
3. Ciphertext size $|\text{RPE.ct}|$: We have $|\text{RPE.ct}| = |\text{kpABE.mpk}| + |\text{kpABE.sk}_L| + |\text{FE.ct}|$. We have $|\text{ABE.mpk}| = \text{poly}(\lambda, |\mathbf{b}|)$ as we showed in the first item. We also have $|\text{kpABE.sk}_L| \leq \text{poly}(\lambda, d) \leq \text{poly}(\lambda)$ by the key compactness of kpABE (3.3). By similar analysis to the first item, full compactness of FE implies $|\text{FE.ct}| \leq \text{poly}(\lambda, |f|, |\mathbf{b}|)$. Therefore, the overall length of the ciphertext is $\text{poly}(\lambda, |f|, |\mathbf{b}|)$.

3.5.2 Security

Now we prove that the above construction of RPE satisfies both function hiding and message hiding security.

Function Hiding

Theorem 3.10. *Assume that F is a secure PRF, SKE is correct and secure, FE and kpABE are secure as per definitions 3.3 and 2.6, respectively. Furthermore, assume $|\mathcal{F}_\lambda| \leq \text{poly}(\lambda)$ and $|\mathcal{M}_\lambda| \leq \text{poly}(\lambda)$. Then the RPE constructed above is function hiding (Def. 3.16).*

Proof. Recall that for function hiding we need $\text{RPE.Enc}(\text{RPE.mpk}, f_0, m, L) \approx_c$

¹⁰We assumed that $|\gamma_1| = |\gamma_2| = |\delta|$. See Footnote 9.

$\text{RPE.Enc}(\text{RPE.mpk}, f_1, m, L)$, where for all the key queries (lb, x) , either $f_0(x) = f_1(x)$ or $\text{lb} \in L$.

The proof proceeds via a sequence of hybrid games between the challenger and a PPT adversary \mathcal{A} .

Hyb₀. This is the real world with $\beta = 0$, i.e. the challenge ciphertext is computed using the function f_0 . We write the complete game here to set up the notations and easy reference in later hybrids.

1. The adversary outputs the challenge functions f_0 and f_1 and the challenge message m ¹¹.
2. The challenger generates $(\text{FE.mpk}, \text{FE.msk}) \leftarrow \text{FE.Setup}(1^\lambda)$, sets $\text{RPE.mpk} = \text{FE.mpk}$ and sends it to the adversary. The challenger then responds to different queries from \mathcal{A} as follows:
 3. **Key Queries :** For each key query (lb, x) , the challenger does the following:
 - Samples random values $\gamma_1, \gamma_2, \delta \leftarrow \mathcal{CT}_{\text{SKE}}$.
 - Defines the circuit $\text{Re-Enc}[\text{lb}, x, \gamma_1, \gamma_2, \delta]$ as in Figure 3.2 and computes $\text{FE.sk}_{\text{lb},x} \leftarrow \text{FE.KeyGen}(\text{FE.msk}, \text{Re-Enc}[\text{lb}, x, \gamma_1, \gamma_2, \delta])$.
 - Returns $\text{RPE.sk}_{\text{lb},x} = \text{FE.sk}_{\text{lb},x}$ to the adversary.
 4. **Challenge Query:** When the adversary outputs the revocation list L for the challenge query, the challenger does the following:
 - Samples a F key K .
 - Generates $(\text{kpABE.mpk}, \text{kpABE.msk}) \leftarrow \text{kpABE.Setup}(1^\lambda)$.
 - Computes FE.ct as

$$\text{FE.Enc}(\text{FE.mpk}, (\text{kpABE.mpk}, f_0, m, K, 0, \perp, \perp)).$$
 - Defines C_L as in Eq. (3.1) and computes $\text{kpABE.sk}_L \leftarrow \text{kpABE.KeyGen}(\text{kpABE.msk}, C_L)$.

¹¹To keep the proofs and notations simple, we let f_0, f_1, m to be given selectively. This is sufficient to achieve security as in Def 3.16, as mentioned in Remark 4.

- Returns $\text{RPE.ct} = (\text{kpABE.mpk}, \text{kpABE.sk}_L, \text{FE.ct})$ to the adversary.

5. In the end, the adversary outputs a bit β' .

Hyb₁. This hybrid is same as the previous hybrid except the following changes:

- The challenger generates $(\text{kpABE.mpk}, \text{kpABE.msk}) \leftarrow \text{kpABE.Setup}(1^\lambda)$ in the beginning of the game after the adversary outputs the challenge functions f_0 and f_1 and the challenge message m . It also samples a SKE secret key SKE.key_1 and a F key K .
- For each key query (lb, x) , γ_1 is computed differently. In particular, the challenger does the following
 - Sets $\tilde{m} = \begin{cases} m & \text{if } f_0(x) = 1 \\ 0 & \text{if } f_0(x) = 0 \end{cases}$ and computes $\text{kpABE.ct}'_{\text{lb}} = \text{kpABE.Enc}(\text{kpABE.mpk}, \text{lb}, \tilde{m}; F(K, (\text{lb}, x)))$.
 - Sets γ_1 as $\text{SKE.Enc}(\text{SKE.key}_1, \text{kpABE.ct}'_{\text{lb}})$.

Hyb₂. This hybrid is same as the previous hybrid except the following changes:

- The challenger samples two SKE secret keys SKE.key_1 and SKE.key_2 in the beginning of the game (after receiving f_0, f_1, m from \mathcal{A}).
- For each key query (lb, x) , γ_2 and δ are computed differently from the previous hybrid. In particular, the challenger does the following:
 - Sets $\text{flag} = \begin{cases} 1 & \text{if } f_0(x) \neq f_1(x) \\ 0 & \text{otherwise.} \end{cases}$
 - Sets $\delta \leftarrow \text{SKE.Enc}(\text{SKE.key}_2, \text{flag})$ and $\gamma_2 \leftarrow \text{SKE.Enc}(\text{SKE.key}_2, \text{kpABE.ct}'_{\text{lb}})$ if $\text{flag} = 1$; else $\gamma_2 \leftarrow \mathcal{CT}_{\text{SKE}}$. We note that $\text{kpABE.ct}'_{\text{lb}}$ and γ_1 are computed as in the previous hybrid.

Hyb₃. This hybrid is same as the previous hybrid except that FE.ct in the challenge

ciphertext is computed as

$$\text{FE.Enc}(\text{FE.mpk}, (\text{kpABE.mpk}, f_0, m, \perp, 1, \text{SKE.key}_1, \text{SKE.key}_2)).$$

Hyb₄. This hybrid is same as the previous hybrid except that for each key query (lb, x) , $\text{kpABE.ct}'_{\text{lb}}$ is computed as follows:

- If $\text{flag} = 1$, $\text{kpABE.ct}'_{\text{lb}} = \text{kpABE.Enc}(\text{kpABE.mpk}, \text{lb}, \tilde{m}; r)$, where $r \leftarrow \{0, 1\}^t$.
- Else, if $\text{flag} = 0$,
 $\text{kpABE.ct}'_{\text{lb}} = \text{kpABE.Enc}(\text{kpABE.mpk}, \text{lb}, \tilde{m}; F(K, (\text{lb}, x)))$.

Hyb₅. This hybrid is same as the previous hybrid except that FE.ct in the challenge ciphertext is computed as

$$\text{FE.Enc}(\text{FE.mpk}, (\text{kpABE.mpk}, f_0, m, K, 2, \perp, \text{SKE.key}_2)).$$

Hyb₆. This hybrid is same as the previous hybrid except that for each key query (lb, x) , γ_1 is set differently as $\gamma_1 \leftarrow \mathcal{CT}_{\text{SKE}}$.

Hyb₇. This hybrid is same as the previous hybrid except that for each such key query (lb, x) where $\text{flag} = 1$, $\text{kpABE.ct}'_{\text{lb}}$ is sampled uniformly from $\mathcal{CT}_{\text{kpABE}}$, i.e.,
 $\text{kpABE.ct}'_{\text{lb}} \leftarrow \mathcal{CT}_{\text{kpABE}}$.

Hyb₈. This hybrid is same as the previous hybrid except that FE.ct in the challenge ciphertext is computed as

$$\text{FE.Enc}(\text{FE.mpk}, (\text{kpABE.mpk}, f_1, m, K, 2, \perp, \text{SKE.key}_2)).$$

We note that in the hybrids hereafter, we rewind the changes made in the preceding hybrids.

Hyb₉. This hybrid is same as the previous hybrid except that for each such key query (lb, x) where $\text{flag} = 1$, $\text{kpABE.ct}'_{\text{lb}}$ is changed back to $\text{kpABE.Enc}(\text{kpABE.mpk}, \text{lb}, \tilde{m}; r)$ for $r \leftarrow \{0, 1\}^t$, where \tilde{m} is now defined as m if $f_1(x) = 1$ and 0 otherwise.

Hyb₁₀. This hybrid is same as the previous hybrid except that for all the key queries (lb, x) , the challenger sets γ_1 as $\text{SKE.Enc}(\text{SKE.key}_1, \text{kpABE.ct}'_{\text{lb}})$.

Hyb₁₁. This hybrid is same as the previous hybrid except that FE.ct in the challenge ciphertext is computed as

$$\text{FE.Enc}(\text{FE.mpk}, (\text{kpABE.mpk}, f_1, m, \perp, 1, \text{SKE.key}_1, \text{SKE.key}_2)).$$

Hyb₁₂. This hybrid is same as the previous hybrid except that for each key query (lb, x) , $\text{kpABE.ct}'_{\text{lb}}$ is computed as $\text{kpABE.Enc}(\text{kpABE.mpk}, \text{lb}, \tilde{m}; F(K, (\text{lb}, x)))$.

Hyb₁₃. This hybrid is same as the previous hybrid except that FE.ct in the challenge ciphertext is computed as

$$\text{FE.Enc}(\text{FE.mpk}, (\text{kpABE.mpk}, f_1, m, K, 0, \perp, \perp)).$$

Hyb₁₄. This hybrid is same as the previous hybrid except that γ_2 and δ_2 are set as $\gamma_2, \delta \leftarrow \mathcal{CT}_{\text{SKE}}$.

Hyb₁₅. This hybrid is same as the previous hybrid except that γ_1 is set as $\gamma_1 \leftarrow \mathcal{CT}_{\text{SKE}}$. Note that this is the real world with $\beta = 1$, i.e., f_1 is encrypted in the challenge ciphertext.

Indistinguishability of hybrids We now show that the consecutive hybrids are indistinguishable.

Claim 3.11. Assume that SKE is secure, then $\text{Hyb}_0 \approx_c \text{Hyb}_1$.

Proof. We show that if \mathcal{A} can distinguish between Hyb_0 and Hyb_1 with non-negligible advantage ϵ , then there exists a PPT adversary \mathcal{B} against the security of SKE scheme with advantage ϵ . The reduction is as follows.

1. The SKE challenger samples $\text{SKE.key}_1 \leftarrow \text{SKE.Setup}(1^\lambda)$ and a bit $\hat{\beta} \leftarrow \{0, 1\}$ and starts the SKE security game with \mathcal{B} .
2. \mathcal{B} invokes \mathcal{A} , which then outputs the challenge functions f_0 and f_1 and the challenge message m .
3. \mathcal{B} generates $(\text{FE.mpk}, \text{FE.msk}) \leftarrow \text{FE.Setup}(1^\lambda)$, sets $\text{RPE.mpk} = \text{FE.mpk}$ and sends RPE.mpk to \mathcal{A} .
 \mathcal{B} also generates $(\text{kpABE.mpk}, \text{kpABE.msk}) \leftarrow \text{kpABE.Setup}(1^\lambda)$ and samples a F key K .

4. **Key Queries :** Whenever \mathcal{A} issues a key query (lb, x) , \mathcal{B} does the following:

- It sets $\tilde{m} = \begin{cases} m & \text{if } f_0(x) = 1 \\ 0 & \text{if } f_0(x) = 0 \end{cases}$
and $\text{kpABE.ct}'_{\text{lb}} = \text{kpABE.Enc}(\text{kpABE.mpk}, \text{lb}, \tilde{m}; F(K, (\text{lb}, x)))$. computes
- It sends $\text{kpABE.ct}'_{\text{lb}}$ as the challenge message to the SKE challenger. The SKE challenger returns $\text{ct}_{\hat{\beta}}$ to \mathcal{B} , where $\text{ct}_0 \leftarrow \mathcal{CT}_{\text{SKE}}$ and $\text{ct}_1 \leftarrow \text{SKE.Enc}(\text{SKE.key}_1, \text{kpABE.ct}'_{\text{lb}})$.
- It sets $\gamma_1 = \text{ct}_{\hat{\beta}}$, samples random values $\gamma_2, \delta \leftarrow \mathcal{CT}_{\text{SKE}}$ and computes the FE key for the circuit $\text{Re-Enc}[\text{lb}, x, \gamma_1, \gamma_2, \delta]$ and returns it as the secret key $\text{RPE.sk}_{\text{lb}, x}$ to \mathcal{A} .

5. **Challenge Query :** When \mathcal{A} outputs the revocation list L for the challenge ciphertext, \mathcal{B} does the following:

- Computes $\text{FE.ct} \leftarrow \text{FE.Enc}(\text{FE.mpk}, (\text{kpABE.mpk}, f_0, m, K, 0, \perp, \perp))$.
- Defines C_L as in Eq. 3.1 and computes $\text{kpABE.sk}_L \leftarrow \text{kpABE.KeyGen}(\text{kpABE.msk}, C_L)$.

- Returns $\text{RPE.ct} = (\text{kpABE.mpk}, \text{kpABE.sk}_L, \text{FE.ct})$ to \mathcal{A} .

6. In the end, \mathcal{A} outputs a bit β' . \mathcal{B} sends β' to the SKE challenger.

We observe that if the SKE challenger samples $\hat{\beta} = 0$, then \mathcal{B} simulated Hyb_0 , else Hyb_1 with \mathcal{A} . Hence, advantage of $\mathcal{B} = |\Pr(\beta' = 1 | \hat{\beta} = 0) - \Pr(\beta' = 1 | \hat{\beta} = 1)| = |\Pr(\beta' = 1 | \text{Hyb}_0) - \Pr(\beta' = 1 | \text{Hyb}_1)| = \epsilon$ (by assumption).

■

Claim 3.12. Assume that SKE is secure. Then $\text{Hyb}_1 \approx_c \text{Hyb}_2$.

Proof. The proof follows the same steps as that for the claim 3.11 and hence omitted. ■

Claim 3.13. Assume that FE satisfies selective security (Def. 3.3) and SKE is correct. Then $\text{Hyb}_2 \approx_c \text{Hyb}_3$.

Proof. We show that if \mathcal{A} can distinguish between the two hybrids with non-negligible advantage ϵ , then there exists a PPT adversary \mathcal{B} against the security of FE with the same advantage ϵ . \mathcal{B} is defined as follows:

1. Upon being invoked by the FE challenger, \mathcal{B} invokes \mathcal{A} . \mathcal{A} outputs the challenge functions f_0 and f_1 and the challenge message m .
2. \mathcal{B} samples two SKE secret keys $\text{SKE.key}_1, \text{SKE.key}_2$, a F key K and generates $(\text{kpABE.mpk}, \text{kpABE.msk}) \leftarrow \text{kpABE.Setup}(1^\lambda)$.
3. It sets $\mu_0 = (\text{kpABE.mpk}, f_0, m, K, 0, \perp, \perp)$ and $\mu_1 = (\text{kpABE.mpk}, f_0, m, \perp, 1, \text{SKE.key}_1, \text{SKE.key}_2)$ and sends (μ_0, μ_1) to the FE challenger as challenge messages.
4. The FE challenger generates $(\text{FE.mpk}, \text{FE.msk}) \leftarrow \text{FE.Setup}(1^\lambda)$, samples $\hat{\beta} \leftarrow \{0, 1\}$. It then computes $\text{FE.ct}_{\hat{\beta}} \leftarrow \text{FE.Enc}(\text{FE.mpk}, \mu_{\hat{\beta}})$ and sends $(\text{FE.mpk}, \text{FE.ct}_{\hat{\beta}})$ to \mathcal{B} .
5. \mathcal{B} sets $\text{RPE.mpk} = \text{FE.mpk}$ and sends RPE.mpk to \mathcal{A} .
6. **Key Queries :** When \mathcal{A} issues a key query (lb, x) , \mathcal{B} does the following:

- It sets $\tilde{m} = \begin{cases} m & \text{if } f_0(x) = 1 \\ 0 & \text{if } f_0(x) = 0 \end{cases}$ and computes $\text{kpABE.ct}'_{\text{lb}} = \text{kpABE.Enc}(\text{kpABE.mpk}, \text{lb}, \tilde{m}; F(K, (\text{lb}, x)))$.
- It sets $\text{flag} = \begin{cases} 1 & \text{if } f_0(x) \neq f_1(x) \\ 0 & \text{otherwise} \end{cases}$.
- It computes $\gamma_1 \leftarrow \text{SKE.Enc}(\text{SKE.key}_1, \text{kpABE.ct}'_{\text{lb}})$, $\delta \leftarrow \text{SKE.Enc}(\text{SKE.key}_2, \text{flag})$ and $\gamma_2 \leftarrow \text{SKE.Enc}(\text{SKE.key}_2, \text{kpABE.ct}'_{\text{lb}})$ if $\text{flag} = 1$; else samples $\gamma_2 \leftarrow \mathcal{CT}_{\text{SKE}}$.
- It defines $\text{Re-Enc}[\text{lb}, x, \gamma_1, \gamma_2, \delta]$ as in Figure 3.2 and sends a key query $\text{Re-Enc}[\text{lb}, x, \gamma_1, \gamma_2, \delta]$ to the FE challenger. The FE challenger returns the secret key $\text{FE.sk}_{\text{lb},x}$ to \mathcal{B} .
- \mathcal{B} returns $\text{RPE.sk}_{\text{lb},x} = \text{FE.sk}_{\text{lb},x}$ to \mathcal{A} .

7. **Challenge Query :** When \mathcal{A} outputs the revocation list L for the challenge ciphertext, \mathcal{B} does the following:

- It defines C_L as in Eq. (3.1) and computes $\text{kpABE.sk}_L \leftarrow \text{kpABE.KeyGen}(\text{kpABE.msk}, C_L)$.
- Returns $\text{RPE.ct} = (\text{kpABE.mpk}, \text{kpABE.sk}_L, \text{FE.ct}_{\hat{\beta}})$ to \mathcal{A} .

8. In the end, \mathcal{A} outputs a bit β' . \mathcal{B} sends β' to the FE challenger.

We observe that if FE challenger chose $\hat{\beta} = 0$, then \mathcal{B} simulated Hyb_2 , else Hyb_3 with \mathcal{A} . Hence, advantage of $\mathcal{B} = |\Pr(\beta' = 1 | \hat{\beta} = 0) - \Pr(\beta' = 1 | \hat{\beta} = 1)| = |\Pr(\beta' = 1 | \text{Hyb}_2) - \Pr(\beta' = 1 | \text{Hyb}_3)| = \epsilon$ (by assumption).

Admissibility of \mathcal{B}

Firstly, we observe that the only key queries that \mathcal{B} issues to the FE challenger are for the Re-Enc functions defined for each key query (lb, x) by \mathcal{A} . Next, we observe that for any function $\text{Re-Enc}[\text{lb}, x, \gamma_1, \gamma_2, \delta]$, $\mu_0 = (\text{kpABE.mpk}, f_0, m, K, 0, \perp, \perp)$, and $\mu_1 = (\text{kpABE.mpk}, f_0, m, \perp, 1, \text{SKE.key}_1, \text{SKE.key}_2)$ the following holds true from the definition of Re-Enc and correctness of SKE decryption,

$$\text{Re-Enc}[\text{lb}, x, \gamma_1, \gamma_2, \delta](\mu_0) = \text{kpABE.ct}_{\text{lb}}$$

$$= \text{kpABE.Enc}(\text{kpABE.mpk}, \text{lb}, \tilde{m}; F(K, (\text{lb}, x)))$$

$$\begin{aligned} \text{Re-Enc}[\text{lb}, x, \gamma_1, \gamma_2, \delta](\mu_1) &= \text{ske.Dec}(\text{ske.key}_1, \gamma_1) \\ &= \text{kpABE.ct}'_{\text{lb}} \\ &= \text{kpABE.Enc}(\text{kpABE.mpk}, \text{lb}, \tilde{m}; F(K, (\text{lb}, x))) \end{aligned}$$

Thus, for all the keys queried to the FE challenger, $\text{Re-Enc}[\text{lb}, x, \gamma_1, \gamma_2, \delta](\mu_0) = \text{Re-Enc}[\text{lb}, x, \gamma_1, \gamma_2, \delta](\mu_1)$. This establishes the admissibility of \mathcal{B} . \blacksquare

Claim 3.14. Assume that F is secure, then $\text{Hyb}_3 \approx_c \text{Hyb}_4$.

Proof. We show that if \mathcal{A} can distinguish between the two hybrids with non-negligible advantage ϵ , then there exists a PPT adversary \mathcal{B} against F security with the same advantage ϵ . The reduction \mathcal{B} is defined as follows:

1. The F challenger samples a F key K and a bit $\hat{\beta} \leftarrow \{0, 1\}$ and starts the game with \mathcal{B} .
2. \mathcal{B} then invokes \mathcal{A} which outputs the challenge functions f_0 and f_1 and the challenge message m .
3. \mathcal{B} samples two SKE secret keys $\text{SKE.key}_1, \text{SKE.key}_2$, generates $(\text{kpABE.mpk}, \text{kpABE.msk}) \leftarrow \text{kpABE.Setup}(1^\lambda)$, $(\text{FE.mpk}, \text{FE.msk}) \leftarrow \text{FE.Setup}(1^\lambda)$, sets $\text{RPE.mpk} = \text{FE.mpk}$ and sends RPE.mpk to \mathcal{A} .
4. **Key Queries :** When \mathcal{A} issues a key query (lb, x) , \mathcal{B} does the following:
 - If $\text{flag} = 0$, it sends an evaluation query for input (lb, x) to the F challenger and gets back $F(K, (\text{lb}, x))$. It then computes $\text{kpABE.ct}'_{\text{lb}} = \text{kpABE.Enc}(\text{kpABE.mpk}, \text{lb}, \tilde{m}; F(K, (\text{lb}, x)))$.
 - If $\text{flag} = 1$, it sends (lb, x) as a challenge query to the F challenger and gets back $r_{\hat{\beta}}$, where $r_0 = F(K, (\text{lb}, x))$ and $r_1 \leftarrow \{0, 1\}^t$. It then computes $\text{kpABE.ct}'_{\text{lb}} = \text{kpABE.Enc}(\text{kpABE.mpk}, \text{lb}, \tilde{m}; r_{\hat{\beta}})$.
 - Computes $\gamma_1 \leftarrow \text{SKE.Enc}(\text{SKE.key}_1, \text{kpABE.ct}'_{\text{lb}})$, $\delta \leftarrow \text{SKE.Enc}(\text{SKE.key}_2, \text{flag})$ and $\gamma_2 \leftarrow \text{SKE.Enc}(\text{SKE.key}_2, \text{kpABE.ct}'_{\text{lb}})$ if $\text{flag} = 1$; else $\gamma_2 \leftarrow C\mathcal{T}_{\text{SKE}}$.

- Defines $\text{Re-Enc}[\text{lb}, x, \gamma_1, \gamma_2, \delta]$ and computes the FE key for the circuit $\text{Re-Enc}[\text{lb}, x, \gamma_1, \gamma_2, \delta]$ and returns it as the secret key $\text{RPE.sk}_{\text{lb}, x}$ to \mathcal{A} .

5. **Challenge Query :** When \mathcal{A} outputs the revocation list L for the challenge ciphertext, \mathcal{B} does the following

- Computes $\text{FE.ct} \leftarrow \text{FE.Enc}(\text{FE.mpk}, (\text{kpABE.mpk}, f_0, m, \perp, 1, \text{SKE.key}_1, \text{SKE.key}_2))$.
- Defines C_L as in Eq. 3.1 and computes $\text{kpABE.sk}_L \leftarrow \text{kpABE.KeyGen}(\text{kpABE.msk}, C_L)$.
- Returns $\text{RPE.ct} = (\text{kpABE.mpk}, \text{kpABE.sk}_L, \text{FE.ct})$ to \mathcal{A} .

6. In the end, \mathcal{A} outputs a bit β' . \mathcal{B} sends β' to the SKE challenger.

We observe that if $\hat{\beta} = 0$, then \mathcal{B} simulated Hyb_3 , else Hyb_4 with \mathcal{A} . Hence, advantage of $\mathcal{B} = |\Pr(\beta' = 1 | \hat{\beta} = 0) - \Pr(\beta' = 1 | \hat{\beta} = 1)| = |\Pr(\beta' = 1 | \text{Hyb}_3) - \Pr(\beta' = 1 | \text{Hyb}_4)| = \epsilon$ (by assumption).

■

Claim 3.15. Assume that FE is selectively secure (as per Def 3.3) and SKE is correct. Then $\text{Hyb}_4 \approx_c \text{Hyb}_5$.

Proof. We show that if \mathcal{A} wins with non-negligible advantage ϵ in distinguishing the two hybrids, then there exists an adversary \mathcal{B} against FE security with the same advantage ϵ . The steps of the reduction are similar as in the proof of the Claim 3.13, with $\mu_0 = (\text{kpABE.mpk}, f_0, m, \perp, 1, \text{SKE.key}_1, \text{SKE.key}_2)$ and $\mu_1 = (\text{kpABE.mpk}, f_0, m, K, 2, \perp, \text{SKE.key}_2)$. Hence, here we only argue the admissibility of \mathcal{B} in the FE security game.

Admissibility of \mathcal{B}

Firstly, we observe that the only key queries that \mathcal{B} issues to the FE challenger are for

the $\text{Re-Enc}[\text{lb}, x, \gamma_1, \gamma_2, \delta]$ functions, where each function is defined corresponding to a key query (lb, x) by \mathcal{A} . Here, $\gamma_1 \leftarrow \text{SKE.Enc}(\text{SKE.key}_1, \text{kpABE.ct}'_{\text{lb}})$, $\delta \leftarrow \text{SKE.Enc}(\text{SKE.key}_2, \text{flag})$ and $\gamma_2 \leftarrow \text{SKE.Enc}(\text{SKE.key}_2, \text{kpABE.ct}'_{\text{lb}})$ if $\text{flag} = 1$; else $\gamma_2 \leftarrow \text{CT}_{\text{SKE}}$. Also $\text{kpABE.ct}'_{\text{lb}} = \text{kpABE.Enc}(\text{kpABE.mpk}, \text{lb}, \tilde{m}; F(K, (\text{lb}, x)))$ if $\text{flag} = 0$; else $\text{kpABE.ct}'_{\text{lb}} \leftarrow \text{kpABE.Enc}(\text{kpABE.mpk}, \text{lb}, \tilde{m}; r)$ where $r \leftarrow \{0, 1\}^t$.

Next, we observe that for any function $\text{Re-Enc}[\text{lb}, x, \gamma_1, \gamma_2, \delta]$, $\mu_0 = (\text{kpABE.mpk}, f_0, m, \perp, 1, \text{SKE.key}_1, \text{SKE.key}_2)$, and $\mu_1 = (\text{kpABE.mpk}, f_0, m, K, 2, \perp, \text{SKE.key}_2)$ the following holds true from the definition of Re-Enc and correctness of SKE decryption,

$$\begin{aligned}
\text{Re-Enc}[\text{lb}, x, \gamma_1, \gamma_2, \delta](\mu_0) &= \text{ske.Dec}(\text{ske.key}_1, \gamma_1) \\
&= \text{kpABE.ct}'_{\text{lb}} \\
&= \begin{cases} \text{kpABE.Enc}(\text{kpABE.mpk}, \text{lb}, \tilde{m}; F(K, (\text{lb}, x))) & \text{if flag} = 0 \\ \text{kpABE.Enc}(\text{kpABE.mpk}, \text{lb}, \tilde{m}; r) & \text{if flag} = 1 \end{cases} \\
\\
\text{Re-Enc}[\text{lb}, x, \gamma_1, \gamma_2, \delta](\mu_1) &= \begin{cases} \text{kpABE.ct}_{\text{lb}} & \text{if flag} = 0 \\ \text{SKE.Dec}(\text{SKE.key}_2, \gamma_2) = \text{kpABE.ct}'_{\text{lb}} & \text{if flag} = 1 \end{cases} \\
&= \begin{cases} \text{kpABE.Enc}(\text{kpABE.mpk}, \text{lb}, \tilde{m}; F(K, (\text{lb}, x))) & \text{if flag} = 0 \\ \text{kpABE.Enc}(\text{kpABE.mpk}, \text{lb}, \tilde{m}; r) & \text{if flag} = 1 \end{cases}
\end{aligned}$$

Thus, $\text{Re-Enc}[\text{lb}, x, \gamma_1, \gamma_2, \delta](\mu_0) = \text{Re-Enc}[\text{lb}, x, \gamma_1, \gamma_2, \delta](\mu_1)$. This establishes the admissibility of \mathcal{B} . ■

Claim 3.16. Assume that SKE is secure, then $\text{Hyb}_5 \approx_c \text{Hyb}_6$.

Proof. The proof is similar to the proof for the claim 3.11 and hence omitted. ■

Claim 3.17. Assume that kpABE is Sel-INDr secure (Def. 2.6), then $\text{Hyb}_6 \approx_c \text{Hyb}_7$.

Proof. To argue indistinguishability between the two hybrids, we define an intermediate hybrid Hyb_{6a} as follows - this hybrid is same as Hyb_6 except that for each such pre-challenge key query (lb, x) where $\text{flag} = 1$, $\text{kpABE.ct}'_{\text{lb}} \leftarrow \mathcal{CT}_{\text{kpABE}}$. The proof then follows from the following two claims:

Claim 3.18 (1). Assuming kpABE is Sel-INDr secure (Def. 2.6), $\text{Hyb}_6 \approx_c \text{Hyb}_{6a}$.

Proof. Let Q_{fpre} be the number of pre-challenge key queries with $\text{flag} = 1$ ¹². Then, we further define the following sub hybrids: for $i = 0$ to Q_{fpre} , define $\text{Hyb}_{6,i}$ which is same as Hyb_6 , except that for the first i key queries with $\text{flag} = 1$, $\text{kpABE.ct}'_{\text{lb}} \leftarrow \mathcal{CT}_{\text{kpABE}}$. Thus, $\text{Hyb}_{6,0} = \text{Hyb}_6$ and $\text{Hyb}_{6,Q_{\text{fpre}}} = \text{Hyb}_{6a}$. Next, we show that for all $i \in [Q_{\text{fpre}}]$, $\text{Hyb}_{6,i-1} \approx_c \text{Hyb}_{6,i}$. In particular, we show that if \mathcal{A} distinguishes between the two hybrids with non-negligible advantage ϵ , then there exists a PPT algorithm \mathcal{B} against Sel-INDr security of kpABE with the same advantage ϵ .

Observe that the two hybrids differ only in the value of $\text{kpABE.ct}'_{\text{lb}}$ used in the computation of $\text{RPE.sk}_{\text{lb},x}$ for the i -th key query (lb, x) with $\text{flag} = 1$; in the former hybrid we have $\text{kpABE.ct}'_{\text{lb}} = \text{kpABE.Enc}(\text{kpABE.mpk}, \text{lb}, \tilde{m}; r)$, while in the latter hybrid, $\text{kpABE.ct}'_{\text{lb}} \leftarrow \mathcal{CT}_{\text{kpABE}}$. Now, we define the reduction \mathcal{B} .

1. \mathcal{B} firstly invokes \mathcal{A} and gets f_0, f_1 and m .
2. \mathcal{B} then samples two SKE secret keys $\text{SKE.key}_1, \text{SKE.key}_2$, a F key K and generates $(\text{FE.mpk}, \text{FE.msk}) \leftarrow \text{FE.Setup}(1^\lambda)$. It sets $\text{RPE.mpk} = \text{FE.mpk}$ and sends it to \mathcal{A} . It then answers different queries from \mathcal{A} as follows:
3. **Key Queries:** For each key query (lb, x) , \mathcal{B} does the following:
 - Sets $\text{flag} = 0$, if $f_0(x) = f_1(x)$; else $\text{flag} = 1$. It computes $\delta = \text{SKE.Enc}(\text{SKE.key}_2, \text{flag})$ and samples $\gamma_1 \leftarrow \mathcal{CT}_{\text{SKE}}$.
 - Computes γ_2 as follows:

¹²Note that we can upper bound Q_{fpre} as $Q_{\text{fpre}} \leq |L|$.

- If $\text{flag} = 0$, $\gamma_2 \leftarrow \mathcal{CT}_{\text{SKE}}$.
 - Else, if $\text{flag} = 1$, then let this be the j -th key query with $\text{flag} = 1$. Then,
 - * For $j < i$, $\gamma_2 \leftarrow \text{SKE.Enc}(\text{SKE.key}_2, \text{kpABE.ct}'_{\text{lb}})$, where $\text{kpABE.ct}'_{\text{lb}} \leftarrow \mathcal{CT}_{\text{kpABE}}$. (Note that this does not require kpABE.mpk).
 - * For $j = i$, \mathcal{B} does the following:
 - Sends lb and \tilde{m} as the challenge attribute and message, respectively, to the kpABE challenger.
 - The kpABE challenger samples $\hat{\beta} \leftarrow \{0, 1\}$ and computes $\text{kpABE.ct} \leftarrow \text{kpABE.Enc}(\text{kpABE.mpk}, \text{lb}, \tilde{m}; r)$, if $\hat{\beta} = 0$, else samples $\text{kpABE.ct} \leftarrow \mathcal{CT}_{\text{kpABE}}$. The kpABE challenger sends $\{\text{kpABE.mpk}, \text{kpABE.ct}\}$ to \mathcal{B} .
 - \mathcal{B} then computes $\gamma_2 \leftarrow \text{SKE.Enc}(\text{SKE.key}_2, \text{kpABE.ct})$.
 - * For $j > i$, $\gamma_2 \leftarrow \text{SKE.Enc}(\text{SKE.key}_2, \text{kpABE.ct}'_{\text{lb}})$, where $\text{kpABE.ct}'_{\text{lb}} \leftarrow \text{kpABE.Enc}(\text{kpABE.mpk}, \text{lb}, \tilde{m}; r)$.
 - Defines $\text{Re-Enc}[\text{lb}, x, \gamma_1, \gamma_2, \delta]$, computes an FE key for this circuit and returns it as the secret key $\text{RPE.sk}_{\text{lb}, x}$ to \mathcal{A} .
4. **Challenge Query:** When \mathcal{A} outputs the revocation list L for the challenge query, \mathcal{B} does the following:
- Computes $\text{FE.ct} \leftarrow \text{FE.Enc}(\text{FE.mpk}, (\text{kpABE.mpk}, f_0, m, K, 2, \perp, \text{SKE.key}_2))$.
 - Defines C_L as in Eq. 3.1 and sends a key query for the circuit C_L to the kpABE challenger. The kpABE challenger returns kpABE.sk_L .
 - Returns $\text{RPE.ct} = (\text{kpABE.mpk}, \text{kpABE.sk}_L, \text{FE.ct})$ to \mathcal{A} .
5. In the end, \mathcal{A} outputs its guess bit β' . \mathcal{B} sends β' to the kpABE challenger as its guess bit.

We observe that if the kpABE challenger chose $\hat{\beta} = 0$, then \mathcal{B} simulated $\text{Hyb}_{6,i-1}$, else $\text{Hyb}_{6,i}$ with \mathcal{A} . Hence, advantage of $\mathcal{B} = |\Pr(\beta' = 1 | \hat{\beta} = 0) - \Pr(\beta' = 1 | \hat{\beta} = 1)| = |\Pr(\beta' = 1 | \text{Hyb}_{6,i-1}) - \Pr(\beta' = 1 | \text{Hyb}_{6,i})| = \epsilon$ (by assumption).

Admissibility of \mathcal{B} : Observe that \mathcal{B} issues a single kpABE key query which is for circuit C_L and the challenge attribute is the label lb corresponding to a key query (lb, x) by \mathcal{A} for which $\text{flag} = 1$. Hence, by the admissibility condition of \mathcal{A} , $\text{lb} \in L$ and hence $C_L(\text{lb}) = 0$. \blacksquare

Claim 3.19 (2). Assume kpABE is Sel-INDr secure (Def. 2.6). Then, $\text{Hyb}_{6a} \approx_c \text{Hyb}_7$.

Proof. To prove the claim, we again consider sub hybrids, $\text{Hyb}_{6a,i}$ for $i = 0$ to $|L|$, defined as follows: let $L_{[1:j]}$ be the set of first j labels in the revocation list L . Then, $\text{Hyb}_{6a,i}$ is same as Hyb_{6a} except the following changes: for any post-challenge key query (lb, x) such that $\text{lb} \in L_{[1:i]}$ and $\text{flag} = 1$, $\text{kpABE.ct}'_{\text{lb}} \leftarrow C\mathcal{T}_{\text{kpABE}}$. Thus, $\text{Hyb}_{6a,0} = \text{Hyb}_{6a}$ and $\text{Hyb}_{6a,|L|} = \text{Hyb}_7$. Next we prove the following claim:

Claim 3.20. Assume kpABE is Sel-INDr secure (Def. 2.6). Then for $i \in [|L|]$, $\text{Hyb}_{6a,i-1} \approx_c \text{Hyb}_{6a,i}$.

Proof. Let lb_i be the i -th label in L . Then we observe that if there is no post-challenge key query (lb, x) such that $\text{flag} = 1$ and $\text{lb} = \text{lb}_i$ then the two hybrids are identical. Else, we show that if \mathcal{A} can distinguish between the two hybrids with non negligible advantage ϵ then there exists a PPT algorithm \mathcal{B} against Sel-INDr security of kpABE security with the same advantage ϵ . \mathcal{B} is defined as follows:

1. \mathcal{B} firstly invokes \mathcal{A} and gets f_0, f_1 and m .
2. It then samples a SKE secret key SKE.key_2 , a F key K and generates FE keys as $(\text{FE.mpk}, \text{FE.msk}) \leftarrow \text{FE.Setup}(1^\lambda)$. It sets $\text{RPE.mpk} = \text{FE.mpk}$ and sends it to \mathcal{A} . It then answers different queries from \mathcal{A} as follows:
 3. **Pre-challenge Key Queries:** For pre-challenge key query (lb, x) , \mathcal{B} does the following:
 - Sets $\text{flag} = 0$ if $f_0(x) = f_1(x)$; else $\text{flag} = 1$ and computes $\delta = \text{SKE.Enc}(\text{SKE.key}_2, \text{flag})$. It also samples $\gamma_1 \leftarrow C\mathcal{T}_{\text{SKE}}$.
 - Computes γ_2 as follows: if $\text{flag} = 0$, $\gamma_2 \leftarrow C\mathcal{T}_{\text{SKE}}$; else $\gamma_2 \leftarrow \text{SKE.Enc}(\text{SKE.key}_2, \text{kpABE.ct}'_{\text{lb}})$, where $\text{kpABE.ct}'_{\text{lb}} \leftarrow C\mathcal{T}_{\text{kpABE}}$.

(Note that this does not require kpABE.mpk).

- Defines $\text{Re-Enc}[\text{lb}, x, \gamma_1, \gamma_2, \delta]$, computes an FE key for this circuit and returns it as the secret key $\text{RPE.sk}_{\text{lb},x}$ to \mathcal{A} .

4. **Challenge Query:** When \mathcal{A} outputs the revocation list $L = \{\text{lb}_1, \dots, \text{lb}_{|L|}\}$ for the challenge ciphertext, \mathcal{B} does the following:

- Sends lb_i to the kpABE challenger as the challenge attribute. The kpABE challenger samples $(\text{kpABE.mpk}, \text{kpABE.msk}) \leftarrow \text{kpABE.Setup}(1^\lambda)$ and $\hat{\beta} \leftarrow \{0, 1\}$, and sends kpABE.mpk to \mathcal{B} .
- Constructs circuit C_L as defined in the construction and sends a key query for C_L to the kpABE challenger and gets kpABE.sk_L in response.
- Computes $\text{FE.ct} \leftarrow \text{FE.Enc}(\text{FE.mpk}, (\text{kpABE.mpk}, f_0, m, K, 2, \perp, \text{SKE.key}_2))$
- Returns $\text{RPE.ct} = (\text{kpABE.mpk}, \text{kpABE.sk}_L, \text{FE.ct})$ to \mathcal{A} .

5. **Post-challenge Key Queries:** For each post-challenge key query (lb, x) , \mathcal{B} computes flag , δ and γ_1 as defined for the hybrid and defines $\text{Re-Enc}[\text{lb}, x, \gamma_1, \gamma_2, \delta]$, where γ_2 is computed as follows:

- if $\text{flag} = 0$, $\gamma_2 \leftarrow \mathcal{CT}_{\text{SKE}}$. Else,
 - if $\text{lb} \in L_{[1:i-1]}$, $\gamma_2 = \text{SKE.Enc}(\text{SKE.key}_2, \text{kpABE.ct}'_{\text{lb}})$, where $\text{kpABE.ct}'_{\text{lb}} \leftarrow \mathcal{CT}_{\text{kpABE}}$.
 - if $\text{lb} = \text{lb}_i$, \mathcal{B} sends challenge query with message \tilde{m} to the kpABE challenger. The kpABE challenger returns a ciphertext $\text{kpABE.ct} = \text{kpABE.Enc}(\text{kpABE.mpk}, \text{lb}_i, \tilde{m}; r)$, if $\hat{\beta} = 0$; else $\text{kpABE.ct} \leftarrow \mathcal{CT}_{\text{kpABE}}$. Then \mathcal{B} computes $\gamma_2 \leftarrow \text{SKE.Enc}(\text{SKE.key}_2, \text{kpABE.ct})$.¹³
 - if $\text{lb} \notin L_{[1:i]}$, then $\gamma_2 \leftarrow \text{SKE.Enc}(\text{SKE.key}_2, \text{kpABE.ct}'_{\text{lb}})$, where $\text{kpABE.ct}'_{\text{lb}} = \text{kpABE.Enc}(\text{kpABE.mpk}, \text{lb}, \tilde{m}; r)$.

\mathcal{B} computes an FE key for $\text{Re-Enc}[\text{lb}, x, \gamma_1, \gamma_2, \delta]$ and returns it as the secret key $\text{RPE.sk}_{\text{lb},x}$ to \mathcal{A} .

6. In the end, \mathcal{A} outputs its guess bit β' . \mathcal{B} sends β' to the kpABE challenger as its guess bit.

We observe that if the kpABE challenger chose $\hat{\beta} = 0$, then \mathcal{B} simulated $\text{Hyb}_{6a,i-1}$, else

¹³In case multiple queries with $\text{lb} = \text{lb}_i$ are made, we need to simulate the ciphertext multiple times. In that case, we rely on multi-challenge version of Sel-INDr , which is easily seen to be equivalent with the single challenge version.

$\text{Hyb}_{6a,i}$ with \mathcal{A} . Hence, advantage of $\mathcal{B} = |\Pr(\beta' = 1 | \hat{\beta} = 0) - \Pr(\beta' = 1 | \hat{\beta} = 1)| = |\Pr(\beta' = 1 | \text{Hyb}_{6a,i-1}) - \Pr(\beta' = 1 | \text{Hyb}_{6a,i})| = \epsilon$ (by assumption).

Admissibility of \mathcal{B} : Observe that \mathcal{B} issues a single ABE key query, which is for circuit C_L and the challenge attribute is $\text{lb}_i \in L$. Hence, by the design of C_L , $C_L(\text{lb}_i) = 0$ as desired. ■

■

■

Claim 3.21. Assume that FE is secure, then $\text{Hyb}_7 \approx_c \text{Hyb}_8$.

Proof. We show that if \mathcal{A} wins with non-negligible advantage ϵ in distinguishing the two hybrids, then there exists an adversary \mathcal{B} against FE security with the same advantage ϵ . The steps of the reduction are similar to the proof of Claim 3.13, with $\mu_0 = (\text{kpABE.mpk}, f_0, m, K, 2, \perp, \text{SKE.key}_2)$ and $\mu_1 = (\text{kpABE.mpk}, f_1, m, K, 2, \perp, \text{SKE.key}_2)$. Hence, here we only argue the admissibility of \mathcal{B} in the FE security game.

Admissibility of \mathcal{B} :

Firstly, we observe that the only key queries that \mathcal{B} issues to the FE challenger are for the $\text{Re-Enc}[\text{lb}, x, \gamma_1, \gamma_2, \delta]$ functions, where each function is defined corresponding to a key query (lb, x) by \mathcal{A} . Here, $\gamma_1 \leftarrow \mathcal{CT}_{\text{SKE}}$, $\delta \leftarrow \text{SKE.Enc}(\text{SKE.key}_2, \text{flag})$ and $\gamma_2 \leftarrow \text{SKE.Enc}(\text{SKE.key}_2, \text{kpABE.ct}'_{\text{lb}})$ if $\text{flag} = 1$; else $\gamma_2 \leftarrow \mathcal{CT}_{\text{SKE}}$, where $\text{kpABE.ct}'_{\text{lb}} \leftarrow \mathcal{CT}_{\text{kpABE}}$ for $\text{flag} = 1$. Next, we observe that for any $\text{Re-Enc}[\text{lb}, x, \gamma_1, \gamma_2, \delta]$ function, $\mu_0 = (\text{kpABE.mpk}, f_0, m, K, 2, \perp, \text{SKE.key}_2)$, and

$\mu_1 = (\text{kpABE.mpk}, f_1, m, K, 2, \perp, \text{SKE.key}_2)$ the following holds true from the definition of Re-Enc and correctness of SKE decryption,

- when $\text{flag} = 0$, i.e. $f_0(x) = f_1(x)$, \tilde{m} in $\text{kpABE.ct}_{\text{lb}} = \text{kpABE.Enc}(\text{kpABE.mpk}, \text{lb}, \tilde{m}; F(K, (\text{lb}, x)))$, computed inside the $\text{Re-Enc}[\text{lb}, x, \gamma_1, \gamma_2, \delta]$ function is same for both f_0 and f_1 , and hence same on both the inputs μ_0 and μ_1 . So,

$$\text{Re-Enc}[\text{lb}, x, \gamma_1, \gamma_2, \delta](\mu_0) = \text{kpABE.ct}_{\text{lb}} \quad (\text{since mode} = 2 \text{ and flag} = 0).$$

$$\text{Re-Enc}[\text{lb}, x, \gamma_1, \gamma_2, \delta](\mu_1) = \text{kpABE.ct}_{\text{lb}} \quad (\text{since mode} = 2 \text{ and flag} = 0).$$

- When $\text{flag} = 1$, i.e. $f_0(x) \neq f_1(x)$, by definition of $\text{Re-Enc}[\text{lb}, x, \gamma_1, \gamma_2, \delta]$ we have

$$\begin{aligned} \text{Re-Enc}[\text{lb}, x, \gamma_1, \gamma_2, \delta](\mu_0) &= \text{ske.Dec}(\text{ske.key}_2, \gamma_2) \\ &= \text{kpABE.ct}'_{\text{lb}} \quad (\text{since mode} = 2 \text{ and flag} = 1). \end{aligned}$$

$$\begin{aligned} \text{Re-Enc}[\text{lb}, x, \gamma_1, \gamma_2, \delta](\mu_1) &= \text{ske.Dec}(\text{ske.key}_2, \gamma_2) \\ &= \text{kpABE.ct}'_{\text{lb}} \quad (\text{since mode} = 2 \text{ and flag} = 1). \end{aligned}$$

This establishes the admissibility of \mathcal{B} . ■

The rest of the hybrids, Hyb_9 to Hyb_{15} , are simply unwinding the previous hybrids and their proofs of indistinguishability are same as their corresponding counterparts in the first set of hybrids and hence, omitted. ■

Message Hiding

Theorem 3.22. *Assume that F is a secure PRF, SKE is correct and secure, FE and kpABE are secure as per definitions 3.3 and 2.6, respectively. Furthermore, assume $|\mathcal{F}_\lambda| \leq \text{poly}(\lambda)$ and $|\mathcal{M}_\lambda| \leq \text{poly}(\lambda)$. Then the construction for RPE satisfies message hiding property as defined in Def. 3.13.*

Proof. Recall that for message hiding we want

$$\text{RPE.Enc}(\text{RPE.mpk}, f, m_0, L) \approx_c \text{RPE.Enc}(\text{RPE.mpk}, f, m_1, L),$$

where for all the key queries (lb, x) , either $f(x) = 0$ or $\text{lb} \in L$. The proof is given via a similar sequence of hybrid games between the challenger and a PPT adversary \mathcal{A} as in the proof of Theorem 3.10. The hybrids are defined as follows:

Hyb₀. This is the real world with $\beta = 0$, i.e., the challenge ciphertext is computed using the message m_0 . We write the complete game here to set up notations and easy reference in later hybrids.

1. The adversary outputs the challenge messages (m_0, m_1) and the challenge function f ¹⁴.
2. The challenger generates $(\text{FE.mpk}, \text{FE.msk}) \leftarrow \text{FE.Setup}(1^\lambda)$, sets $\text{RPE.mpk} = \text{FE.mpk}$ and sends RPE.mpk to the adversary.
3. **Key Queries** : When adversary issues a key query on (lb, x) , the challenger does the following:
 - Samples random values $\gamma_1, \gamma_2, \delta \leftarrow \mathcal{CT}_{\text{SKE}}$.
 - Defines the circuit $\text{Re-Enc}[\text{lb}, x, \gamma_1, \gamma_2, \delta]$ as in the Figure 3.2.
 - Computes $\text{FE.sk}_{\text{lb}, x} \leftarrow \text{FE.KeyGen}(\text{FE.msk}, \text{Re-Enc}[\text{lb}, x, \gamma_1, \gamma_2, \delta])$.
 - Returns $\text{RPE.sk}_{\text{lb}, x} = \text{FE.sk}_{\text{lb}, x}$ to the adversary.
4. **Challenge Query** : When the adversary outputs the revocation list L for the challenge ciphertext, the challenger does the following:
 - Samples a F key K .
 - Generates $(\text{kpABE.mpk}, \text{kpABE.msk}) \leftarrow \text{kpABE.Setup}(1^\lambda)$.
 - Computes FE.ct as
$$\text{FE.Enc}(\text{FE.mpk}, (\text{kpABE.mpk}, f, m_0, K, 0, \perp, \perp)).$$
 - Defines C_L as in Eq. 3.1 and computes $\text{kpABE.sk}_L \leftarrow \text{kpABE.KeyGen}(\text{kpABE.msk}, C_L)$.

¹⁴To keep the proofs and notations simple, we let f, m_0, m_1 to be output selectively. This is sufficient to achieve security as in Def 3.13, as mentioned in Remark 4.

- Returns $\text{RPE.ct} = (\text{kpABE.mpk}, \text{kpABE.sk}_L, \text{FE.ct})$ to the adversary.

5. In the end, the adversary outputs a bit β' .

Hyb₁. This hybrid is same as the previous hybrid except the following changes:

- The challenger generates $(\text{kpABE.mpk}, \text{kpABE.msk}) \leftarrow \text{kpABE.Setup}(1^\lambda)$ in the beginning of the game after the adversary outputs (m_0, m_1) and f . It also samples a SKE secret key SKE.key_1 and a F key K .
- For each key query (lb, x) , γ_1 is computed differently. In particular, the challenger does the following:
 - Sets $\tilde{m}_0 = \begin{cases} m_0 & \text{if } f(x) = 1, \\ 0 & \text{if } f(x) = 0. \end{cases}$ and computes $\text{kpABE.ct}'_{\text{lb}} = \text{kpABE.Enc}(\text{kpABE.mpk}, \text{lb}, \tilde{m}_0; F(K, (\text{lb}, x)))$.
 - Sets γ_1 as $\text{SKE.Enc}(\text{SKE.key}_1, \text{kpABE.ct}'_{\text{lb}})$.

Hyb₂. This hybrid is same as the previous hybrid except the following:

- The challenger samples two SKE secret keys SKE.key_1 and SKE.key_2 .
- For each key query (lb, x) , γ_2 and δ are computed differently as follows:
 - Set $\text{flag} = \begin{cases} 1 & \text{if } f(x) = 1, \\ 0 & \text{otherwise.} \end{cases}$
 - Set $\delta \leftarrow \text{SKE.Enc}(\text{SKE.key}_2, \text{flag})$ and $\gamma_2 \leftarrow \text{SKE.Enc}(\text{SKE.key}_2, \text{kpABE.ct}'_{\text{lb}})$ if $\text{flag} = 1$, else $\gamma_2 \leftarrow \mathcal{CT}_{\text{SKE}}$.

Hyb₃. This hybrid is same as the previous hybrid except that FE.ct in the challenge ciphertext is computed as

$$\text{FE.Enc}(\text{FE.mpk}, (\text{kpABE.mpk}, f, m_0, \perp, 1, \text{SKE.key}_1, \text{SKE.key}_2)).$$

Hyb₄. This hybrid is same as the previous hybrid except that for each key query (lb, x)

$\text{kpABE.ct}'_{\text{lb}}$ is computed as follows:

- If $\text{flag} = 1$, $\text{kpABE.Enc}(\text{kpABE.mpk}, \text{lb}, \tilde{m}_0; r)$, where $r \leftarrow \{0, 1\}^t$,
- Else, if $\text{flag} = 0$, $\text{kpABE.ct}'_{\text{lb}} \leftarrow \text{kpABE.Enc}(\text{kpABE.mpk}, \text{lb}, \tilde{m}_0; F(K, (\text{lb}, x)))$. (This is same as in the previous hybrid).

Hyb₅. This hybrid is same as the previous hybrid except that FE.ct in the challenge ciphertext is computed as

$$\text{FE.Enc}(\text{FE.mpk}, (\text{kpABE.mpk}, f, m_0, K, 2, \perp, \text{SKE.key}_2)).$$

Hyb₆. This hybrid is same as the previous hybrid except that for all the key queries (lb, x) , $\gamma_1 \leftarrow \mathcal{CT}_{\text{SKE}}$.

Hyb₇. This hybrid is same as the previous hybrid except that for each such key query (lb, x) where $\text{flag} = 1$, $\text{kpABE.ct}'_{\text{lb}} \leftarrow \mathcal{CT}_{\text{kpABE}}$.

Hyb₈. This hybrid is same as the previous hybrid except that FE.ct in the challenge ciphertext is computed as

$$\text{FE.Enc}(\text{FE.mpk}, (\text{kpABE.mpk}, f, m_1, K, 2, \perp, \text{SKE.key}_2)).$$

We note that the hybrids hereafter are unwinding the changes made in the previous hybrids.

Hyb₉. This hybrid is same as the previous hybrid except that for each key query (lb, x) with $\text{flag} = 1$, $\text{kpABE.ct}'_{\text{lb}}$ is changed back to $\text{kpABE.Enc}(\text{kpABE.mpk}, \text{lb}, \tilde{m}_1; r)$, where $r \leftarrow \{0, 1\}^t$.

Hyb₁₀. This hybrid is same as the previous hybrid except that for each key query (lb, x) ,
 $\gamma_1 \leftarrow \text{SKE.Enc}(\text{SKE.key}_1, \text{kpABE.ct}'_{\text{lb}}).$

Hyb₁₁. This hybrid is same as the previous hybrid except that FE.ct in the challenge ciphertext is computed as

$$\text{FE.Enc}(\text{FE.mpk}, (\text{kpABE.mpk}, f, m_1, \perp, 1, \text{SKE.key}_1, \text{SKE.key}_2)).$$

Hyb₁₂. This hybrid is same as the previous hybrid except that for each key query (lb, x) ,
 $\text{kpABE.ct}'_{\text{lb}}$ is computed as $\text{kpABE.Enc}(\text{kpABE.mpk}, \text{lb}, \tilde{m}_1; F(K, (\text{lb}, x))).$

Hyb₁₃. This hybrid is same as the previous hybrid except that FE.ct in the challenge ciphertext is computed as

$$\text{FE.Enc}(\text{FE.mpk}, (\text{kpABE.mpk}, f, m_1, K, 0, \perp, \perp)).$$

Hyb₁₄. This hybrid is same as the previous hybrid except that the γ_2, δ are now sampled uniformly from $\mathcal{CT}_{\text{SKE}}$ for all the key queries.

Hyb₁₅. This hybrid is same as the previous hybrid except that γ_1 is now sampled uniformly from $\mathcal{CT}_{\text{SKE}}$ for all the key queries. This is the real world where the message m_1 is encrypted in the challenge ciphertext.

Indistinguishability of hybrids: The indistinguishability between the consecutive hybrids is argued in the same way as that in the proof of Theorem 3.10. Therefore, here we give only a brief sketch.

Hyb₀ \approx_c Hyb₁ \approx_c Hyb₂ from SKE security. Hyb₂ \approx_c Hyb₃ due to FE security and SKE correctness and Hyb₃ \approx_c Hyb₄ follows from PRF security. Hyb₄ \approx_c Hyb₅ follows from FE security and SKE correctness and Hyb₅ \approx_c Hyb₆ follows again from the SKE security.

$\text{Hyb}_6 \approx_c \text{Hyb}_7$ follows from selective security of ABE. We observe that in both Hyb_6 and Hyb_7 , $\text{ABE.ct}'_{\text{lb}}$ is not used when $\text{flag} = 0$ and when $\text{flag} = 1$, it is sampled from $C\mathcal{T}_{\text{kpABE}}$ directly which can be efficiently done without using kpABE.mpk . This lets the reduction go through. The steps of reduction are same as in the proof of Claim 3.17. $\text{Hyb}_7 \approx_c \text{Hyb}_8$ follows again from the security of FE and SKE correctness. In particular, we observe that here the FE challenge messages are $\mu_0 = (\text{kpABE.mpk}, f, m_0, K, 2, \perp, \text{SKE.key}_2)$ and $\mu_1 = (\text{kpABE.mpk}, f, m_1, K, 2, \perp, \text{SKE.key}_2)$. For every $\text{Re-Enc}[\text{lb}, x, \gamma_1, \gamma_2, \delta]$ function (corresponding to RPE key query (lb, x)) for which FE key is generated, we have the following:

- When $\text{flag} = 0$, this implies $f(x) = 0$, which in turn implies that $\tilde{m}_0 = \tilde{m}_1 = 0$. Hence,

$$\begin{aligned}
& \text{Re-Enc}[\text{lb}, x, \gamma_1, \gamma_2, \delta](\mu_0) \\
&= \text{Re-Enc}[\text{lb}, x, \gamma_1, \gamma_2, \delta](\text{kpABE.mpk}, f, m_0, K, 2, \perp, \text{SKE.key}_2) \\
&= \text{kpABE.Enc}(\text{kpABE.mpk}, \text{lb}, 0; F(K, (\text{lb}, x))) \\
&= \text{Re-Enc}[\text{lb}, x, \gamma_1, \gamma_2, \delta](\text{kpABE.mpk}, f, m_1, K, 2, \perp, \text{SKE.key}_2) \\
&= \text{Re-Enc}[\text{lb}, x, \gamma_1, \gamma_2, \delta](\mu_1).
\end{aligned}$$

- When $\text{flag} = 1$,

$$\begin{aligned}
& \text{Re-Enc}[\text{lb}, x, \gamma_1, \gamma_2, \delta](\mu_0) \\
&= \text{Re-Enc}[\text{lb}, x, \gamma_1, \gamma_2, \delta](\text{kpABE.mpk}, f, m_0, K, 2, \perp, \text{SKE.key}_2) \\
&= \text{SKE.Dec}(\text{SKE.key}_2, \gamma_2) \\
&= \text{Re-Enc}[\text{lb}, x, \gamma_1, \gamma_2, \delta](\text{kpABE.mpk}, f, m_1, K, 2, \perp, \text{SKE.key}_2) \\
&= \text{Re-Enc}[\text{lb}, x, \gamma_1, \gamma_2, \delta](\mu_1).
\end{aligned}$$

This satisfies the admissibility condition for FE security. The rest of the hybrids undo the changes made so far to get to the real world with $\beta = 1$ and the arguments for indistinguishability are same as their counterparts in the first set of hybrids. ■

3.5.3 Alternate Construction using LOT

Here, we consider an alternative construction using LOT. Compared to our construction in Sec. 3.5.1, the construction here can only handle the case where the number of users is

polynomially bounded and only achieves selective security. On the other hand, it can be based on FE and LOT rather than FE and kpABE with specific properties. Note that LOT can be based on more diverse assumptions than kpABE and this leads to an instantiation without LWE in particular.

The construction is similar to that in Sec. 3.5.1 except that we use LOT in place of ABE, which brings in the following changes in the KeyGen, Enc, and Dec algorithms:

- We use $\text{LOT} = (\text{LOT.crsGen}, \text{LOT.Hash}, \text{LOT.Send}, \text{LOT.Receive})$ instead of kpABE.
- The function in Figure 3.2, for which FE key is generated now takes as input LOT objects crs and digest , instead of kpABE.mpk and computes $\text{LOT.ct}_{\text{lb}} = \text{LOT.Send}(\text{crs}, \text{digest}, \text{lb}, 0, \tilde{m}; F(K, (\text{lb}, x)))$, instead of $\text{kpABE.ct}_{\text{lb}}$.
- The encryption algorithm changes as follows:
 $\text{RPE.Enc}(\text{RPE.mpk}, f, m, L) \rightarrow \text{RPE.ct}$. The encryption algorithm does the following:
 - Parse $\text{RPE.mpk} = \text{FE.mpk}$ and sample a F key $K \leftarrow \{0, 1\}^\lambda$.
 - Generate $\text{crs} \leftarrow \text{LOT.crsGen}(1^\lambda)$.
 - Compute $(\text{digest}, \hat{D}) \leftarrow \text{LOT.Hash}(\text{crs}, D)$, where D is a binary vector of length N (the number of users) and is 1 at positions corresponding to non-revoked labels, i.e. $D[\text{lb}'] = 1$ iff $\text{lb}' \notin L$.
 - Compute $\text{FE.ct} \leftarrow \text{FE.Enc}(\text{FE.mpk}, (\text{crs}, \text{digest}, f, m, K, 0, \perp, \perp))$.
 - Output $\text{RPE.ct} = (\text{crs}, \text{FE.ct})$.
- The algorithm for decryption also changes accordingly as follows:
 $\text{RPE.Dec}(\text{RPE.sk}_{\text{lb}, x}, \text{RPE.ct}, L) \rightarrow m'$. The decryption algorithm does the following:
 - Parse $\text{RPE.ct} = (\text{crs}, \text{FE.ct})$ and $\text{RPE.sk}_{\text{lb}, x} = \text{FE.sk}_{\text{lb}, x}$.
 - Define D from L as described in the encryption algorithm and compute $(\text{digest}, \hat{D}) \leftarrow \text{LOT.Hash}(\text{crs}, D)$.
 - Compute $\text{LOT.ct}' = \text{FE.Dec}(\text{FE.sk}_{\text{lb}, x}, \text{FE.ct})$.
 - Compute $m' = \text{LOT.Receive}^{\hat{D}}(\text{crs}, \text{LOT.ct}', \text{lb})$.

– Output m' .

We note that the above construction works when the identity space is of polynomial size.

Next we sketch the correctness, efficiency and security of the above construction.

Correctness. Firstly, for any label lb , attribute x , and function f such that $f(x) = 1$, we have $\text{FE.Dec}(\text{FE.sk}_{\text{lb},x}, \text{FE.ct}) = \text{LOT.ct}_{\text{lb}}$, where $\text{LOT.ct}_{\text{lb}} = \text{LOT.Send}(\text{crs}, \text{digest}, \text{lb}, 0, m; F(K, (\text{lb}, x)))$, by the correctness of FE and the definition of Re-Enc. Next we observe that if $\text{lb} \notin L$, then $D[\text{lb}] = 1$ and hence from the correctness of LOT scheme it follows that $\text{LOT.Receive}^{\hat{D}}(\text{crs}, \text{LOT.ct}_{\text{lb}}, \text{lb}) = m_{D[\text{lb}]} = m_1 = m$.

So the decryption correctly recovers the message when $f(x) = 1$ and $\text{lb} \notin L$.

Efficiency Here we argue that the above construction using LOT achieves optimal parameters. Namely, we show that the size of each parameter is independent from $|L|$. We note that $|f|$ refers to the description size of the function, not the circuit size that implements the function. When $|x|$ is very long and f has succinct description, the former can be much shorter than the latter.

1. Public key size $|\text{RPE.mpk}|$: We have $|\text{RPE.mpk}| = |\text{FE.mpk}|$. Since we assumed that FE is fully compact (Def. 3.4), the length of FE.mpk only depends on the input length of Re-Enc. We have that the input length is $|\text{crs}| + |\text{digest}| + |f| + |m| + |K| + |\text{mode}| + 2|\text{SKE.key}| = |\text{crs}| + |\text{digest}| + |f| + O(\lambda)$. We have $|\text{crs}| + |\text{digest}| = \text{poly}(\lambda)$ by the efficiency of LOT (Theorem 3.8). The total length of the public key is therefore $\text{poly}(\lambda, |f|)$.
2. Secret key size $|\text{RPE.sk}_{\text{lb},x}|$: We have $|\text{RPE.sk}_{\text{lb},x}| = |\text{FE.sk}_{\text{lb},x}|$. Since the size of the latter is polynomially dependent on the size of Re-Enc, we evaluate its size. We can see that the size of Re-Enc is polynomial in the total length of the input and the hardwired values. The length of the input is bounded by $\text{poly}(\lambda, |f|)$ as analyzed in the above item. The length of the hardwired values are $|\text{lb}| + |x| + |\gamma_1| + |\gamma_2| + |\delta|$. We have $|\gamma_1| + |\gamma_2| + |\delta| = 3|\gamma_2|$ ¹⁵ and $|\gamma_2| = \text{poly}(\lambda, \text{LOT.ct}_{\text{lb}}) = \text{poly}(\log |D|, \lambda) = \text{poly}(\lambda)$. Therefore, the size of Re-Enc is $\text{poly}(\lambda, |f|) + |\text{lb}| + |x|$ and so is the size of the secret key.

¹⁵We assumed that $|\gamma_1| = |\gamma_2| = |\delta|$. See Footnote 9.

3. Ciphertext size $|\text{RPE.ct}|$: We have $|\text{RPE.ct}| = |\text{crs}| + |\text{FE.ct}|$. We have $|\text{crs}| = \text{poly}(\lambda)$. Also, by similar analysis to the first item, full compactness of FE implies $|\text{FE.ct}| \leq \text{poly}(\lambda, |f|)$. Therefore, the overall length of the ciphertext is $\text{poly}(\lambda, |f|)$.

Security We show that the above construction satisfies the function hiding (Def. 3.17) and the message hiding (Def. 3.14) properties. The key difference here is that we only achieve selective security w.r.t the revoke list L .

Function Hiding The security proof for function hiding will follow the same sequence of hybrids as in the Theorem 3.10 except the following differences :

- In Hyb_0 , the challenger generates and uses LOT parameters to compute the challenge ciphertext. Concretely, the challenger computes FE.ct as $\text{FE.Enc}(\text{FE.mpk}, (\text{crs}, \text{digest}, f_0, m, K, 0, \perp, \perp))$ and returns $\text{RPE.ct} = (\text{crs}, \text{FE.ct})$ to the adversary.
- In Hyb_1 to Hyb_3 , the challenger computes $\text{LOT.ct}'_{\text{lb}} = \text{LOT.Send}(\text{crs}, \text{digest}, \text{lb}, 0, \tilde{m}; F(K, (\text{lb}, x)))$ instead of $\text{ABE.ct}'_{\text{lb}}$ while answering the key queries.
- Similarly, in Hyb_4 to Hyb_6 , for each key query (lb, x) , instead of $\text{ABE.ct}'_{\text{lb}}$, $\text{LOT.ct}'_{\text{lb}}$ is computed as follows:
 - If $\text{flag} = 1$, $\text{LOT.ct}'_{\text{lb}} = \text{LOT.Enc}(\text{crs}, \text{digest}, \text{lb}, 0, \tilde{m}; r)$, where $r \leftarrow \{0, 1\}^t$.
 - Else, if $\text{flag} = 0$, $\text{LOT.ct}'_{\text{lb}} = \text{LOT.Enc}(\text{crs}, \text{digest}, \text{lb}, 0, \tilde{m}; F(K, (\text{lb}, x)))$.
- In Hyb_7 , when $\text{flag} = 1$, $\text{LOT.ct}'_{\text{lb}}$ is simulated using LOTSim , i.e., $\text{LOT.ct}'_{\text{lb}} \leftarrow \text{LOTSim}(D, \text{lb}, 0)$. We observe that when $\text{flag} = 1$, $\text{lb} \in L$ (by admissibility), so $D[\text{lb}] = 0$ and we have $m_{D[\text{lb}]} = m_0 = 0$. Hence, the indistinguishability between Hyb_6 and Hyb_7 follows from the sender privacy of LOT (Def. 3.10).

We note that after Hyb_8 , we rewind the changes made in the preceding hybrids accordingly. The reason why we only achieve selective security w.r.t the revoke list L is that while answering the key queries, computation of $\text{LOT.ct}'_{\text{lb}}$ uses digest . This digest is computed using the database D , which in turn is derived using the revoke list L .

Message Hiding The proof for message hiding is given via the same sequence of hybrids as in the Theorem 3.22. The differences in the hybrids are similar to the case of function hiding as highlighted in the above paragraph, where we use LOT parameters to compute the challenge ciphertext and $\text{LOT.ct}'_{\text{lb}}$ instead of $\text{ABE.ct}'_{\text{lb}}$ while answering the key queries.

3.6 REVOCABLE MIXED FUNCTIONAL ENCRYPTION

3.6.1 Definition

A revocable mixed functional encryption (RMFE) scheme with input domain $\mathcal{X} = \{\mathcal{X}_\lambda\}_{\lambda \in [\mathbb{N}]}$, a function family $\mathcal{F} = \{\mathcal{F}_\lambda\}_{\lambda \in [\mathbb{N}]}$ where $\mathcal{F}_\lambda = \{f : \mathcal{X}_\lambda \rightarrow \{0, 1\}\}$, a label space $\mathcal{L} = \{\mathcal{L}_\lambda\}_{\lambda \in [\mathbb{N}]}$ has the following syntax.

$\text{Setup}(1^\lambda) \rightarrow (\text{mpk}, \text{msk})$. The setup algorithm takes as input the security parameter λ and outputs a master public key mpk and a master secret key msk .

$\text{KeyGen}(\text{msk}, \text{lb}, x) \rightarrow \text{sk}_{\text{lb}, x}$. The key generation algorithm takes as input the master secret key msk , a label $\text{lb} \in \mathcal{L}_\lambda$ and an input $x \in \mathcal{X}_\lambda$. It outputs a secret key $\text{sk}_{\text{lb}, x}$.

$\text{PK-Enc}(\text{mpk}, L) \rightarrow \text{ct}$. The public key encryption algorithm takes as input the master public key mpk and a revocation list $L \subseteq \mathcal{L}_\lambda$ and outputs a ciphertext ct .

$\text{SK-Enc}(\text{msk}, f, L) \rightarrow \text{ct}$. The secret key encryption algorithm takes as input the master secret key msk , a function $f \in \mathcal{F}_\lambda$ and a revocation list $L \subseteq \mathcal{L}_\lambda$, and outputs a ciphertext ct .

$\text{Dec}(\text{sk}_{\text{lb}, x}, L, \text{ct}) \rightarrow \{0, 1\}$. The decryption algorithm takes the secret key $\text{sk}_{\text{lb}, x}$, a revocation list $L \subseteq \mathcal{L}_\lambda$ and a ciphertext ct and outputs a bit.

Definition 3.19 (Correctness). A RMFE scheme is said to be correct if there exists

negligible functions $\text{negl}_1(\cdot), \text{negl}_2(\cdot)$ such that for all $\lambda \in \mathbb{N}$, the following holds

$$\Pr \left[\begin{array}{l} (\text{mpk}, \text{msk}) \leftarrow \text{Setup}(1^\lambda); \\ \text{Dec}(\text{sk}_{\text{lb},x}, L, \text{ct}) = 1 : \text{sk}_{\text{lb},x} \leftarrow \text{KeyGen}(\text{msk}, \text{lb}, x); \\ \text{ct} \leftarrow \text{PK-Enc}(\text{mpk}, L) \end{array} \right] \geq 1 - \text{negl}_1(\lambda).$$

$$\text{lb} \notin L \Rightarrow \Pr \left[\begin{array}{l} (\text{mpk}, \text{msk}) \leftarrow \text{Setup}(1^\lambda); \\ \text{Dec}(\text{sk}_{\text{lb},x}, L, \text{ct}) = f(x) : \text{sk}_{\text{lb},x} \leftarrow \text{KeyGen}(\text{msk}, \text{lb}, x); \\ \text{ct} \leftarrow \text{SK-Enc}(\text{msk}, f, L) \end{array} \right] \geq 1 - \text{negl}_2(\lambda).$$

Security. Here we define the security requirements of RMFE scheme.

Definition 3.20 (q -query Mode Hiding). Let $q(\cdot)$ be any fixed polynomial. A RMFE scheme satisfies q -query mode hiding security if for every PPT adversary \mathcal{A} , there exists a negligible function $\text{negl}(\cdot)$ such that for every $\lambda \in \mathbb{N}$,

$$\Pr \left[\begin{array}{l} (\text{mpk}, \text{msk}) \leftarrow \text{Setup}(1^\lambda); \\ f, L \leftarrow \mathcal{A}^{\text{KeyGen}(\text{msk}, \cdot, \cdot), \text{SK-Enc}(\text{msk}, \cdot, \cdot)}(\text{mpk}); \\ \beta^{\text{Prime}} = \beta : \beta \leftarrow \{0, 1\}; \text{ct}_0 \leftarrow \text{SK-Enc}(\text{msk}, f, L); \\ \text{ct}_1 \leftarrow \text{PK-Enc}(\text{mpk}, L); \\ \beta^{\text{Prime}} \leftarrow \mathcal{A}^{\text{KeyGen}(\text{msk}, \cdot, \cdot), \text{SK-Enc}(\text{msk}, \cdot, \cdot)}(\text{ct}_\beta) \end{array} \right] \leq \frac{1}{2} + \text{negl}(\lambda)$$

where \mathcal{A} can make at most $q(\lambda)$ queries to the $\text{SK-Enc}(\text{msk}, \cdot, \cdot)$ oracle and is admissible only if for all the key queries (lb, x) to the $\text{KeyGen}(\text{msk}, \cdot, \cdot)$ oracle, $f(x) = 1$.

Definition 3.21 (q -query Selective Function Hiding). Let $q(\cdot)$ be any fixed polynomial. A RMFE scheme satisfies q -query selective function hiding security if for every PPT adversary \mathcal{A} , there exists a negligible function $\text{negl}(\cdot)$ such that for every $\lambda \in \mathbb{N}$,

$$\Pr \left[\begin{array}{l} L \leftarrow \mathcal{A}(1^\lambda); \\ (\text{mpk}, \text{msk}) \leftarrow \text{Setup}(1^\lambda); \\ \beta^{\text{Prime}} = \beta : (f_0, f_1) \leftarrow \mathcal{A}^{\text{KeyGen}(\text{msk}, \cdot, \cdot), \text{SK-Enc}(\text{msk}, \cdot, \cdot)}(\text{mpk}); \\ \beta \leftarrow \{0, 1\}; \text{ct}_\beta \leftarrow \text{SK-Enc}(\text{msk}, f_\beta, L); \\ \beta^{\text{Prime}} \leftarrow \mathcal{A}^{\text{KeyGen}(\text{msk}, \cdot, \cdot), \text{SK-Enc}(\text{msk}, \cdot, \cdot)}(\text{ct}_\beta) \end{array} \right] \leq \frac{1}{2} + \text{negl}(\lambda)$$

where \mathcal{A} can make at most $q(\lambda)$ queries to the $\text{SK-Enc}(\text{msk}, \cdot, \cdot)$ oracle and for all the key queries (lb, x) to the $\text{KeyGen}(\text{msk}, \cdot, \cdot)$ oracle, either $f_0(x) = f_1(x)$ or $\text{lb} \in L$.

Remark 5. We note that when the function space \mathcal{F}_λ is polynomially small and q is a constant, a variant of Definition 3.21 where the adversary outputs the challenge functions (f_0, f_1) and the SK-Enc query functions $\{\tilde{f}_i\}_{i \in [q]}$ at the beginning of the game, before the $\text{Setup}(1^\lambda)$ algorithm is run, is equivalent to Definition 3.21 where the adversary adaptively outputs the challenge functions (f_0, f_1) and can make SK-Enc queries adaptively, with polynomial loss. Similar comment also applies to Definition 3.20. We will use these simplifications in the security proofs.

3.6.2 Construction

In this section we give a construction of 1-query secure RMFE scheme, with input space $\mathcal{X} = \{\mathcal{X}_\lambda\}_\lambda$, a function family $\mathcal{F} = \{\mathcal{F}_\lambda\}_\lambda$ where $\mathcal{F}_\lambda = \{f : \mathcal{X}_\lambda \rightarrow \{0, 1\}\}$ and a label space $\mathcal{L} = \{\mathcal{L}_\lambda\}_\lambda$. We assume that the size of $|\mathcal{F}_\lambda|$ is bounded by some polynomial in λ , which will suffice for our purpose.

Our scheme uses the following building blocks:

1. A 2-bounded semi-adaptive simulation based function-message private (Definition 3.6) SKFE scheme $\text{SKFE} = (\text{SKFE.Setup}, \text{SKFE.KeyGen}, \text{SKFE.Enc}, \text{SKFE.Dec})$ that supports the function class \mathcal{F} . This can be instantiated from one-way functions (Lemma 3.2).
2. A key-policy ABE scheme $\text{kpABE} = (\text{kpABE.Setup}, \text{kpABE.Enc}, \text{kpABE.KeyGen}, \text{kpABE.Dec})$ for the circuit class $\mathcal{C}_{\ell(\lambda), d(\lambda)}$ with message space $\{0, 1\}^\lambda$ satisfying Sel-IND security (Definition 2.5) and efficiency properties described in Theorem 3.3. We set $\ell(\lambda) = \ell_{\text{lb}} + \log(\lambda) + 1$ and $d(\lambda) = \omega(\log \lambda)$, where ℓ_{lb} is the label length.¹⁶ This can be instantiated from the LWE assumption (Theorem 3.3).
3. A lockable obfuscation scheme $\text{LO} = (\text{LO.Obf}, \text{LO.Eval})$ with lock space $\{0, 1\}^\lambda$ that supports circuits of the form CC defined in Fig. 3.3. As we will analyze later, the circuit is of fixed polynomial size in λ and $|f|$, where $|f|$ is the description size of the function $f \in \mathcal{F}$. This can be instantiated from the LWE assumption (Theorem 3.7).

Below we describe our construction of a 1-query secure RMFE scheme $\text{RMFE} =$

¹⁶Concretely, we can choose $d(\lambda) = \Theta(\log \lambda \log \log \lambda)$ for example.

(RMFE.Setup, RMFE.KeyGen, RMFE.PK-Enc, RMFE.SK-Enc, RMFE.Dec).

$\text{RMFE.Setup}(1^\lambda) \rightarrow (\text{RMFE.mpk}, \text{RMFE.msk})$. The setup algorithm does the following:

- Generate $\text{SKFE.msk} \leftarrow \text{SKFE.Setup}(1^\lambda)$.
- Generate $(\text{kpABE.mpk}, \text{kpABE.msk}) \leftarrow \text{kpABE.Setup}(1^\lambda)$.
- Output $\text{RMFE.mpk} = \text{kpABE.mpk}$ and $\text{RMFE.msk} = (\text{SKFE.msk}, \text{kpABE.mpk}, \text{kpABE.msk})$.

$\text{RMFE.KeyGen}(\text{RMFE.msk}, \text{lb}, x) \rightarrow \text{RMFE.sk}_{\text{lb},x}$. The key generation algorithm does the following:

- Parse $\text{RMFE.msk} = (\text{SKFE.msk}, \text{kpABE.mpk}, \text{kpABE.msk})$.
- For all $j \in [\lambda], b \in \{0, 1\}$, sample $K_{j,b}, R_{j,b} \leftarrow \{0, 1\}^\lambda$.
Denote $K = \{K_{j,b}\}_{j \in [\lambda], b \in \{0,1\}}$ and $R = \{R_{j,b}\}_{j \in [\lambda], b \in \{0,1\}}$.

- Compute

$$\text{SKFE.ct} \leftarrow \text{SKFE.Enc}(\text{SKFE.msk}, (x, K, R)).$$

- For all $j \in [\lambda], b \in \{0, 1\}$, compute

$$\text{kpABE.ct}_{\text{lb},j,b} \leftarrow \text{kpABE.Enc}(\text{kpABE.mpk}, (\text{lb}, j, b), K_{j,b}).$$

- Output $\text{RMFE.sk}_{\text{lb},x} = (\text{SKFE.ct}, \text{kpABE.mpk}, \{(\text{lb}, j, b), \text{kpABE.ct}_{\text{lb},j,b}\}_{j \in [\lambda], b \in \{0,1\}})$.

$\text{RMFE.PK-Enc}(\text{RMFE.mpk}, L) \rightarrow \text{RMFE.ct}$. The public key encryption algorithm does the following:

- Computes a simulated code $\text{RMFE.ct} \leftarrow \text{LO.Sim}(1^\lambda, 1^{|\text{CC}|})$ ¹⁷.
- It outputs RMFE.ct as the ciphertext.

¹⁷Here, CC represents the maximum possible size of $\text{CC}[\cdot, \cdot]$ circuit defined in Figure 3.3.

$\text{RMFE.SK-Enc}(\text{RMFE.msk}, f, L) \rightarrow \text{RMFE.ct}$. The secret key encryption algorithm does the following:

- Parse $\text{RMFE.msk} = (\text{SKFE.msk}, \text{kpABE.mpk}, \text{kpABE.msk})$, and sample a tag $\mathbf{z} \leftarrow \{0, 1\}^\lambda$ and a lock value $\alpha \leftarrow \{0, 1\}^\lambda$.
- For all $j \in [\lambda]$, compute $\text{kpABE.sk}_{L,j,z_j} \leftarrow \text{kpABE.KeyGen}(\text{kpABE.msk}, C_{L,j,z_j})$, where the function C_{L,j,z_j} has L, j and z_j hardwired and is defined as follows :
On input $(\text{lb}, i, b) \in \mathcal{L}_\lambda \times [\lambda] \times \{0, 1\}$,

$$C_{L,j,z_j}(\text{lb}, i, b) = \begin{cases} 1 & \text{if } (\text{lb} \notin L) \wedge (i = j) \wedge (b = z_j) \\ 0 & \text{otherwise.} \end{cases} \quad (3.2)$$

- Compute $\text{SKFE.sk} \leftarrow \text{SKFE.KeyGen}(\text{SKFE.msk}, P_{f,\mathbf{z},\alpha})$, where the function $P_{f,\mathbf{z},\alpha}$ has f, \mathbf{z}, α hardwired and is defined as follows :
On input $x \in \mathcal{X}_\lambda, K = \{K_{j,b}\}_{j \in [\lambda], b \in \{0,1\}}, R = \{R_{j,b}\}_{j \in [\lambda], b \in \{0,1\}}$,

$$P_{f,\mathbf{z},\alpha}(x, K, R) = \begin{cases} \bigoplus_j K_{j,z_j} \oplus \alpha & \text{if } f(x) = 0 \\ \bigoplus_j R_{j,z_j} & \text{if } f(x) = 1. \end{cases} \quad (3.3)$$

- Construct function $\text{CC}[\text{SKFE.sk}, \{\text{kpABE.sk}_{L,j,z_j}\}_{j \in [\lambda]}]$, with SKFE.sk and $\{\text{kpABE.sk}_{L,j,z_j}\}_{j \in [\lambda]}$ hardwired and is defined as in Figure 3.3.
- Output $\text{RMFE.ct} \leftarrow \text{LO.Obf}(\text{CC}[\text{SKFE.sk}, \{\text{kpABE.sk}_{L,j,z_j}\}_{j \in [\lambda]}], \alpha)$.

$\text{RMFE.Dec}(\text{RMFE.sk}_{\text{lb},x}, \text{RMFE.ct}, L) \rightarrow \{0, 1\}$. The decryption algorithm does the following:

- Parse $\text{RMFE.sk}_{\text{lb},x} = (\text{SKFE.ct}, \text{kpABE.mpk}, \{(\text{lb}, j, b), \text{kpABE.ct}_{\text{lb},j,b}\}_{j \in [\lambda], b \in \{0,1\}})$ and $\text{RMFE.ct} = \widetilde{\text{CC}}$, where $\widetilde{\text{CC}}$ is regarded as an obfuscated circuit of LO.

- For all $j \in [\lambda], b \in \{0, 1\}$, compute

$$\text{kpABE.off}_{\text{lb},j,b} \leftarrow \text{kpABE.Dec}^{\text{off}}(\text{kpABE.mpk}, C_{L,j,b}, (\text{lb}, j, b)).$$

- Compute

$$y = \text{LO.Eval}(\widetilde{\text{CC}}, (\text{SKFE.ct}, \{\text{kpABE.ct}_{\text{lb},j,b}, \text{kpABE.off}_{\text{lb},j,b}\}_{j \in [\lambda], b \in \{0,1\}})).$$

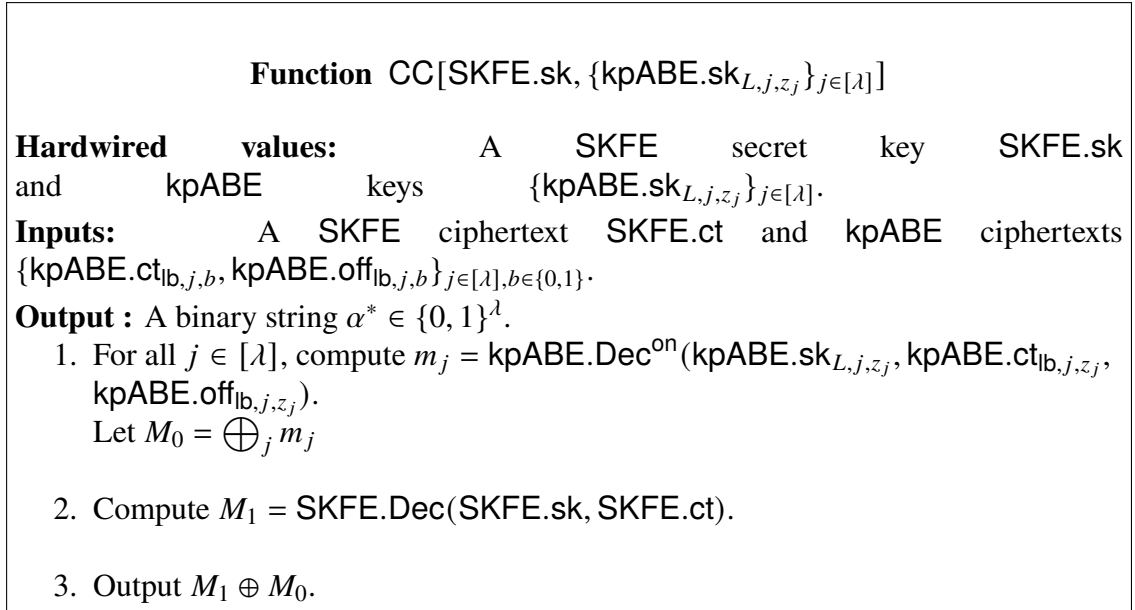


Figure 3.3: Compute and Compare function CC

- Output 1 if $y = \perp$, else output 0.

Remark 6. We note that by performing the part of the ABE decryption that uses $C_{L,j,b}$, outside of CC, we do not need to provide $C_{L,j,b}$ (or L) as input to CC. Instead, we provide $\{\text{kpABE.off}_{\text{lb},j,b}\}_{j \in [\lambda], b \in \{0,1\}}$ whose size is independent of the size of $C_{L,j,b}$ (and thus that of L). This helps us in getting succinct ciphertext.

Correctness. We prove the correctness via the following theorem.

Theorem 3.23. *Suppose kpABE, LO and SKFE are correct and LO is secure, then the above construction of RMFE satisfies correctness as defined in Def. 3.19.*

Proof. We consider the following two cases:

1. Public Encryption Correctness.

For $\text{RMFE.ct} \leftarrow \text{RMFE.PK-Enc}(\text{RMFE.mpk}, L)$, we have that $\text{RMFE.ct} = \text{LO.Sim}(1^\lambda, 1^{|\text{CC}|})$. Firstly, we note that from LO security, RMFE.ct is indistinguishable from $\text{RMFE.ct}'$ computed as $\text{LO.Obf}(C, \alpha)$, for any circuit C of the same size as that used by the simulator and has output of length λ . Now, since $\alpha \leftarrow \{0, 1\}^\lambda$ has high entropy, for all but negligible inputs w , $C(w) \neq \alpha$. So, from the correctness of LO, it follows that with all but negligible probability

$$\text{LO.Eval}(\text{RMFE.ct}, (\text{SKFE.ct}, \{\text{kpABE.ct}_{\text{lb},j,b}, \text{kpABE.off}_{\text{lb},j,b}\}_{j \in [\lambda], b \in \{0,1\}})) = \perp.$$

Hence the RMFE.Dec algorithm outputs 1 with all but negligible probability.

2. Secret Encryption Correctness.

For $\text{RMFE.ct} \leftarrow \text{RMFE.SK-Enc}(\text{RMFE.msk}, f, L)$, we have $\text{RMFE.ct} = \text{LO.Obf}(\text{CC}[\text{SKFE.sk}, \{\text{kpABE.sk}_{L,j,z_j}\}_{j \in [\lambda]}], \alpha)$. Now consider the following steps of $\text{LO.Eval}(\text{RMFE.ct}, (\text{SKFE.ct}, \{\text{kpABE.ct}_{\text{lb},j,b}, \text{kpABE.off}_{\text{lb},j,b}\}_{j \in [\lambda], b \in \{0,1\}}))$

- Since $\text{lb} \notin L$, by correctness of kpABE , for all $j \in [\lambda]$, $M_0 = \oplus_j m_j = \oplus_j K_{j,z_j}$ with all but negligible probability.
- By correctness of SKFE , if $f(x) = 0$, we have $M_1 = \bigoplus_j K_{j,z_j} \oplus \alpha$, else $M_1 = \bigoplus_j R_{j,z_j}$.
- We have $M_0 \oplus M_1 = \alpha$ if $f(x) = 0$.

So, by the correctness of LO , LO.Eval outputs 1 if $f(x) = 0$, \perp otherwise with all but negligible probability. Hence the RMFE.Dec algorithm, by construction, outputs 0 if $f(x) = 0$ and 1 if $f(x) = 1$ with all but negligible probability.

This proves the correctness of the above construction. \blacksquare

Efficiency. Here we argue that our construction achieves optimal parameters. Namely, we show that the sizes of the parameters are independent of $|L|$. We first observe that $C_{L,j,b}$ as defined in Eq (3.2) can be implemented with depth $d = \omega(\log \lambda)$, since the membership check $\text{lb} \stackrel{?}{\in} L$ can be done with depth $\log(|\text{lb}| \cdot |L|) = \log(\text{poly}(\lambda))$ and the equality check can be done with depth $\log(|j|) = \log \log \lambda$. We then bound the size of parameters.

1. Public key size $|\text{RMFE.mpk}|$: By the efficiency property of kpABE (Theorem 3.3), we have $|\text{RMFE.mpk}| = |\text{kpABE.mpk}| = \text{poly}(\lambda, d, |\text{lb}|) = \text{poly}(\lambda, |\text{lb}|)$.
2. Secret key size $|\text{RMFE.sk}_{\text{lb},x}|$: We have $|\text{RMFE.sk}_{\text{lb},x}| = |\text{SKFE.ct}| + |\text{kpABE.mpk}| + 2 \cdot \lambda(|\text{lb}, j, b|) + |\text{kpABE.ct}|$. The first term can be bounded by $\text{poly}(\lambda, |f|, |x|)$, the second is $\text{poly}(\lambda, |\text{lb}|)$, and the last terms is $\text{poly}(\lambda, |\text{lb}|)$. Therefore, the total size is $\text{poly}(\lambda, |f|, |x|, |\text{lb}|)$.
3. Ciphertext size $|\text{RMFE.ct}|$: We first bound the size of the circuit CC defined in Fig. 3.3. The dominant operations in the circuit is the decryption of SKFE and the online decryption of kpABE . We can see that the former is implemented by a circuit of size $\text{poly}(\lambda, |x|, |f|)$ by the efficiency of SKFE . The latter can be implemented by a circuit of size $\text{poly}(\lambda, \ell, d) = \text{poly}(\lambda, \text{lb})$ by the online efficiency

of kpABE (Theorem 3.3). Therefore, the total size CC is $\text{poly}(\lambda, |f|, |x|, |\mathbf{lb}|)$. By the efficiency of LO , the size of the ciphertext is $\text{poly}(\lambda, |f|, |x|, |\mathbf{lb}|)$ as well.

4. Depth of the circuit implementing RMFE.Dec : We also evaluate the depth of the circuit implementing RMFE.Dec and show that it is independent from $|L|$, since it will be used later in Sec. 3.7. We first observe that the dominant operations in RMFE.Dec are the offline decryption of kpABE and the evaluation of CC . The depth of the former can be bounded by $\text{poly}(\lambda, \ell, \text{depth}(C_{L,j,b})) \leq \text{poly}(\lambda, |\mathbf{lb}|)$ by Theorem 3.3. The depth of the latter can be bounded by its size and thus is $\text{poly}(\lambda, |f|, |x|, |\mathbf{lb}|)$ as we have seen in the previous item. The total depth is thus bounded by $\text{poly}(\lambda, |f|, |x|, |\mathbf{lb}|)$, which is independent from $|L|$.

3.6.3 Security

In this section we show that our construction of RMFE scheme satisfies all the security requirements.

We will use the following notations in the security proof:

For $X \in \{K, R\}$,

- $X := \{X_{j,b}\}_{j \in [\lambda], b \in \{0,1\}}$, $X_b := \{X_{j,b}\}_{j \in [\lambda]}$, $X_j := \{X_{j,b}\}_{b \in \{0,1\}}$.
- For any vector $\mathbf{z} \in \{0, 1\}^\lambda$, $X_{\mathbf{z}} := \{X_{j,z_j}\}_{j \in [\lambda]}$.

Mode hiding

Theorem 3.24. *Assume that SKFE and LO are secure as per definitions 3.6 and 3.8, respectively. Furthermore, assume $|\mathcal{F}_\lambda| \leq \text{poly}(\lambda)$. Then the RMFE construction satisfies 1-query mode hiding security as per Definition 3.20.*

Proof. Recall that for mode hiding, we need

$$\text{RMFE.SK-Enc}(\text{RMFE.msk}, f^*, L^*) \approx_c \text{RMFE.PK-Enc}(\text{RMFE.mpk}, L^*),$$

where for all key queries (\mathbf{lb}, x) , $f^*(x) = 1$.

The proof proceeds via the following sequence of hybrid games between the challenger and a PPT adversary \mathcal{A} .

Hyb_0 : This is the real world with $\beta = 0$, where the challenge ciphertext for (f^*, L^*) is computed using the RMFE.SK-Enc algorithm. We write the complete game here to set up the notations and easy reference in later hybrids.

1. The adversary outputs the challenge function f^* and the SK-Enc query function \tilde{f} ¹⁸.
2. The challenger generates $\text{SKFE.msk} \leftarrow \text{SKFE.Setup}(1^\lambda)$, $(\text{kpABE.mpk}, \text{kpABE.msk}) \leftarrow \text{kpABE.Setup}(1^\lambda)$, sets $\text{RMFE.mpk} = \text{kpABE.mpk}$ and $\text{RMFE.msk} = (\text{SKFE.msk}, \text{kpABE.mpk}, \text{kpABE.msk})$. It sends RMFE.mpk to \mathcal{A} .
3. **Key Queries:** For each key query (lb, x) , the challenger computes SKFE.ct and $\text{kpABE.ct}_{\text{lb},j,b}$ as in the construction and returns $\text{RMFE.sk}_{\text{lb},x} = (\text{SKFE.ct}, \text{kpABE.mpk}, \{(\text{lb}, j, b), \text{kpABE.ct}_{\text{lb},j,b}\}_{j \in [\lambda], b \in \{0,1\}})$ to \mathcal{A} .
4. **Challenge Query:** When the adversary outputs L^* for the challenge query, the challenger does the following:
 - Samples a tag $\mathbf{z}^* \leftarrow \{0, 1\}^\lambda$ and a lock value $\alpha^* \leftarrow \{0, 1\}^\lambda$.
 - Computes kpABE secret keys $\{\text{kpABE.sk}_{L^*,j,z_j^*}\}_{j \in [\lambda]}$ and a SKFE secret key SKFE.sk^* as in the construction.
 - Constructs $\text{CC}[\text{SKFE.sk}^*, \{\text{kpABE.sk}_{L^*,j,z_j^*}\}_{j \in [\lambda]}]$ as defined in Figure 3.3 and returns $\text{RMFE.ct}^* \leftarrow \text{LO.Obf}(\text{CC}[\text{SKFE.sk}^*, \{\text{kpABE.sk}_{L^*,j,z_j^*}\}_{j \in [\lambda]}], \alpha^*)$ to the adversary \mathcal{A} .
5. **SK-Enc Query:** When the adversary outputs \bar{L} for the SK-Enc query, the challenger does the following:
 - Samples a tag $\bar{\mathbf{z}} \leftarrow \{0, 1\}^\lambda$ and a lock value $\bar{\alpha} \leftarrow \{0, 1\}^\lambda$.
 - Computes kpABE secret keys $\{\text{kpABE.sk}_{\bar{L},j,\bar{z}_j}\}_{j \in [\lambda]}$ and a SKFE secret key $\text{SKFE.s}\bar{\mathbf{k}}$ as in the construction.
 - Constructs $\text{CC}[\text{SKFE.s}\bar{\mathbf{k}}, \{\text{kpABE.sk}_{\bar{L},j,\bar{z}_j}\}_{j \in [\lambda]}]$ and returns $\text{RMFE.c}\bar{\mathbf{t}} \leftarrow \text{LO.Obf}(\text{CC}[\text{SKFE.s}\bar{\mathbf{k}}, \{\text{kpABE.sk}_{\bar{L},j,\bar{z}_j}\}_{j \in [\lambda]}], \bar{\alpha})$ to the adversary \mathcal{A} .
6. In the end, \mathcal{A} outputs a bit β' .

¹⁸To keep the proofs and notations simple, we let f^* and \tilde{f} to be given selectively. This is sufficient to achieve security as in Def 3.20, as mentioned in Remark 5.

Hyb₁ : This hybrid is same as the previous hybrid except the following changes:

1. The challenger samples $\alpha^*, \bar{\alpha}, \mathbf{z}^*, \bar{\mathbf{z}}$ in the beginning of the game after the adversary outputs f^*, \bar{f} .
2. The challenger then computes $\text{SKFE.sk}^* \leftarrow \text{SKFE.Sim}^{\text{SK}}(\text{SKFE.msk}, 1^{|P_{f^*, \mathbf{z}^*, \alpha^*}|})$ and $\text{SKFE.s}\bar{\mathbf{k}} \leftarrow \text{SKFE.Sim}^{\text{SK}}(\text{SKFE.msk}, 1^{\text{poly}(|P_{\bar{f}, \bar{\mathbf{z}}, \bar{\alpha}}|)})$ in this order using SKFE simulators.
3. The key generation, challenge and SK-Enc queries are answered as follows:
 - For each key query (lb, x) , the SKFE ciphertext in $\text{RMFE.sk}_{\text{lb}, x}$ is computed as $\text{SKFE.Sim}^{\text{CT}}(\text{SKFE.msk}, \bigoplus_j R_{j, \mathbf{z}_j^*}, P_{\bar{f}, \bar{\mathbf{z}}, \bar{\alpha}}(x, K, R))$. Note that $P_{f^*, \mathbf{z}^*, \alpha^*}(x, K, R) = \bigoplus_j R_{j, \mathbf{z}_j^*}$, due to admissibility requirement.
 - To answer the challenge and the SK-Enc queries, SKFE.sk^* and $\text{SKFE.s}\bar{\mathbf{k}}$ generated by SKFE simulators in Step 2 are used for generating RMFE.ct^* and $\text{RMFE.ct}\bar{}$, respectively.

Hyb₂ : This hybrid is same as the previous hybrid except that the challenger uses LO.Sim to generate the challenge ciphertext.

$$\text{RMFE.ct}^* = \text{LO.Sim}(1^\lambda, 1^{|\text{CC}|})^{19}.$$

Hyb₃ : This hybrid is same as the previous hybrid except that the challenger uses SKFE.Enc and SKFE.KeyGen to generate SKFE ciphertexts and keys respectively. Formally,

$$\text{SKFE.s}\bar{\mathbf{k}} = \text{SKFE.KeyGen}(\text{SKFE.msk}, P_{\bar{f}, \bar{\mathbf{z}}, \bar{\alpha}})$$

For each key query (lb, x) ,

$$\text{SKFE.ct} = \text{SKFE.Enc}(\text{SKFE.msk}, (K, R, x)),$$

where vectors K and R are freshly sampled for each key as defined in the

¹⁹Here, $|\text{CC}|$ represents the maximum size of the circuit $\text{CC}[\cdot, \cdot]$ defined in Figure 3.3.

construction. This is the real world with $\beta = 1$, where the challenge ciphertext is computed using the RMFE.PK-Enc algorithm.

Indistinguishability of hybrids

Claim 3.25. Assume that SKFE is secure (Def. 3.6), then $\text{Hyb}_0 \approx_c \text{Hyb}_1$.

Proof. We show that if there exists a PPT adversary \mathcal{A} who can distinguish between Hyb_0 and Hyb_1 with non-negligible advantage ϵ , then there exists a PPT adversary \mathcal{B} against the security of SKFE scheme with the same advantage ϵ . The reduction is as follows.

1. Upon being invoked by the SKFE challenger, \mathcal{B} invokes \mathcal{A} which outputs f^*, \bar{f} .
2. \mathcal{B} samples $\alpha^*, \bar{\alpha}, \mathbf{z}^*, \bar{\mathbf{z}} \leftarrow \{0, 1\}^\lambda$.
3. \mathcal{B} defines the functions $P_{f^*, \mathbf{z}^*, \alpha^*}$ and $P_{\bar{f}, \bar{\mathbf{z}}, \bar{\alpha}}$ as defined in Eq. (3.3) and sends it to the SKFE challenger in this order as key queries. The challenger generates $\text{SKFE.msk} \leftarrow \text{SKFE.Setup}(1^\lambda)$ and samples $\hat{\beta} \leftarrow \{0, 1\}$.
 If $\hat{\beta} = 0$, it computes $\text{SKFE.sk}^* = \text{SKFE.KeyGen}(\text{SKFE.msk}, P_{f^*, \mathbf{z}^*, \alpha^*})$ and $\text{SKFE.s}\bar{\mathbf{k}} = \text{SKFE.KeyGen}(\text{SKFE.msk}, P_{\bar{f}, \bar{\mathbf{z}}, \bar{\alpha}})$
 else it computes $\text{SKFE.sk}^* = \text{SKFE.Sim}^{\text{SK}}(\text{SKFE.msk}, 1^{|P_{f^*, \mathbf{z}^*, \alpha^*}|})$,
 $\text{SKFE.s}\bar{\mathbf{k}} = \text{SKFE.Sim}^{\text{SK}}(\text{SKFE.msk}, 1^{|P_{\bar{f}, \bar{\mathbf{z}}, \bar{\alpha}}|})$ and sends $\{\text{SKFE.sk}^*, \text{SKFE.s}\bar{\mathbf{k}}\}$ to \mathcal{B} .
4. \mathcal{B} generates $(\text{kpABE.mpk}, \text{kpABE.msk}) \leftarrow \text{kpABE.Setup}(1^\lambda)$, sets $\text{RMFE.mpk} = \text{kpABE.mpk}$ and sends RMFE.mpk to \mathcal{A} .
5. **Key Queries:** For each key query (lb, x) , \mathcal{B} does the following:
 - Samples $K_{j,b}, R_{j,b} \leftarrow \{0, 1\}^\lambda, \forall j \in [\lambda], b \in \{0, 1\}$.
 - Computes $\text{kpABE.ct}_{\text{lb}, j, b} \leftarrow \text{kpABE.Enc}(\text{kpABE.msk}, (\text{lb}, j, b), K_{j,b})$ for $j \in [\lambda], b \in \{0, 1\}$.
 - It sends (x, K, R) as challenge message to the SKFE challenger. The challenger returns $\text{SKFE.ct}_{\hat{\beta}}$, where $\text{SKFE.ct}_0 = \text{SKFE.Enc}(\text{SKFE.msk}, (x, K, R))$ and $\text{SKFE.ct}_1 = \text{SKFE.Sim}^{\text{CT}}(\text{SKFE.msk}, \bigoplus_j R_{j, z_j^*}, P_{\bar{f}, \bar{\mathbf{z}}, \bar{\alpha}}(x, K, R))$.
 - Returns

$\text{RMFE.sk}_{\text{lb},x} = (\text{SKFE.ct}_{\beta}, \text{kpABE.mpk}, \{(\text{lb}, j, b), \text{kpABE.ct}_{\text{lb},j,b}\}_{j,b})$
to \mathcal{A} .

6. **Challenge Query:** When \mathcal{A} outputs L^* for the challenge query, \mathcal{B} does the following:

- Computes kpABE secret keys $\{\text{kpABE.sk}_{L^*,j,z_j^*}\}_{j \in [\lambda]}$.
- Constructs $\text{CC}[\text{SKFE.sk}^*, \{\text{kpABE.sk}_{L^*,j,z_j^*}\}_{j \in [\lambda]}]$ as defined in Figure 3.3 and returns $\text{RMFE.ct}^* \leftarrow \text{LO.Obf}(\text{CC}[\text{SKFE.sk}^*, \{\text{kpABE.sk}_{L^*,j,z_j^*}\}_{j \in [\lambda]}], \alpha^*)$ to \mathcal{A} .

7. **SK-Enc Query:** When the adversary outputs \bar{L} for the SK-Enc query, \mathcal{B} does the following:

- Computes kpABE secret keys $\{\text{kpABE.sk}_{\bar{L},j,\bar{z}_j}\}_{j \in [\lambda]}$.
- Constructs $\text{CC}[\text{SKFE.sk}, \{\text{kpABE.sk}_{\bar{L},j,\bar{z}_j}\}_{j \in [\lambda]}]$ and returns $\text{RMFE.ct} \leftarrow \text{LO.Obf}(\text{CC}[\text{SKFE.sk}, \{\text{kpABE.sk}_{\bar{L},j,\bar{z}_j}\}_{j \in [\lambda]}], \bar{\alpha})$ to \mathcal{A} .

8. In the end, \mathcal{A} outputs a bit β' . \mathcal{B} forwards β' as its guess bit to the SKFE challenger.

We observe that if $\hat{\beta} = 0$, then \mathcal{B} simulated Hyb_0 , else Hyb_1 with \mathcal{A} . Hence, advantage of $\mathcal{B} = |\Pr(\beta' = 1 | \hat{\beta} = 0) - \Pr(\beta' = 1 | \hat{\beta} = 1)| = |\Pr(\beta' = 1 | \text{Hyb}_0) - \Pr(\beta' = 1 | \text{Hyb}_1)| = \epsilon$ (by assumption). \blacksquare

Claim 3.26. Assume that LO is secure (Def. 3.8), then $\text{Hyb}_1 \approx_c \text{Hyb}_2$.

Proof. We show that if there exists a PPT adversary \mathcal{A} who can distinguish between Hyb_1 and Hyb_2 with non-negligible advantage ϵ , then there exists a PPT adversary \mathcal{B} against the security of LO scheme with the same advantage ϵ . The reduction is as follows

1. Upon being invoked by the LO challenger, \mathcal{B} invokes \mathcal{A} . \mathcal{A} outputs f^*, \bar{f} .
2. \mathcal{B} samples $\bar{\alpha}, \mathbf{z}^*, \bar{\mathbf{z}} \leftarrow \{0, 1\}^\lambda$. It also defines the function $P_{\bar{f}, \bar{\mathbf{z}}, \bar{\alpha}}$.
3. \mathcal{B} generates $\text{SKFE.msk} \leftarrow \text{SKFE.Setup}(1^\lambda)$, $(\text{kpABE.mpk}, \text{kpABE.msk}) \leftarrow \text{kpABE.Setup}(1^\lambda)$ and sets $\text{RMFE.mpk} = \text{kpABE.mpk}$ and $\text{RMFE.msk} = (\text{SKFE.msk}, \text{kpABE.mpk}, \text{kpABE.msk})$ and sends RMFE.mpk to \mathcal{A} .
4. **Key Queries:** For each key query (lb, x) , \mathcal{B} does the following:

- Samples $K_{j,b}, R_{j,b} \leftarrow \{0, 1\}^\lambda, \forall j \in [\lambda], b \in \{0, 1\}$.
- Computes $\text{kpABE.ct}_{\text{lb},j,b} \leftarrow \text{kpABE.Enc}(\text{kpABE.mpk}, (\text{lb}, j, b), K_{j,b})$
 $\forall j \in [\lambda], b \in \{0, 1\}$ and $\text{SKFE.ct} \leftarrow \text{SKFE.Sim}^{\text{CT}}(\text{SKFE.msk}, \bigoplus_j R_{j,z_j^*}, P_{\bar{f}, \bar{z}, \bar{\alpha}}(x, K, R))$.
- Returns $\text{RMFE.sk}_{\text{lb},x} = (\text{SKFE.ct}, \text{kpABE.mpk}, \{(\text{lb}, j, b), \text{kpABE.ct}_{\text{lb},j,b}\}_{j,b})$ to \mathcal{A} .

5. **Challenge Query:** When the adversary outputs L^* for the challenge ciphertext, \mathcal{B} does the following:

- Computes kpABE secret keys $\{\text{kpABE.sk}_{L^*,j,z_j^*}\}_{j \in [\lambda]}$.
- Constructs a function $\text{CC}[\text{SKFE.sk}^*, \{\text{kpABE.sk}_{L^*,j,z_j^*}\}_{j \in [\lambda]}]$ as defined in Figure 3.3, where $\text{SKFE.sk}^* \leftarrow \text{SKFE.Sim}^{\text{SK}}(\text{SKFE.msk}, 1^{|P_{f^*,z^*,\alpha^*}|})$ ²⁰.
- It sends $\text{CC}[\text{SKFE.sk}^*, \{\text{kpABE.sk}_{L^*,j,z_j^*}\}_{j \in [\lambda]}]$ to the LO challenger. The challenger samples a lock value $\alpha^* \leftarrow \{0, 1\}^\lambda$ and $\hat{\beta} \leftarrow \{0, 1\}$, computes and return $\text{Obf}_{\hat{\beta}}$, where $\text{Obf}_0 = \text{LO.Obf}(\text{CC}[\text{SKFE.sk}^*, \{\text{kpABE.sk}_{L^*,j,z_j^*}\}_{j \in [\lambda]}], \alpha^*)$ and $\text{Obf}_1 = \text{LO.Sim}(1^\lambda, 1^{|\text{CC}|})$.
- It sets $\text{RMFE.ct}^* = \text{Obf}_{\hat{\beta}}$ and sends it to the adversary \mathcal{A} .

6. **SK-Enc Query:** When the adversary outputs \bar{L} for the SK-Enc query, \mathcal{B} does the following:

- Computes kpABE secret keys $\{\text{kpABE.sk}_{\bar{L},j,\bar{z}_j}\}_{j \in [\lambda]}$.
- Constructs $\text{CC}[\text{SKFE.s}\bar{k}, \{\text{kpABE.sk}_{\bar{L},j,\bar{z}_j}\}_{j \in [\lambda]}]$, where $\text{SKFE.s}\bar{k} \leftarrow \text{SKFE.Sim}^{\text{SK}}(\text{SKFE.msk}, 1^{|P_{\bar{f},\bar{z},\bar{\alpha}}|})$.
- Computes $\text{RMFE.ct} \leftarrow \text{LO.Obf}(\text{CC}[\text{SKFE.s}\bar{k}, \{\text{kpABE.sk}_{\bar{L},j,\bar{z}_j}\}_{j \in [\lambda]}], \bar{\alpha})$ and returns it to the adversary \mathcal{A} .

7. In the end, \mathcal{A} outputs a bit β' . \mathcal{B} forwards β' as its guess bit to the LO challenger.

We observe that if $\hat{\beta} = 0$, then \mathcal{B} simulated Hyb_1 , else Hyb_2 with \mathcal{A} . Hence, advantage of $\mathcal{B} = |\Pr(\beta' = 1 | \hat{\beta} = 0) - \Pr(\beta' = 1 | \hat{\beta} = 1)| = |\Pr(\beta' = 1 | \text{Hyb}_1) - \Pr(\beta' = 1 | \text{Hyb}_2)| = \epsilon$ (by assumption). \blacksquare

²⁰The size of P_{f^*,z^*,α^*} is independent of any specific lock value and hence can be computed without the knowledge of α^* .

Claim 3.27. Assume that SKFE is secure (Def. 3.6), then $\text{Hyb}_2 \approx_c \text{Hyb}_3$.

The proof of this claim is similar to that of Claim 3.25, hence omitted. \blacksquare

Function Hiding

Theorem 3.28. Assume SKFE is secure (Def. 3.6), kpABE satisfies Sel-IND security (Def. 2.5). Furthermore, assume $|\mathcal{F}_\lambda| \leq \text{poly}(\lambda)$. Then, the RMFE construction satisfies 1-query function hiding as defined in Definition 3.21.

Proof. We prove the theorem via the following sequence of hybrids.

Hyb_0 : This is the real world with $\beta = 0$. We summarize the steps of the game here to set up the notations used in the later hybrids.

1. The adversary outputs the challenge query f_0, f_1, L^* and the SK-Enc query function \tilde{f} in the beginning of the game²¹. The challenger then does the following:
 - Generates $\text{SKFE.msk} \leftarrow \text{SKFE.Setup}(1^\lambda)$,
 $(\text{kpABE.mpk}, \text{kpABE.msk}) \leftarrow \text{kpABE.Setup}(1^\lambda)$, sets
 $\text{RMFE.mpk} = \text{kpABE.mpk}$ and
 $\text{RMFE.msk} = (\text{SKFE.msk}, \text{kpABE.mpk}, \text{kpABE.msk})$.
 - It also samples a tag $\mathbf{z}^* \in \{0, 1\}^\lambda$ and a lock value α^* .
 - Computes $\text{SKFE.sk}^* \leftarrow \text{SKFE.KeyGen}(\text{SKFE.msk}, P_{f_0, \mathbf{z}^*, \alpha^*})$ and $\{\text{kpABE.sk}_{L^*, j, z_j^*}\}_{j \in [\lambda]}$ as defined in the construction.
 - Returns RMFE.mpk and the challenge ciphertext $\text{RMFE.ct}^* \leftarrow \text{LO.Obf}(\text{CC}[\text{SKFE.sk}^*, \{\text{kpABE.sk}_{L^*, j, z_j^*}\}_{j \in [\lambda]}], \alpha^*)$ to \mathcal{A} .
2. **Key Queries:** For each key query (lb, x) , the challenger returns $\text{RMFE.sk}_{\text{lb}, x} = (\text{SKFE.ct}, \{(\text{lb}, j, b), \text{kpABE.ct}_{\text{lb}, j, b}\}_{j \in [\lambda], b \in \{0, 1\}})$, where $\text{SKFE.ct}, \{\text{kpABE.ct}_{\text{lb}, j, b}\}_{j \in [\lambda], b \in \{0, 1\}}$ are computed as in the construction.
3. **SK-Enc Query :** When the adversary outputs \tilde{L} for the SK-Enc query, the challenger does the following:

²¹To keep the proofs and notations simple, we let f_0, f_1 and \tilde{f} to be given selectively. This is sufficient to achieve security as in Def 3.20, as mentioned in Remark 5.

- Samples a tag $\bar{z} \leftarrow \{0, 1\}^\lambda$ and a lock value $\bar{\alpha} \leftarrow \{0, 1\}^\lambda$.
- Computes kpABE secret keys $\{\text{kpABE.sk}_{\bar{L},j,\bar{z}_j}\}_{j \in [\lambda]}$ and a SKFE secret key $\text{SKFE.s}\bar{\text{k}}$ as in the construction.
- Constructs $\text{CC}[\text{SKFE.s}\bar{\text{k}}, \{\text{kpABE.sk}_{\bar{L},j,\bar{z}_j}\}_{j \in [\lambda]}]$ and returns $\text{RMFE.ct} \leftarrow \text{LO.Obf}(\text{CC}[\text{SKFE.s}\bar{\text{k}}, \{\text{kpABE.sk}_{\bar{L},j,\bar{z}_j}\}_{j \in [\lambda]}], \bar{\alpha})$ to \mathcal{A} .

4. In the end, \mathcal{A} outputs a bit β' .

Hyb₁ : This hybrid is same as the previous hybrid except the following:

1. The challenger samples $\alpha^*, \bar{\alpha}, \mathbf{z}^*, \bar{\mathbf{z}}$ and defines the functions $P_{f_0, \mathbf{z}^*, \alpha^*}$ and $P_{\bar{f}, \bar{\mathbf{z}}, \bar{\alpha}}$ in the beginning of the game.
2. The SKFE ciphertexts and keys are computed using SKFE simulators as:
 - For each key query (lb, x) , the SKFE ciphertext in $\text{RMFE.sk}_{\text{lb},x}$ is computed as $\text{SKFE.Sim}^{\text{CT}}(\text{SKFE.msk}, P_{f_0, \mathbf{z}^*, \alpha^*}(x, K, R), P_{\bar{f}, \bar{\mathbf{z}}, \bar{\alpha}}(x, K, R))$.
 - The SKFE secret keys SKFE.sk^* and $\text{SKFE.s}\bar{\text{k}}$ in RMFE.ct^* and RMFE.ct are computed as $\text{SKFE.Sim}^{\text{SK}}(\text{SKFE.msk}, 1^{|P_{f_0, \mathbf{z}^*, \alpha^*}|})$ and $\text{SKFE.Sim}^{\text{SK}}(\text{SKFE.msk}, 1^{|P_{\bar{f}, \bar{\mathbf{z}}, \bar{\alpha}}|})$, respectively.

Hyb₂ : In this hybrid, for each key query (lb, x) with $\text{lb} \in L^*$, for all $j \in [\lambda]$, $\text{kpABE.ct}_{\text{lb},j,1-\bar{z}_j}$ in $\text{RMFE.sk}_{\text{lb},x}$ is computed as $\text{kpABE.Enc}(\text{kpABE.mpk}, (\text{lb}, j, 1 - \bar{z}_j), 0^\lambda)$.

Hyb₃ : In this hybrid, for each key query (lb, x) with $\text{lb} \in L^*$, K -values are chosen differently as follows: let $i \in [\lambda]$ be the first position where $z_i^* \neq \bar{z}_i$. If no such i exists, then the challenger aborts the game. Else, it samples $\{K_{j,b}\}_{j \in [\lambda] \setminus \{i\}, b \in \{0,1\}}$, K_{i,\bar{z}_i} , $\{R_{j,b}\}_{j \in [\lambda], b \in \{0,1\}}$ uniformly randomly as in the original game. It then sets $K_{i,1-\bar{z}_i} = \bigoplus_{j \in [\lambda] \setminus \{i\}} K_{j,z_j^*} \oplus \alpha^* \oplus \bigoplus_{j \in [\lambda]} R_{j,z_j^*}$.

Hyb₄ : This hybrid is same as the previous hybrid, except that the SKFE ciphertexts

and keys in RMFE secret keys and ciphertexts are now generated with function f_1 in place of f_0 as follows:

- For each key query (lb, x) , the SKFE ciphertext in $\text{RMFE.sk}_{\text{lb},x}$ is computed as $\text{SKFE.Sim}^{\text{CT}}(\text{SKFE.msk}, P_{f_1, \mathbf{z}^*, \alpha^*}(x, K, R), P_{\bar{f}, \bar{\mathbf{z}}, \bar{\alpha}}(x, K, R))$.
- The SKFE secret key SKFE.sk^* in RMFE.ct^* is computed as $\text{SKFE.Sim}^{\text{SK}}(\text{SKFE.msk}, 1^{|P_{f_1, \mathbf{z}^*, \alpha^*}|})$.

Now, from here we undo the changes made in the previous hybrids.

Hyb₅ : In this hybrid, for each key query (lb, x) with $\text{lb} \in L^*$, $K_{i, 1-\bar{z}_i}$ (where i is the first position where \mathbf{z}^* and $\bar{\mathbf{z}}$ differ) is also sampled randomly.

Hyb₆ : In this hybrid, for each such key query (lb, x) with $\text{lb} \in L^*$, $\text{kpABE.Enc}(\text{kpABE.mpk}, (\text{lb}, j, 1 - \bar{z}_j), 0^\lambda)$ is changed back to $\text{kpABE.Enc}(\text{kpABE.mpk}, (\text{lb}, j, 1 - \bar{z}_j), K_{j, 1-\bar{z}_j})$ in $\text{RMFE.sk}_{\text{lb},x}$.

Hyb₇ : In this hybrid SKFE ciphertexts in RMFE secret keys and SKFE secret keys in RMFE ciphertexts are computed as in the real world. That is,

- For each key query (lb, x) , $\text{RMFE.sk}_{\text{lb},x} = (\text{SKFE.Enc}(\text{SKFE.msk}, K, R), \text{kpABE.mpk}, \{(\text{lb}, j, b), \text{kpABE.Enc}(\text{kpABE.mpk}, (\text{lb}, j, b), K_{j,b})\}_{j \in [\lambda], b \in \{0,1\}})$.
- $\text{RMFE.ct}^* = \text{LO.Obf}(\text{CC}[\text{SKFE.sk}^*, \{\text{kpABE.sk}_{L^*, j, z_j^*}\}_{j \in [\lambda]}], \alpha^*)$,
 $\text{RMFE.ct} = \text{LO.Obf}(\text{CC}[\text{SKFE.sk}, \{\text{kpABE.sk}_{L, j, \bar{z}_j}\}_{j \in [\lambda]}], \bar{\alpha})$, where
 $\text{SKFE.sk}^* = \text{SKFE.KeyGen}(\text{SKFE.msk}, P_{f_1, \mathbf{z}^*, \alpha^*})$ and
 $\text{SKFE.sk} = \text{SKFE.KeyGen}(\text{SKFE.msk}, P_{\bar{f}, \bar{\mathbf{z}}, \bar{\alpha}})$

Note that this hybrid corresponds to the real world for $\beta = 1$.

Indistinguishability of the hybrids. Now we show that the consecutive hybrids are computationally indistinguishable for any PPT adversary \mathcal{A} .

Claim 3.29. Assume that SKFE is secure (Def. 3.6). Then, $\text{Hyb}_0 \approx_c \text{Hyb}_1$.

Proof. The proof follows similar steps as the proof for Claim 3.25, hence omitted. ■

Claim 3.30. Assume that kpABE is selectively secure. Then, $\text{Hyb}_1 \approx_c \text{Hyb}_2$.

Proof. To prove the claim we consider the following sub hybrids between Hyb_1 and Hyb_2 . Let $L^* = \{\text{lb}_1, \dots, \text{lb}_{|L^*|}\}$ with some fixed ordering between the labels in L^* . Let $L_{[1:k]} \subseteq L^*$ denote the set of first k labels in L^* , i.e. $L_{[1:k]}^* = \{\text{lb}_1, \dots, \text{lb}_k\}$. Then for all $1 \leq i \leq |L^*|$ and $0 \leq \tau \leq \lambda$, define $\text{Hyb}_{1,(i,\tau)}$: which is same as Hyb_1 except that for any key query (lb, x) , $\{\text{kpABE.ct}_{\text{lb},j,b}\}_{j \in [\lambda], b \in \{0,1\}}$ is computed differently as follows:

$$\text{kpABE.ct}_{\text{lb},j,b} = \text{kpABE.Enc}(\text{kpABE.mpk}, (\text{lb}, j, b), W),$$

where

$$W = \begin{cases} K_{i,b} & \text{if } (b = \bar{z}_j) \vee (\text{lb} \notin L_{[1:i]}) \vee (\text{lb} = \text{lb}_i \wedge j > \tau), \\ 0^\lambda & \text{otherwise, i.e. } (b = 1 - \bar{z}_j) \wedge (\text{lb} \in L_{[1:i-1]}) \vee (\text{lb} = \text{lb}_i \wedge j \leq \tau). \end{cases}$$

Then, we observe that $\text{Hyb}_{1,(1,0)} = \text{Hyb}_1$, $\text{Hyb}_{1,(|L^*|,\lambda)} = \text{Hyb}_2$ and $\text{Hyb}_{1,(i-1,\lambda)} = \text{Hyb}_{1,(i,0)}$.

Hence, all we need to show is that for all $i \in [|L^*|]$, $\tau \in [\lambda]$,

$$\text{Hyb}_{1,(i,\tau-1)} \approx_c \text{Hyb}_{1,(i,\tau)}.$$

This follows from the **Sel-IND** security of kpABE . In particular, if there is no key query issued for lb_i , then the hybrids are identical. On the other hand, if there is a key query (lb, x) , such that $\text{lb} = \text{lb}_i$, then we show that if \mathcal{A} can distinguish between the two hybrids with non-negligible advantage ϵ then we can design a PPT algorithm \mathcal{B} with the same advantage ϵ against **Sel-IND** security of kpABE . \mathcal{B} is defined as follows:

1. Upon being invoked by the kpABE challenger, \mathcal{B} invokes \mathcal{A} which outputs f_0, f_1, L^*, \bar{f} . Let $L^* = \{\text{lb}_1, \dots, \text{lb}_{|L^*|}\}$.
2. \mathcal{B} sends $(\text{lb}_i, \tau, 1 - \bar{z}_\tau)$ as challenge attribute to the kpABE challenger. The kpABE challenger samples $\hat{\beta} \leftarrow \{0, 1\}$ and $(\text{kpABE.mpk}, \text{kpABE.msk}) \leftarrow \text{kpABE.Setup}(1^\lambda)$ and sends kpABE.mpk to \mathcal{B} .

3. \mathcal{B} generates $\text{SKFE.msk} \leftarrow \text{SKFE.Setup}(1^\lambda)$ and sets $\text{RMFE.mpk} = \text{kpABE.mpk}$. It also samples $\alpha^*, \bar{\alpha}, \mathbf{z}^*, \bar{\mathbf{z}}$ and defines $P_{f_0, \mathbf{z}^*, \alpha^*}$ and $P_{\bar{f}, \bar{\mathbf{z}}, \bar{\alpha}}$. It then does the following:

- For each $j \in [\lambda]$, defines the circuit $C_{L^*, j, \mathbf{z}_j^*}$ and sends a key query for $C_{L^*, j, \mathbf{z}_j^*}$ to the kpABE challenger. The kpABE challenger returns $\text{kpABE.sk}_{L^*, j, \mathbf{z}_j^*}$.
- Computes $\text{SKFE.sk}^* \leftarrow \text{SKFE.Sim}^{\text{SK}}(\text{SKFE.msk}, 1^{|P_{f_0, \mathbf{z}^*, \alpha^*}|})$.
- Computes $\text{RMFE.ct}^* \leftarrow \text{LO.Obf}(\text{CC}[\text{SKFE.sk}^*, \{\text{kpABE.sk}_{L^*, j, \mathbf{z}_j^*}\}_{j \in [\lambda]}])$.
- Returns RMFE.mpk and RMFE.ct^* to \mathcal{A} .

4. **Key Queries:** For each key query (lb, x) , \mathcal{B} does the following:

- Computes $\text{SKFE.ct} \leftarrow \text{SKFE.Sim}^{\text{CT}}(\text{SKFE.msk}, P_{f_0, \mathbf{z}^*, \alpha^*}(x, K, R), P_{\bar{f}, \bar{\mathbf{z}}, \bar{\alpha}}(x, K, R))$.
- To compute kpABE ciphertext, $\text{kpABE.sk}_{\text{lb}, j, b}$ for $j \in [\lambda], b \in \{0, 1\}$,
 - if $(\text{lb} = \text{lb}_i) \wedge j = \tau \wedge b = 1 - \bar{\mathbf{z}}_j^{22}$, then \mathcal{B} sends challenge query with messages $\mu_0 = K_{j, 1 - \bar{\mathbf{z}}_j}$ and $\mu_1 = 0^\lambda$. The kpABE challenger returns $\text{kpABE.ct} = \text{kpABE.Enc}(\text{kpABE.mpk}, (\text{lb}_i, \tau, 1 - \bar{\mathbf{z}}_\tau), \mu_\beta)$, which \mathcal{B} sets as $\text{kpABE.ct}_{\text{lb}_i, \tau, 1 - \bar{\mathbf{z}}_\tau}$.
 - else, \mathcal{B} computes $\text{kpABE.ct}_{\text{lb}, j, b}$ on its own using kpABE.mpk as defined for $\text{Hyb}_{1, (i, \tau - 1)}$ (same for $\text{Hyb}_{1, (i, \tau)}$).
- Returns $\text{RMFE.sk}_{\text{lb}, x} = (\text{SKFE.ct}, \text{kpABE.mpk}, \{(\text{lb}, j, b), \text{kpABE.ct}_{\text{lb}, j, b}\}_{j, b})$ to \mathcal{A} .

5. **SK-Enc Query:** When \mathcal{A} outputs \bar{L} as part of SK-Enc query, \mathcal{B} does the following:

- Computes $\text{SKFE.s}\bar{\mathbf{k}} \leftarrow \text{SKFE.Sim}^{\text{SK}}(\text{SKFE.msk}, 1^{|P_{\bar{f}, \bar{\mathbf{z}}, \bar{\alpha}}|})$.
- For each $j \in [\lambda]$, defines circuit $C_{\bar{L}, j, \bar{\mathbf{z}}_j}$ and sends a kpABE key query for $C_{\bar{L}, j, \bar{\mathbf{z}}_j}$. The kpABE challenger returns $\text{kpABE.sk}_{\bar{L}, j, \bar{\mathbf{z}}_j}$.
- Computes $\text{RMFE.c}\bar{\mathbf{t}} \leftarrow \text{LO.Obf}(\text{CC}[\text{SKFE.s}\bar{\mathbf{k}}, \{\text{kpABE.sk}_{\bar{L}, j, \bar{\mathbf{z}}_j}\}_{j \in [\lambda]}])$.

²²If there are more than one key queries with $\text{lb} = \text{lb}_i$, then we use multi-challenge version of Sel-IND security of kpABE which can easily be shown equivalent to the one defined in Definition 2.5.

- Returns $\text{RMFE}.\overline{\text{ct}}$ to \mathcal{A} .

6. In the end, \mathcal{A} outputs a bit β' . \mathcal{B} forwards β' as its guess bit to the kpABE challenger.

We observe that if $\hat{\beta} = 0$, then \mathcal{B} simulated $\text{Hyb}_{1,(i,\tau-1)}$, else $\text{Hyb}_{1,(i,\tau)}$ with \mathcal{A} . Hence, advantage of \mathcal{B} = $|\Pr(\beta' = 1 | \hat{\beta} = 0) - \Pr(\beta' = 1 | \hat{\beta} = 1)| = |\Pr(\beta' = 1 | \text{Hyb}_{1,(i,\tau-1)}) - \Pr(\beta' = 1 | \text{Hyb}_{1,(i,\tau)})| = \epsilon$ (by assumption).

Admissibility of \mathcal{B} : Firstly, we observe that \mathcal{B} issues key queries for only the following set of circuits: $\{C_{L^*,j,z_j^*}\}_{j \in [\lambda]}$, $\{C_{\bar{L},j,\bar{z}_j}\}_{j \in [\lambda]}$ and the challenge attribute is $(\text{lb}_i, \tau, 1 - \bar{z}_\tau)$, where $\text{lb}_i \in L^*$. Next, we note that

- $C_{L^*,j,z_j^*}(\text{lb}_i, \tau, 1 - \bar{z}_\tau) = 0$ for all $j \in [\lambda]$, because $\text{lb}_i \in L^*$.
- $C_{\bar{L},j,\bar{z}_j}(\text{lb}_i, \tau, 1 - \bar{z}_\tau) = 0$ for all $j \neq \tau$, and $C_{\bar{L},\tau,\bar{z}_\tau}(\text{lb}_i, \tau, 1 - \bar{z}_\tau) = 0$, because $\bar{z}_\tau \neq 1 - \bar{z}_\tau$.

This establishes the admissibility of \mathcal{B} . ■

Claim 3.31. Hyb_2 and Hyb_3 are statistically indistinguishable.

Proof. Firstly, we observe that $\Pr[\mathbf{z}^* = \bar{\mathbf{z}}] = 1/2^\lambda$. Hence, with probability $1 - 1/2^\lambda$, the challenger does not abort in Hyb_3 . Next we show that if the game is not aborted, the two hybrids are statistically indistinguishable in the view of the adversary. In this case, the only difference between the two hybrids is the following: for each key query (lb, x) with $\text{lb} \in L^*$, the value of $K_{i,1-\bar{z}_i}$ (i.e. K_{i,z_i^*}), where i is the first index such that $z_i^* \neq \bar{z}_i$, are computed differently. In Hyb_2 , K_{i,z_i^*} is sampled uniformly, while in Hyb_3 , it is computed as $K_{i,z_i^*} = \bigoplus_{j \in [\lambda] \setminus \{i\}} (K_{j,z_j^*} \oplus R_{j,z_j^*}) \oplus R_{i,z_i^*} \oplus \alpha^*$. However, note that if $f_0(x) = 1$, then K_{i,z_i^*} is not used in anywhere because of the change we introduced in Hyb_2 and hence does not affect the adversary's view. On the other hand, if $f_0(x) = 0$, R_{i,z_i^*} is not used anywhere, and hence K_{i,z_i^*} is uniformly random in the adversary's view because of the randomness of R_{i,z_i^*} . ■

Claim 3.32. Hyb_3 and Hyb_4 are identical in the view of the adversary.

Proof. Observe that the two hybrids differ only in the computation (simulation) of SKFE ciphertexts and secret keys. In particular,

- $\text{SKFE.Sim}^{\text{SK}}(\text{SKFE.msk}, 1^{|P_{f_0, \mathbf{z}^*, \alpha^*}|})$ is changed to $\text{SKFE.Sim}^{\text{SK}}(\text{SKFE.msk}, 1^{|P_{f_1, \mathbf{z}^*, \alpha^*}|})$ in the computation of RMFE.ct^* . This is just a conceptual change and both the computations are exactly the same, since $|P_{f_0, \mathbf{z}^*, \alpha^*}| = |P_{f_1, \mathbf{z}^*, \alpha^*}|$.²³
- For each key query (lb, x) , $\text{SKFE.Sim}^{\text{CT}}(\text{SKFE.msk}, P_{f_0, \mathbf{z}^*, \alpha^*}(x, K, R), P_{\tilde{f}, \tilde{\mathbf{z}}, \tilde{\alpha}}(x, K, R))$ is changed to $\text{SKFE.Sim}^{\text{CT}}(\text{SKFE.msk}, P_{f_1, \mathbf{z}^*, \alpha^*}(x, K, R), P_{\tilde{f}, \tilde{\mathbf{z}}, \tilde{\alpha}}(x, K, R))$. The change is again only conceptual and does not affect the actual computation as we argue below:

- For $f_0(x) = f_1(x)$: there is no change.
- For $f_0(x) \neq f_1(x)$: let $f_0(x) = 1$ and $f_1(x) = 0$. Then, $\text{lb} \in L^*$ and thus

$$\begin{aligned}
 P_{f_0, \mathbf{z}^*, \alpha^*}(x, K, R) &= \bigoplus_{j \in [\lambda]} R_{j, z_j^*} \\
 P_{f_1, \mathbf{z}^*, \alpha^*}(x, K, R) &= \bigoplus_{j \in [\lambda] \setminus \{i\}} K_{j, z_j^*} \oplus \alpha^* \oplus K_{i, z_i^*} \\
 &= \bigoplus_{j \in [\lambda] \setminus \{i\}} K_{j, z_j^*} \oplus \alpha^* \oplus \left(\bigoplus_{j \in [\lambda] \setminus \{i\}} (K_{j, z_j^*} \oplus R_{j, z_j^*}) \oplus R_{i, z_i^*} \right) \oplus \alpha^* \\
 &= \bigoplus_{j \in [\lambda]} R_{j, z_j^*}.
 \end{aligned}$$

- The same argument works for $f_0(x) = 0$ and $f_1(x) = 1$.

■

Indistinguishability between the rest of the hybrids can be argued in the same way as their counterparts in the previous set of hybrids. In particular, proofs for indistinguishability between Hyb_4 and Hyb_5 is same as the proof of claim 3.31, Hyb_5 and Hyb_6 is same as the proof for claim 3.30 and Hyb_6 and Hyb_7 is same as the proof for claim 3.29. ■

²³we can use padding to make the sizes equal.

3.7 SECRET KEY RPE FROM EVASIVE AND TENSOR LWE

In this section we construct a secret-key RPE scheme from evasive and tensor LWE, followed by the efficiency and security analysis of our scheme.

3.7.1 Construction

We give a construction of the secret-key RPE scheme $\text{RPE} = (\text{RPE.Setup}, \text{RPE.KeyGen}, \text{RPE.Broadcast}, \text{RPE.Enc}, \text{RPE.Dec})$ for an attribute space $\mathcal{X} = \{\mathcal{X}_\lambda\}_\lambda$, a function family $\mathcal{F} = \{\mathcal{F}_\lambda\}_\lambda$ where $\mathcal{F}_\lambda = \{f : \mathcal{X}_\lambda \rightarrow \{0, 1\}\}$, a label space $\mathcal{L} = \{\mathcal{L}_\lambda\}_\lambda$ and a message space $\mathcal{M} = \{\mathcal{M}_\lambda\}_\lambda$. We assume that the size of $|\mathcal{F}_\lambda|$ is bounded by some polynomial in λ , which will suffice for our purpose. Our scheme uses the following building blocks:

1. An RMFE scheme $\text{RMFE} = (\text{RMFE.Setup}, \text{RMFE.KeyGen}, \text{RMFE.PK-Enc}, \text{RMFE.SK-Enc}, \text{RMFE.Dec})$ with attribute space \mathcal{X} , label space \mathcal{L} and function family \mathcal{F} . We instantiate it by our construction in Sec. 3.6.2 based on LWE.
2. A CP-ABE scheme $\text{cpABE} = (\text{cpABE.Setup}, \text{cpABE.Enc}, \text{cpABE.KeyGen}, \text{cpABE.Dec})$ with message space \mathcal{M} satisfying VerSel-IND security (Def. 2.8) that supports the circuit class $\mathcal{C}_{\ell(\lambda), d(\lambda)}$ and the efficiency property listed in Theorem 3.4. We set $\ell(\lambda) = |\text{RMFE.sk}_{\text{lb}, x}|$ and $d(\lambda)$ to be the upper bound on the depth of the circuit $C_{L, \text{RMFE.ct}}$ in Eq. (3.4). Looking ahead, we show that the depth is bounded by some polynomial $\text{poly}(\lambda, |f|, |x|, |\text{lb}|)$. We choose d so that it is larger than the value. We can instantiate the scheme by the one proposed by [169] based on tensor and evasive LWE.

We describe our construction below.

$\text{RPE.Setup}(1^\lambda) \rightarrow (\text{RPE.mpk}, \text{RPE.msk})$. The setup algorithm takes as input the security parameter λ and does the following:

- Generates $(\text{RMFE.mpk}, \text{RMFE.msk}) \leftarrow \text{RMFE.Setup}(1^\lambda)$ and $(\text{cpABE.mpk}, \text{cpABE.msk}) \leftarrow \text{cpABE.Setup}(1^\lambda)$.
- Outputs $\text{RPE.mpk} = (\text{RMFE.mpk}, \text{cpABE.mpk})$ and $\text{RPE.msk} = (\text{RMFE.msk}, \text{cpABE.msk})$.

$\text{RPE.KeyGen}(\text{RPE.msk}, \text{lb}, x) \rightarrow \text{RPE.sk}_{\text{lb}, x}$. The key generation algorithm takes as input the master secret key RPE.msk , a label $\text{lb} \in \mathcal{L}$ and an attribute $x \in \mathcal{X}$

and does the following:

- Parse $\text{RPE.msk} = (\text{RMFE.msk}, \text{cpABE.msk})$.
- Compute $\text{RMFE.sk}_{\text{lb},x} \leftarrow \text{RMFE.KeyGen}(\text{RMFE.msk}, \text{lb}, x)$.
- Set $\text{att} = ((\text{lb}, x), \text{RMFE.sk}_{\text{lb},x})$.
- Compute $\text{cpABE.sk}_{\text{lb},x} \leftarrow \text{cpABE.KeyGen}(\text{cpABE.msk}, \text{att})$.
- Output $\text{RPE.sk}_{\text{lb},x} = (\text{att}, \text{cpABE.sk}_{\text{lb},x})$.

$\text{RPE.Broadcast}(\text{RPE.mpk}, m, L) \rightarrow \text{RPE.ct}$. The broadcast algorithm takes as input the master public key RPE.mpk , a message m , and a revocation list $L \subseteq \mathcal{L}$ and does the following:

- Parse $\text{RPE.mpk} = (\text{RMFE.mpk}, \text{cpABE.mpk})$.
- Compute $\text{RMFE.ct} \leftarrow \text{RMFE.PK-Enc}(\text{RMFE.mpk}, L)$.
- Construct circuit $C_{L,\text{RMFE.ct}}$, with L and RMFE.ct hardwired, defined as follows:
On input (lb, x) and $\text{RMFE.sk}_{\text{lb},x}$,

$$\begin{aligned} &C_{L,\text{RMFE.ct}}((\text{lb}, x), \text{RMFE.sk}_{\text{lb},x}) \\ &= (\text{lb} \notin L) \wedge (\text{RMFE.Dec}(\text{RMFE.sk}_{\text{lb},x}, \text{RMFE.ct}, L) = 1). \end{aligned} \tag{3.4}$$

- Compute $\text{cpABE.ct} \leftarrow \text{cpABE.Enc}(\text{cpABE.mpk}, C_{L,\text{RMFE.ct}}, m)$.
- Output $\text{RPE.ct} = (\text{RMFE.ct}, \text{cpABE.ct})$.

$\text{RPE.Enc}(\text{RPE.msk}, f, m, L) \rightarrow \text{RPE.ct}$. The encryption algorithm takes as input the master secret key, a function f , a message m , and a revocation list $L \subseteq \mathcal{L}$ and does the following:

- Parse $\text{RPE.msk} = (\text{RMFE.msk}, \text{cpABE.msk})$.
- Compute $\text{RMFE.ct} \leftarrow \text{RMFE.SK-Enc}(\text{RMFE.msk}, f, L)$.

- Construct circuit $C_{L, \text{RMFE.ct}}$ from L and RMFE.ct as defined in Eq. (3.4).
- Compute $\text{cpABE.ct} \leftarrow \text{cpABE.Enc}(\text{cpABE.mpk}, C_{L, \text{RMFE.ct}}, m)$.
- Output $\text{RPE.ct} = (\text{RMFE.ct}, \text{cpABE.ct})$.

$\text{RPE.Dec}(\text{RPE.sk}_{\text{lb}, x}, \text{RPE.ct}, L) \rightarrow m$. The decryption algorithm takes as input the secret key $\text{RPE.sk}_{\text{lb}, x}$, a ciphertext RPE.ct , and a revocation list L and does the following:

- Parse $\text{RPE.sk}_{\text{lb}, x} = (\text{att} = ((\text{lb}, x), \text{RMFE.sk}_{\text{lb}, x}), \text{cpABE.sk}_{\text{lb}, x})$ and $\text{RPE.ct} = (\text{RMFE.ct}, \text{cpABE.ct})$.
- Construct circuit $C_{L, \text{RMFE.ct}}$ from L and RMFE.ct as defined in Eq. (3.4).
- Compute $\text{cpABE.Dec}(\text{cpABE.mpk}, \text{cpABE.sk}_{\text{lb}, x}, \text{att}, \text{cpABE.ct}, C_{L, \text{RMFE.ct}})$ and output m .

Correctness. First, we show that $C_{L, \text{RMFE.ct}} \in \mathcal{C}_{\ell, d}$. In particular, it suffices to bound the depth of the circuit d by some fixed polynomial $\text{poly}(\lambda)$. We first observe that checking whether $\text{lb} \in L$ or not can be done with depth $|\log(|\text{lb}| \cdot |L|)| = \log(\text{poly}(\lambda)) \leq \lambda$. We also have that the depth of RMFE.Dec is bounded by $\text{poly}(\lambda, |f|, |x|, |\text{lb}|)$ as we saw in Sec. 3.6.2. Therefore, the total depth is bounded by $\text{poly}(\lambda, |f|, |x|, |\text{lb}|)$.

Theorem 3.33. *Suppose RMFE and cpABE are correct, then the above construction of secret-key RPE satisfies correctness (Def. 3.12).*

Proof. For any (lb, x) , function $f \in \mathcal{F}_\lambda$ and a revocation list $L \subseteq \mathcal{L}_\lambda$ such that $f(x) = 1, \forall x \in \mathcal{X}_\lambda$ and $\text{lb} \notin L$, consider the following two cases:

1. **Broadcast Correctness:** For any ciphertext $\text{RPE.ct} \leftarrow \text{RPE.Broadcast}(\text{RPE.mpk}, m, L)$, we have $\text{RPE.ct} = (\text{RMFE.ct}, \text{cpABE.ct})$, where $\text{RMFE.ct} \leftarrow \text{RMFE.PK-Enc}(\text{RMFE.mpk}, L)$. So, by the public encryption correctness of RMFE scheme, with all but negligible probability

$$\text{RMFE.Dec}(\text{RMFE.sk}_{\text{lb}, x}, \text{RMFE.ct}, L) = 1.$$

Since $\text{lb} \notin L$, we have $C_{L, \text{RMFE.ct}}(\text{att}) = 1$, where $\text{att} = ((\text{lb}, x), \text{RMFE.sk}_{\text{lb}, x})$. Hence by the correctness of cpABE scheme we have that with all but negligible

probability

$$\text{cpABE.Dec}(\text{cpABE.mpk}, \text{cpABE.sk}_{\text{lb},x}, \text{att}, \text{cpABE.ct}, C_{L,\text{RMFE.ct}}) = m.$$

2. **Encryption Correctness:** For any ciphertext $\text{RPE.ct} \leftarrow \text{RPE.Enc}(\text{RPE.msk}, f, m, L)$, we have $\text{RPE.ct} = (\text{RMFE.ct}, \text{cpABE.ct})$, where $\text{RMFE.ct} \leftarrow \text{RMFE.SK-Enc}(\text{RMFE.msk}, f, L)$. If $(\text{lb} \notin L) \wedge f(x) = 1$, then by the correctness of RMFE.SK-Enc algorithm, we have with all but negligible probability

$$\text{RMFE.Dec}(\text{RMFE.sk}_{\text{lb},x}, \text{RMFE.ct}, L) = 1.$$

Furthermore, $\text{lb} \notin L$ implies $C_{L,\text{RMFE.ct}}((\text{lb}, x), \text{RMFE.sk}_{\text{lb},x}) = 1$. Hence by the correctness of cpABE scheme, we have that with all but negligible probability

$$\text{cpABE.Dec}(\text{cpABE.mpk}, \text{cpABE.sk}_{\text{lb},x}, \text{att}, \text{cpABE.ct}, C_{L,\text{RMFE.ct}}) = m.$$

Hence the above construction of secret-key RPE satisfies correctness. \blacksquare

Efficiency. Here we argue that our construction achieves optimal parameters. Namely, we show that all the parameters are independent from $|L|$.

1. **Public key size $|\text{RPE.mpk}|$:** We have $|\text{RPE.mpk}| = |\text{cpABE.mpk}| + |\text{RMFE.mpk}|$. The former is bounded by $\text{poly}(\lambda, \ell, d) = \text{poly}(\lambda, |f|, |x|, |\text{lb}|)$ using Theorem 3.4, where we additionally used $\ell = |\text{RMFE.sk}_{\text{lb},x}| = \text{poly}(\lambda, |f|, |x|, |\text{lb}|)$ and $d = \text{poly}(\lambda, |f|, |x|, |\text{lb}|)$.
2. **Secret key size $|\text{RPE.sk}_{\text{lb},x}|$:** We have $|\text{RPE.sk}_{\text{lb},x}| = |\text{att}| + |\text{cpABE.sk}_{\text{lb},x}|$. We have $|\text{att}| = |\text{lb}| + |x| + |\text{RMFE.sk}_{\text{lb},x}| = \text{poly}(\lambda, |f|, |x|, |\text{lb}|)$ and $|\text{cpABE.sk}_{\text{lb},x}| = \text{poly}(\ell, d) = \text{poly}(\lambda, |f|, |x|, |\text{lb}|)$ by Theorem 3.4. Therefore, the overall size is $\text{poly}(\lambda, |f|, |x|, |\text{lb}|)$.
3. **Ciphertext size $|\text{RPE.ct}|$:** We have $|\text{RPE.ct}| = |\text{RMFE.ct}| + |\text{cpABE.ct}|$. The former is bounded by $\text{poly}(\lambda, |f|, |x|, |\text{lb}|)$ and the latter is bounded by $\text{poly}(\lambda, d) = \text{poly}(\lambda, |f|, |x|, |\text{lb}|)$ by Theorem 3.4. Therefore, the overall size is $\text{poly}(\lambda, |f|, |x|, |\text{lb}|)$.

3.7.2 Security

In this section we show that our construction of secret-key RPE is secure.

Message Hiding Security

Theorem 3.34. *Assume that cpABE is very selectively secure (Def. 2.8), RMFE is correct and $|\mathcal{F}| \leq \text{poly}(\lambda)$. Then RPE scheme satisfies 1-query very selective message hiding security (Def. 3.15).*

Proof. Recall that in the message hiding security game, we want

$$\text{RPE.Enc}(\text{RPE.msk}, f, m_0, L) \approx_c \text{RPE.Enc}(\text{RPE.msk}, f, m_1, L),$$

where for all the key queries (lb, x) to the $\text{RPE.KeyGen}(\text{RPE.msk}, \cdot, \cdot)$ oracle, either $f(x) = 0$ or $\text{lb} \in L$. We show that if there exists an adversary \mathcal{A} who has non-negligible advantage ϵ in the selective message hiding security game, then there exists a PPT adversary \mathcal{B} against the security of cpABE scheme with the same advantage ϵ . The reduction is as follows.

1. \mathcal{B} first runs \mathcal{A} . \mathcal{A} outputs the key queries $\{(\text{lb}_1, x_1), (\text{lb}_2, x_2), \dots, (\text{lb}_Q, x_Q)\}$, f and L ²⁴.
2. \mathcal{B} generates $(\text{RMFE.mpk}, \text{RMFE.msk}) \leftarrow \text{RMFE.Setup}(1^\lambda)$ and computes $\text{RMFE.ct} \leftarrow \text{RMFE.SK-Enc}(\text{RMFE.msk}, f, L)$.
3. \mathcal{B} computes $\text{RMFE.sk}_{\text{lb}_i, x_i} \leftarrow \text{RMFE.KeyGen}(\text{RMFE.msk}, \text{lb}_i, x_i)$ for all $i \in [Q]$. It also constructs the circuit $C_{L, \text{RMFE.ct}}$ as defined in the construction and sends $(C_{L, \text{RMFE.ct}}, \{(\text{lb}_i, x_i), \text{RMFE.sk}_{\text{lb}_i, x_i}\}_{i \in [Q]})$ as the challenge function and key attributes to the cpABE challenger.
 The cpABE challenger generates $(\text{cpABE.mpk}, \text{cpABE.msk}) \leftarrow \text{cpABE.Setup}(1^\lambda)$ and keys $\text{cpABE.sk}_{\text{lb}_i, x_i} \leftarrow \text{cpABE.KeyGen}(\text{cpABE.msk}, (\text{lb}_i, x_i), \text{RMFE.sk}_{\text{lb}_i, x_i})$ and returns $(\text{cpABE.mpk}, \{\text{cpABE.sk}_{\text{lb}_i, x_i}\}_{i \in [Q]})$ to \mathcal{B} . \mathcal{B} sets $\text{RPE.mpk} = (\text{RMFE.mpk}, \text{cpABE.mpk})$ and sends $(\text{RPE.mpk}, \{\text{cpABE.sk}_{\text{lb}_i, x_i}\}_{i \in [Q]})$ to \mathcal{A} .
4. **Challenge Query:** When \mathcal{A} sends the challenge messages (m_0, m_1) , \mathcal{B} forwards it to the cpABE challenger. The cpABE challenger samples a bit $\hat{\beta} \leftarrow \{0, 1\}$ and returns $\text{cpABE.ct}_{\hat{\beta}} \leftarrow \text{cpABE.Enc}(\text{cpABE.mpk}, C_{L, \text{RMFE.ct}}, m_{\hat{\beta}})$ to \mathcal{B} . \mathcal{B} sends $\text{RPE.ct} = (\text{RMFE.ct}, \text{cpABE.ct}_{\hat{\beta}})$ to \mathcal{A} .

²⁴To keep the proofs simple, we let f to be given selectively. This is sufficient to achieve security as in Def 3.14, as mentioned in Remark 4.

5. **Encryption Query:** When \mathcal{A} makes the encryption query $(\bar{f}, \bar{m}, \bar{L})$, \mathcal{B} does the following:

- Computes $\text{RMFE.ct} \leftarrow \text{RMFE.SK-Enc}(\text{RMFE.msk}, \bar{f}, \bar{L})$ and constructs the circuit $C_{\bar{L}, \text{RMFE.ct}}$ as defined in the construction.
- Computes $\text{cpABE.ct} \leftarrow \text{cpABE.Enc}(\text{cpABE.mpk}, C_{\bar{L}, \text{RMFE.ct}}, \bar{m})$.
- Returns $\text{RPE.ct} = (\text{RMFE.ct}, \text{cpABE.ct})$ to \mathcal{A} .

6. In the end, \mathcal{A} outputs a bit β' .

Observe that if the cpABE challenger chose $\hat{\beta} = 0$, then \mathcal{B} simulated the real world where m_0 was encrypted, else it simulated the real world where m_1 was encrypted with \mathcal{A} .

Hence, advantage of \mathcal{B} is $|\Pr[\beta' = 1 | \hat{\beta} = 0] - \Pr[\beta' = 1 | \hat{\beta} = 1]| = |\Pr[\beta' = 1 | \text{RPE.ct} = \text{RPE.Enc}(\text{RPE.msk}, f, m_0, L)] - \Pr[\beta' = 1 | \text{RPE.ct} = \text{RPE.Enc}(\text{RPE.msk}, f, m_1, L)]| = \epsilon$ (by assumption).

Admissibility of \mathcal{B} . We observe that for the challenge circuit $C_{L, \text{RMFE.ct}}$ and for all key queries $(\text{lb}_i, x_i), i \in [Q]$, queried by \mathcal{B} , we have $C_{L, \text{RMFE.ct}}((\text{lb}_i, x_i), \text{RMFE.sk}_{\text{lb}_i, x_i}) = 0$ as either (i) $\text{lb}_i \in L$, or (ii) $f(x_i) = 0$, which implies $\text{RMFE.Dec}(\text{RMFE.sk}_{\text{lb}_i, x_i}, \text{RMFE.ct}) = 0$ (due to RMFE correctness), by the admissibility condition on \mathcal{A} . So, if \mathcal{A} is admissible then so is \mathcal{B} . ■

Function Hiding Security

Theorem 3.35. Assume RMFE satisfies 1-query selective function hiding security (Def. 3.21), then the RPE scheme satisfies 1-query selective function hiding security (Def. 3.17).

Proof. Recall that in the function hiding security game, we want

$$\text{RPE.Enc}(\text{RPE.msk}, f_0, m, L) \approx_c \text{RPE.Enc}(\text{RPE.msk}, f_1, m, L),$$

where for all the key queries (lb, x) to the $\text{RPE.KeyGen}(\text{RPE.msk}, \cdot, \cdot)$ oracle, either $f_0(x) = f_1(x)$ or $\text{lb} \in L$.

We show that if there exists an adversary \mathcal{A} who has non-negligible advantage ϵ in the 1-query selective function hiding security game, then there exists a PPT adversary \mathcal{B} against the 1-query selective function-hiding security of RMFE scheme with the same advantage ϵ . The reduction is as follows.

1. \mathcal{B} runs \mathcal{A} and gets the revocation list L .
2. \mathcal{B} sends L as the challenge revocation list to the RMFE challenger. The challenger generates $(\text{RMFE.mpk}, \text{RMFE.msk})$ and returns RMFE.mpk to \mathcal{B} .
3. \mathcal{B} generates $(\text{cpABE.mpk}, \text{cpABE.msk}) \leftarrow \text{cpABE.Setup}(1^\lambda)$. It sets $\text{RPE.mpk} = (\text{RMFE.mpk}, \text{cpABE.mpk})$ and sends it to \mathcal{A} .
4. **Key Queries:** On each key query (lb, x) , \mathcal{B} does the following:
 - It sends a key query (lb, x) to the RMFE challenger and gets back $\text{RMFE.sk}_{\text{lb}, x}$.
 - Sets $\text{att} = ((\text{lb}, x), \text{RMFE.sk}_{\text{lb}, x})$ and computes $\text{cpABE.sk}_{\text{lb}, x} \leftarrow \text{cpABE.KeyGen}(\text{cpABE.msk}, \text{att})$.
 - It returns $\text{RPE.sk}_{\text{lb}, x} = (\text{att}, \text{cpABE.sk}_{\text{lb}, x})$ to \mathcal{A} .
5. **Challenge Query:** When \mathcal{A} sends the challenge functions (f_0, f_1) and message m , \mathcal{B} does the following:
 - Sends (f_0, f_1) as the challenge functions to the RMFE challenger. The challenger samples $\hat{\beta} \leftarrow \{0, 1\}$, computes $\text{RMFE.ct}_{\hat{\beta}} \leftarrow \text{RMFE.SK-Enc}(\text{RMFE.msk}, f_{\hat{\beta}}, L)$ and returns $\text{RMFE.ct}_{\hat{\beta}}$ to \mathcal{B} .
 - Constructs the circuit $C_{L, \text{RMFE.ct}_{\hat{\beta}}}$ as defined in the construction and computes $\text{cpABE.ct} \leftarrow \text{cpABE.Enc}(\text{cpABE.mpk}, C_{L, \text{RMFE.ct}_{\hat{\beta}}}, m)$.
 - Returns $\text{RPE.ct} = (\text{RMFE.ct}_{\hat{\beta}}, \text{cpABE.ct})$ to \mathcal{A} .
6. **Encryption Query:** When \mathcal{A} makes an encryption query $(\bar{f}, \bar{m}, \bar{L})$, \mathcal{B} does the following:
 - Sends a SK-Enc query (\bar{f}, \bar{L}) to the RMFE challenger and gets back RMFE.ct .

- Constructs the circuit $C_{\bar{L}, \text{RMFE.ct}}$ as defined in the construction.
- Computes $\text{cpABE.ct} \leftarrow \text{cpABE.Enc}(\text{cpABE.mpk}, C_{\bar{L}, \text{RMFE.ct}}, \bar{m})$.
- Returns $\text{RPE.ct} = (\text{RMFE.ct}, \text{cpABE.ct})$ to \mathcal{A} .

7. In the end, the adversary outputs a bit β' .

Observe that if the RMFE challenger chose $\hat{\beta} = 0$, then \mathcal{B} simulated the real world where f_0 was encrypted, else it simulated the real world where f_1 was encrypted with \mathcal{A} . Hence, advantage of \mathcal{B} is $|\Pr[\beta' = 1 | \hat{\beta} = 0] - \Pr[\beta' = 1 | \hat{\beta} = 1]| = |\Pr[\beta' = 1 | \text{RPE.ct} = \text{RPE.Enc}(\text{RPE.msk}, f_0, m, L)] - \Pr[\beta' = 1 | \text{RPE.ct} = \text{RPE.Enc}(\text{RPE.msk}, f_1, m, L)]| = \epsilon$ (by assumption).

Admissibility of \mathcal{B} . First, we note that since \mathcal{A} is allowed to make only one query, $(\bar{f}, \bar{m}, \bar{L})$ to the $\text{RPE.Enc}(\text{msk}, \cdot, \cdot, \cdot)$ oracle, \mathcal{B} also makes only one query, (\bar{f}, \bar{L}) , to the $\text{RMFE.SK-Enc}(\text{msk}, \cdot, \cdot)$ oracle. Next, we observe that since \mathcal{A} is restricted to make key queries (lb, x) such that either $f_0(x) = f_1(x)$ or $\text{lb} \in L$, thus \mathcal{B} also issues key queries (lb, x) to RMFE challenger such that either $f_0(x) = f_1(x)$ or $\text{lb} \in L$. Hence, if \mathcal{A} is admissible, then so is \mathcal{B} . ■

Broadcast Security

Theorem 3.36. *Assume RMFE satisfies 1-query mode hiding security (Def. 3.20), then the RPE scheme satisfies 1-query selective broadcast security (Def. 3.18).*

Proof. Recall that in the broadcast security game, we want

$$\text{RPE.Enc}(\text{RPE.msk}, f, m, L) \approx_c \text{RPE.Broadcast}(\text{RPE.mpk}, m, L),$$

where $f(x) = 1, \forall x \in \mathcal{X}$.

We show that if there exists an adversary \mathcal{A} who has non-negligible advantage ϵ in the broadcast security game, then there exists a PPT adversary \mathcal{B} against the mode-hiding

security of RMFE scheme with the same advantage ϵ . The reduction is as follows.

1. \mathcal{B} runs \mathcal{A} and gets the revocation list L .
2. \mathcal{B} sends L as the challenge revocation list to the RMFE challenger. The challenger generates $(\text{RMFE.mpk}, \text{RMFE.msk})$ and returns RMFE.mpk to \mathcal{B} .
3. \mathcal{B} generates $(\text{cpABE.mpk}, \text{cpABE.msk}) \leftarrow \text{cpABE.Setup}(1^\lambda)$. It sets $\text{RPE.mpk} = (\text{RMFE.mpk}, \text{cpABE.mpk})$ and sends it to \mathcal{A} .
4. **Key Queries:** On each key query (lb, x) , \mathcal{B} does the following:
 - It sends a key query (lb, x) to the RMFE challenger and gets back $\text{RMFE.sk}_{\text{lb},x}$.
 - Sets $\text{att} = ((\text{lb}, x), \text{RMFE.sk}_{\text{lb},x})$ and computes $\text{cpABE.sk}_{\text{lb},x} \leftarrow \text{cpABE.KeyGen}(\text{cpABE.msk}, \text{att})$.
 - It returns $\text{RPE.sk}_{\text{lb},x} = (\text{att}, \text{cpABE.sk}_{\text{lb},x})$ to \mathcal{A} .
5. **Challenge Query:** When \mathcal{A} sends the challenge function f and message m , \mathcal{B} does the following:
 - Sends f as the challenge function to the RMFE challenger. The challenger samples $\hat{\beta} \leftarrow \{0, 1\}$, and returns $\text{RMFE.ct}_{\hat{\beta}}$ to \mathcal{B} , where $\text{RMFE.ct}_0 \leftarrow \text{RMFE.SK-Enc}(\text{RMFE.msk}, f, L)$ and $\text{RMFE.ct}_1 \leftarrow \text{RMFE.PK-Enc}(\text{RMFE.mpk}, L)$.
 - Constructs the circuit $C_{L, \text{RMFE.ct}_{\hat{\beta}}}$ as defined in the construction and computes $\text{cpABE.ct} \leftarrow \text{cpABE.Enc}(\text{cpABE.mpk}, C_{L, \text{RMFE.ct}_{\hat{\beta}}}, m)$.
 - Returns $\text{RPE.ct} = (\text{RMFE.ct}_{\hat{\beta}}, \text{cpABE.ct})$ to \mathcal{A} .
6. **Encryption Query:** When \mathcal{A} makes an encryption query $(\bar{f}, \bar{m}, \bar{L})$, \mathcal{B} does the following:
 - Sends a SK-Enc query (\bar{f}, \bar{L}) to the RMFE challenger and gets back RMFE.ct .
 - Constructs the circuit $C_{\bar{L}, \text{RMFE.ct}}$ as defined in the construction.
 - Computes $\text{cpABE.ct} \leftarrow \text{cpABE.Enc}(\text{cpABE.mpk}, C_{\bar{L}, \text{RMFE.ct}}, \bar{m})$.
 - Returns $\text{RPE.ct} = (\text{RMFE.ct}, \text{cpABE.ct})$ to \mathcal{A} .
7. In the end, the adversary outputs a bit β' .

Observe that if the RMFE challenger chose $\hat{\beta} = 0$, then \mathcal{B} simulated the real world where the ciphertext was computed using Enc algorithm else it simulated the real world where the ciphertext was computed using Broadcast algorithm, with \mathcal{A} .

Hence, the advantage of \mathcal{B} is $|\Pr[\beta' = 1 | \hat{\beta} = 0] - \Pr[\beta' = 1 | \hat{\beta} = 1]| = |\Pr[\beta' = 1 | \text{RPE.ct} = \text{RPE.Enc}(\text{RPE.msk}, f, m, L)] - \Pr[\beta' = 1 | \text{RPE.ct} = \text{RPE.Broadcast}(\text{RPE.msk}, m, L)]| = \epsilon$ (by assumption).

Admissibility of \mathcal{B} . First, we note that since \mathcal{A} is allowed to make only one query, $(\bar{f}, \bar{m}, \bar{L})$ to the $\text{RPE.Enc}(\text{msk}, \cdot, \cdot, \cdot)$ oracle, \mathcal{B} also makes only one query, (\bar{f}, \bar{L}) , to the $\text{RMFE.SK-Enc}(\text{msk}, \cdot, \cdot)$ oracle. Next, we observe that since \mathcal{A} is restricted to output f , such that $f(x) = 1$ for all $x \in \mathcal{X}$, this implies \mathcal{B} also issues such f as the challenge query to the RMFE challenger in the above mode hiding security game. Thus, if \mathcal{A} is admissible, then so is \mathcal{B} . ■

3.8 EMBEDDED IDENTITY TRACE AND REVOKE

In this section, we define different variants of an embedded identity trace and revoke system (EITR). Our definitions extend the different notions that Goyal et al. [113] introduced, in the context of embedded identity traitor tracing, to incorporate the revocation list. Concretely, we define three variants of an EITR scheme: 1) Indexed EITR, 2) Bounded EITR, and 3) Unbounded EITR. Our goal is Unbounded EITR and other variants are introduced as intermediate goals. We show a construction of indexed EITR from RPE in Sec. 3.9, bounded EITR from indexed EITR in Sec. 3.10, and unbounded EITR from bounded EITR in Sec. 3.11. These implications hold for both secret key and public key settings.

An EITR scheme consists of five polynomial time algorithms—Setup, KeyGen, Enc, Dec and Trace. The syntax of the EITR variants mentioned above differs only in the inputs to the Setup, KeyGen and Trace algorithms.

Here we give a unified definition and later specify the distinctness of the three variants.

Consider a general identity space \mathcal{GID} , a label space \mathcal{L} and a message space \mathcal{M} . An embedded identity trace and revoke scheme $\text{EITR} = (\text{Setup}, \text{KeyGen}, \text{Enc}, \text{Dec}, \text{Trace})$ has the following syntax:

$\text{Setup}(1^\lambda, \text{params}_1) \rightarrow (\text{mpk}, \text{msk})$. The setup algorithm takes as input the security parameter λ and parameters params_1 . It outputs a master public key mpk and a master secret key msk .

$\text{KeyGen}(\text{msk}, \text{lb}, \text{gid}) \rightarrow \text{sk}_{\text{lb}, \text{gid}}$. The key generation algorithm takes as input the master secret key msk , a label $\text{lb} \in \mathcal{L}$, and a general identity $\text{gid} \in \mathcal{GID}$. It outputs a secret key $\text{sk}_{\text{lb}, \text{gid}}$.

$\text{Enc}(\text{mpk}, m, L) \rightarrow \text{ct}$. The encryption algorithm takes as input the master public key mpk , a message $m \in \mathcal{M}$, a revocation list $L \subseteq \mathcal{L}$ and outputs a ciphertext ct .

$\text{Dec}(\text{sk}_{\text{lb}, \text{gid}}, \text{ct}, L) \rightarrow y$. The decryption algorithm takes as input a secret key $\text{sk}_{\text{lb}, \text{gid}}$, a ciphertext ct , and a revocation list L and outputs $y \in \mathcal{M} \cup \{\perp\}$.

$\text{Trace}^D(\text{tk}, \text{params}_2, m_0, m_1, L) \rightarrow T$. The tracing algorithm takes as input a tracing key tk , parameter params_2 , two messages m_0, m_1 , a revocation list L and has an oracle access to a decoder D . It outputs a set of traitors T .

The above syntax captures both public key and secret key trace EITR schemes. For any *public tracing* EITR scheme, we have $\text{tk} = \text{mpk}$ and in the *secret tracing* EITR scheme, $\text{tk} = \text{msk}$. We now describe the properties satisfied by an EITR scheme.

Definition 3.22 (Correctness). An EITR scheme is said to be correct if there exists a negligible function $\text{negl}(\cdot)$ such that for all $\lambda \in \mathbb{N}$, any label $\text{lb} \in \mathcal{L}$, and any revocation

list $L \subseteq \mathcal{L}$ such that $\text{lb} \notin L$, the following holds

$$\Pr \left[\begin{array}{l} (\text{mpk}, \text{msk}) \leftarrow \text{Setup}(1^\lambda, \text{params}_1); \\ \text{Dec}(\text{sk}_{\text{lb}, \text{gid}}, \text{ct}, L) = m : \text{sk}_{\text{lb}, \text{gid}} \leftarrow \text{KeyGen}(\text{msk}, \text{lb}, \text{gid}); \\ \text{ct} \leftarrow \text{Enc}(\text{mpk}, m, L) \end{array} \right] \geq 1 - \text{negl}(\lambda).$$

Definition 3.23 (ind-cpa Security). An EITR scheme is said to be ind-cpa secure if for every stateful PPT adversary \mathcal{A} , there exists a negligible function $\text{negl}(\cdot)$ such that for every $\lambda \in \mathbb{N}$, the following holds

$$\Pr \left[\begin{array}{l} \text{params}_1 \leftarrow \mathcal{A}(1^\lambda); \\ (\text{mpk}, \text{msk}) \leftarrow \text{Setup}(1^\lambda, \text{params}_1); \\ (m_0, m_1, L) \leftarrow \mathcal{A}^{\text{KeyGen}(\text{msk}, \cdot, \cdot)}(\text{mpk}); \\ \beta \leftarrow \{0, 1\}; \text{ct}_\beta \leftarrow \text{Enc}(\text{mpk}, m_\beta, L) \end{array} \right] \leq \frac{1}{2} + \text{negl}(\lambda)$$

where the adversary has the access to the $\text{KeyGen}(\text{msk}, \cdot, \cdot)$ oracle which has msk hardwired and \mathcal{A} is admissible only if for all the key generation queries (lb, gid) to the KeyGen oracle, $\text{lb} \in L$.

Definition 3.24 (Very Selective ind-cpa Security). The very selective ind-cpa security of an EITR scheme is defined in the same way as Def. 3.23, except that the adversary outputs the revocation list L and all the key queries $\{\text{lb}_i, \text{gid}_i\}_{i \in [Q]}$, where Q is the number of key queries made, along with params_1 before the Setup algorithm is run.

In the following, we specify the space \mathcal{GID} , inputs $\text{params}_1, \text{gid}$ and params_2 to the Setup , KeyGen and Trace algorithm, respectively, and then define the secure tracing guarantee of each of the EITR notions separately. We let $\mathcal{ID} = \{0, 1\}^\kappa$ denote the identity space.

Remark 7. We assume that there exists an efficiently computable mapping $\text{map} : \mathcal{GID} \rightarrow \mathcal{L}$, that *uniquely* maps an $\text{gid} \in \mathcal{GID}$ to a label $\text{lb} \in \mathcal{L}$. This can be easily ensured, for e.g. by making label lb a part of id . We further note that in real world applications one may also want to ensure that any label lb is associated with at most one id . This can be achieved by using a collision resistant hash function.

Experiment $\text{Exp-Ind-TR}_{\mathcal{A},\epsilon}(\lambda)$

- $1^\kappa, 1^{n_{\text{ind}}} \leftarrow \mathcal{A}(1^\lambda)$
- $(\text{mpk}, \text{msk}) \leftarrow \text{Setup}(1^\lambda, 1^\kappa, n_{\text{ind}})$
- $(D, m_0, m_1, L) \leftarrow \mathcal{A}^{\text{KeyGen}(\text{msk}, \cdot, \cdot, \cdot)}(\text{mpk})$
- $T \leftarrow \text{Trace}^D(\text{tk}, 1^{1/\epsilon(\lambda)}, m_0, m_1, L)$

Here \mathcal{A} has the oracle access to $\text{KeyGen}(\text{msk}, \cdot, \cdot, \cdot)$, which has msk hardwired and on query $(\text{lb}, \text{id}, i)$, it outputs $\text{sk}_{\text{lb}, \text{id}, i}$. The adversary is admissible only if it makes at most one key query for each index.

Figure 3.4: Security Experiment: Exp-Ind-TR

3.8.1 Indexed Trace and Revoke with Embedded Identity

In an indexed EITR scheme, the key is generated w.r.t. an identity and an index. We have $\mathcal{GID} = \mathcal{ID} \times [n_{\text{ind}}]$, where $[n_{\text{ind}}]$ is the index space for $n_{\text{ind}} \in \mathbb{N}$, $\text{params}_1 = (1^\kappa, n_{\text{ind}})$, $\text{gid} = (\text{id}, i) \in \mathcal{ID} \times [n_{\text{ind}}]$, and $\text{params}_2 = y$ for some $y > 1$.

We now define the secure tracing requirement.

Definition 3.25 (Secure Tracing). Let $\text{Ind-TR} = (\text{Setup}, \text{KeyGen}, \text{Enc}, \text{Dec}, \text{Trace})$ be an indexed EITR scheme. For any non-negligible function $\epsilon(\cdot)$ and stateful PPT adversary \mathcal{A} , define an experiment $\text{Exp-Ind-TR}_{\mathcal{A},\epsilon}(\lambda)$ as in Figure 3.4.

Let S be the set of key queries $(\text{lb}, \text{id}, i)$ queried by \mathcal{A} and $S_{\mathcal{ID}} = \{\text{id} : \exists \text{lb} \in \mathcal{L}, i \in [n_{\text{ind}}] \text{ s.t. } (\text{lb}, \text{id}, i) \in S\}$.

Consider the following probabilistic events and their corresponding probabilities:

- **Good-Decoder** : $\Pr[D(\text{ct}) = b : b \leftarrow \{0, 1\}, \text{ct} \leftarrow \text{Enc}(\text{mpk}, m_b, L)] \geq 1/2 + \epsilon(\lambda)$
 $\Pr\text{-Good-Decoder}_{\mathcal{A},\epsilon}(\lambda) = \Pr[\text{Good-Decoder}]$.
- **Cor-Tr** : $|T| > 0, (T \subseteq S_{\mathcal{ID}}) \wedge (T_{\text{lb}} \cap L = \emptyset)$, where $T_{\text{lb}} = \{\text{lb} \in \mathcal{L} : \exists i \in [n_{\text{ind}}], \exists \text{id} \in T, (\text{lb}, \text{id}, i) \in S\}$.
 $\Pr\text{-Cor-Tr}_{\mathcal{A},\epsilon}(\lambda) = \Pr[\text{Cor-Tr}]$.
- **Fal-Tr** : $(T \not\subseteq S_{\mathcal{ID}}) \wedge (T_{\text{lb}} \cap L = \emptyset)$
 $\Pr\text{-Fal-Tr}_{\mathcal{A},\epsilon}(\lambda) = \Pr[\text{Fal-Tr}]$.

The Ind-TR scheme is said to satisfy secure tracing property if for every PPT adversary \mathcal{A} and non-negligible function $\epsilon(\cdot)$, there exists negligible functions $\text{negl}_1(\cdot)$ and $\text{negl}_2(\cdot)$, such that for all $\lambda \in \mathbb{N}$, the following holds:

$$\text{Pr-Fal-Tr}_{\mathcal{A},\epsilon}(\lambda) \leq \text{negl}_1(\lambda), \quad \text{Pr-Cor-Tr}_{\mathcal{A},\epsilon}(\lambda) \geq \text{Pr-Good-Decoder}_{\mathcal{A},\epsilon}(\lambda) - \text{negl}_2(\lambda).$$

Definition 3.26 (Very Selective Secure Tracing). The very selective secure tracing of an indexed EITR scheme is defined in the same way as Def. 3.25, except that the adversary outputs the challenge revocation list L and all the key queries $(\text{lb}, \text{id}, i)$ along with $(1^\kappa, 1^{n_{\text{ind}}})$ in the beginning of the $\text{Exp-Ind-TR}_{\mathcal{A},\epsilon}(\lambda)$.

Remark 8. In the definition above, **Setup** takes n_{ind} as an input in the binary form rather than in the unary form. This indicates that **Setup** runs in polynomial time, even if $n_{\text{ind}} = 2^{\text{poly}(\lambda)}$. On the other hand, the adversary outputs n_{ind} in the unary form in the security game (Fig. 3.4). This indicates that we consider the security game only for the case where $n_{\text{ind}} = \text{poly}(\lambda)$. Looking ahead, the former property is necessary when we convert bounded EITR into unbounded EITR in Sec. 3.11.

3.8.2 Bounded Trace and Revoke with Embedded Identity

In a bounded EITR scheme, we have $\mathcal{GID} = \mathcal{ID}$, $\text{params}_1 = (1^\kappa, 1^{n_{\text{bd}}})$, where $n_{\text{bd}} \in \mathbb{N}$ is the bound on the number of key queries that an adversary can make in the *correct trace experiment* game, $\text{gid} = \text{id}$ for $\text{id} \in \mathcal{ID}$, and $\text{params}_2 = y$ for some $y > 1$.

We now define the secure tracing requirement.

Definition 3.27 (Secure Tracing). Let BD-TR = (Setup, KeyGen, Enc, Dec, Trace) be a bounded EITR scheme. For any non-negligible function $\epsilon(\cdot)$ and stateful PPT adversary \mathcal{A} , define an experiment $\text{Exp-BD-TR}_{\mathcal{A},\epsilon}(\lambda)$ as in Figure 3.5.

Let S be the set of key queries made by \mathcal{A} and $S_{\mathcal{ID}} = \{\text{id} : \exists \text{lb} \in \mathcal{L} \text{ s.t. } (\text{lb}, \text{id}) \in S\}$.

Consider the following probabilistic events and their corresponding probabilities :

- Good-Decoder : $\Pr[D(\text{ct}) = b : b \leftarrow \{0, 1\}, \text{ct} \leftarrow \text{Enc}(\text{mpk}, m_b, L)] \geq 1/2 + \epsilon(\lambda)$
 $\text{Pr-Good-Decoder}_{\mathcal{A},\epsilon}(\lambda) = \Pr[\text{Good-Decoder} \wedge |S_{\mathcal{ID}}| \leq n_{\text{ind}}].$

Experiment $\text{Exp-BD-TR}_{\mathcal{A},\epsilon}(\lambda)$

- $1^\kappa, 1^{n_{\text{bd}}} \leftarrow \mathcal{A}(1^\lambda)$.
- $(\text{mpk}, \text{msk}) \leftarrow \text{Setup}(1^\lambda, 1^\kappa, n_{\text{bd}})$.
- $(D, m_0, m_1, L) \leftarrow \mathcal{A}^{\text{KeyGen}(\text{msk}, \cdot, \cdot)}(\text{mpk})$.
- $T \leftarrow \text{Trace}^D(\text{tk}, 1^{1/\epsilon}, m_0, m_1, L)$.

Here \mathcal{A} has the oracle access to $\text{KeyGen}(\text{msk}, \cdot, \cdot)$, which has msk hardwired and on query (lb, id) , it outputs $\text{sk}_{\text{lb}, \text{id}}$.

Figure 3.5: Security Experiment: Exp-BD-TR

- **Cor-Tr** : $|T| > 0, (T \subseteq S_{\text{ID}}) \wedge (T_{\text{lb}} \cap L = \phi)$, where $T_{\text{lb}} = \{\text{lb} \in \mathcal{L} : \exists \text{id} \in T, (\text{lb}, \text{id}) \in S\}$. $\text{Pr-Cor-Tr}_{\mathcal{A},\epsilon}(\lambda) = \Pr[\text{Cor-Tr}]$.
- **Fal-Tr** : $(T \not\subseteq S_{\text{ID}}) \wedge (T_{\text{lb}} \cap L = \phi)$
 $\text{Pr-Fal-Tr}_{\mathcal{A},\epsilon}(\lambda) = \Pr[\text{Fal-Tr}]$.

The scheme is said to satisfy secure tracing if for every PPT adversary \mathcal{A} and non-negligible function $\epsilon(\cdot)$, there exists negligible functions $\text{negl}_1(\cdot)$ and $\text{negl}_2(\cdot)$, such that for all $\lambda \in \mathbb{N}$, the following holds:

$$\text{Pr-Fal-Tr}_{\mathcal{A},\epsilon}(\lambda) \leq \text{negl}_1(\lambda), \quad \text{Pr-Cor-Tr}_{\mathcal{A},\epsilon}(\lambda) \geq \text{Pr-Good-Decoder}_{\mathcal{A},\epsilon}(\lambda) - \text{negl}_2(\lambda).$$

Definition 3.28 (Very Selective Secure Tracing). The very selective secure tracing of a bounded EITR scheme is defined in the same way as Def. 3.27, except that the adversary outputs the challenge revocation list L and all the key queries (lb, id) along with $(1^\kappa, 1^{n_{\text{bd}}})$ in the beginning of the $\text{Exp-Ind-TR}_{\mathcal{A},\epsilon}(\lambda)$.

Remark 9. We point out that the bound on the number of key queries by the adversary is only required for the correct trace guarantee. The false trace guarantee will hold even if the adversary exceeds n_{bd} key queries— this is essential for the transformation from bounded-EITR scheme to unbounded-EITR scheme.

Experiment $\text{Exp-TR}_{\mathcal{A},\epsilon,p}(\lambda)$

- $1^\kappa \leftarrow \mathcal{A}(1^\lambda)$.
- $(\text{mpk}, \text{msk}) \leftarrow \text{Setup}(1^\lambda, 1^\kappa)$.
- $(D, m_0, m_1, L) \leftarrow \mathcal{A}^{\text{KeyGen}(\text{msk}, \cdot, \cdot)}(\text{mpk})$.
- $T \leftarrow \text{Trace}^D(\text{tk}, 1^{1/\epsilon(\lambda)}, p(\lambda), m_0, m_1, L)$.

Here \mathcal{A} has the oracle access to $\text{KeyGen}(\text{msk}, \cdot, \cdot)$, which has msk hardwired and on query (lb, id) , it outputs $\text{sk}_{\text{lb}, \text{id}}$.

Figure 3.6: Security Experiment: Exp-TR

3.8.3 Unbounded Trace and Revoke with Embedded Identity

In an unbounded EITR scheme, we have $\mathcal{GID} = \mathcal{ID}$, $\text{params}_1 = 1^\kappa$, $\text{gid} = \text{id}$ for $\text{id} \in \mathcal{ID}$, and $\text{params}_2 = (y, Q_{\text{bd}})$ where $y > 1$, $Q_{\text{bd}} \in \mathbb{N}$.

We now define the secure tracing requirement.

Definition 3.29 (Secure Tracing). Let $\text{TR} = (\text{Setup}, \text{KeyGen}, \text{Enc}, \text{Dec}, \text{Trace})$ be an unbounded EITR scheme. For any non-negligible function $\epsilon(\cdot)$, polynomial $p(\cdot)$ and stateful PPT adversary \mathcal{A} , define the experiment $\text{Exp-TR}_{\mathcal{A},\epsilon,p}(\lambda)$ as in Figure 3.6.

Let S be the set of key queries made by \mathcal{A} and $S_{\mathcal{ID}} = \{\text{id} : \exists \text{lb} \in \mathcal{L} \text{ s.t. } (\text{lb}, \text{id}) \in S\}$.

Consider the following probabilistic events and their corresponding probabilities :

- **Good-Decoder** : $\Pr[D(\text{ct}) = b : b \leftarrow \{0, 1\}, \text{ct} \leftarrow \text{Enc}(\text{pk}, m_b, L) \geq 1/2 + \epsilon(\lambda)]$
 $\Pr\text{-Good-Decoder}_{\mathcal{A},\epsilon,p}(\lambda) = \Pr[\text{Good-Decoder} \wedge |S_{\mathcal{ID}}| \leq p(\lambda)]$.
- **Cor-Tr** : $|T| > 0, (T \subseteq S_{\mathcal{ID}}) \wedge (T_{\text{lb}} \cap L = \emptyset)$, where $T_{\text{lb}} = \{\text{lb} \in \mathcal{L} : \exists \text{id} \in T, (\text{lb}, \text{id}) \in S\}$.
 $\Pr\text{-Cor-Tr}_{\mathcal{A},\epsilon,p}(\lambda) = \Pr[\text{Cor-Tr}]$.
- **Fal-Tr** : $(T \not\subseteq S_{\mathcal{ID}}) \wedge (T_{\text{lb}} \cap L = \emptyset)$
 $\Pr\text{-Fal-Tr}_{\mathcal{A},\epsilon,p}(\lambda) = \Pr[\text{Fal-Tr}]$.

The scheme is said to satisfy secure traitor tracing property if for every PPT adversary \mathcal{A} and non-negligible function $\epsilon(\cdot)$, there exists negligible functions $\text{negl}_1(\cdot)$ and $\text{negl}_2(\cdot)$,

such that for all $\lambda \in \mathbb{N}$, the following holds:

$$\text{Pr-Fal-Tr}_{\mathcal{A},\epsilon,p}(\lambda) \leq \text{negl}_1(\lambda), \quad \text{Pr-Cor-Tr}_{\mathcal{A},\epsilon,p}(\lambda) \geq \text{Pr-Good-Decoder}_{\mathcal{A},\epsilon,p}(\lambda) - \text{negl}_2(\lambda).$$

Definition 3.30 (Very Selective Secure Tracing). The very selective secure tracing of an unbounded EITR scheme is defined in the same way as Def. 3.29, except that the adversary outputs the challenge revocation list L and all the key queries (lb, id) along with 1^κ in the beginning of the $\text{Exp-TR}_{\mathcal{A},\epsilon}(\lambda)$.

Remark 10. [Remark 10.4, [113]]. Note that here the trace algorithm takes an additional parameter Q_{bd} . In the correct trace definition, we require that as long as the tracing algorithm uses a bound greater than the number of keys queried, the tracing algorithm must identify at least one traitor. However, the false trace guarantee should hold for all polynomially bounded Q_{bd} values. In particular, even if the number of keys queried is more than the bound used in tracing, the trace algorithm must not output an identity that was not queried. We can show that this definition implies the ‘standard’ tracing definition where the trace algorithm does not take this bound as input. One simply needs to run this bounded-version of trace with increasing powers of two until the trace algorithm outputs at least one traitor.

3.9 INDEXED TRACE AND REVOKE WITH EMBEDDED IDENTITY

In this section we construct an indexed (secret/public tracing)-EITR scheme from a (secret/public key)-RPE scheme. We present the construction and proofs for the secret trace setting primarily and also outline the differences in the public trace setting simultaneously.

3.9.1 Construction

Consider the identity space $\mathcal{ID} = \{0, 1\}^\kappa$ and index bound n_{ind} . For our purpose, it suffices to assume $n_{\text{ind}} \leq 2^{2\lambda}$. Let $\text{RPE} = (\text{RPE.Setup}, \text{RPE.KeyGen}, \text{RPE.Broadcast}, \text{RPE.Enc}, \text{RPE.Dec})$ be a

Inputs: an identity id and an index i .

Hardwired Values : Indices $j \in [n_{\text{ind}}]$, $\ell \in [\kappa]$, a bit $b \in \{0, 1\}$.

Output : 0/1.

$$f[j, \ell, b](\text{id}, i) = \begin{cases} 1 & \text{if } (i > j) \vee (i = j \wedge \ell = \perp) \vee (i = j \wedge \text{id}_\ell = 1 - b) \\ 0 & \text{otherwise} \end{cases}$$

Figure 3.7: Comparison Function $f[j, \ell, b]$

(secret/public key) RPE scheme²⁵ with attribute space $\mathcal{ID} \times [n_{\text{ind}}]$, and supports the function class $\mathcal{F} = \{f[j, \ell, b]\}_{j \in [n_{\text{ind}}], \ell \in [\kappa], b \in \{0, 1\}}$, where $f[j, \ell, b] : \mathcal{ID} \times [n_{\text{ind}}] \rightarrow \{0, 1\}$ is as defined in figure 3.9.1.

We construct an indexed (secret/public tracing)-EITR scheme with identity space $\mathcal{ID} = \{0, 1\}^\kappa$ and index bound n_{ind} as follows :

$\text{Setup}(1^\lambda, 1^\kappa, n_{\text{ind}}) \rightarrow (\text{mpk}, \text{msk})$. The setup algorithm does the following:

- Samples $(\text{RPE.mpk}, \text{RPE.msk}) \leftarrow \text{RPE.Setup}(1^\lambda)$.
- Outputs $\text{mpk} = \text{RPE.mpk}$ and $\text{msk} = \text{RPE.msk}$.

$\text{KeyGen}(\text{msk}, (\text{lb}, \text{id}, i)) \rightarrow \text{sk}_{\text{lb}, \text{id}, i}$. The key generation algorithm does the following:

- Parse $\text{msk} = \text{RPE.msk}$.
- Sets $x = (\text{id}, i)$ and runs $\text{RPE.sk}_{\text{lb}, \text{id}, i} \leftarrow \text{RPE.KeyGen}(\text{RPE.msk}, \text{lb}, x)$.
- Outputs $\text{sk}_{\text{lb}, \text{id}, i} = \text{RPE.sk}_{\text{lb}, \text{id}, i}$.

$\text{Enc}(\text{mpk}, m, L) \rightarrow \text{ct}$. The encryption algorithm does the following:

- Parse $\text{mpk} = \text{RPE.mpk}$.
- Compute $\text{RPE.ct} \leftarrow \text{RPE.Broadcast}(\text{RPE.mpk}, m, L)$. (In public trace setting, the algorithm computes $\text{RPE.ct} \leftarrow \text{RPE.Enc}(\text{RPE.mpk}, f[1, \perp, 0]$,

²⁵Public key RPE does not have the Broadcast algorithm.

$m, L))$.

- Outputs $\text{ct} = \text{RPE.ct}$.

$\text{Dec}(\text{sk}_{\text{lb}, \text{id}, i}, \text{ct}, L) \rightarrow m'$. The decryption algorithm does the following:

- Parse $\text{sk}_{\text{lb}, \text{id}, i}$ as $\text{RPE.sk}_{\text{lb}, \text{id}, i}$ and ct as RPE.ct .
- Computes and outputs $m' \leftarrow \text{RPE.Dec}(\text{RPE.sk}_{\text{lb}, \text{id}, i}, \text{RPE.ct}, L)$.

$\text{Trace}^D(\text{tk}, y, m_0, m_1, L) \rightarrow T$. Parse tk as $\text{msk} = \text{RPE.msk}$ (in the public trace setting tk is $\text{mpk} = \text{RPE.mpk}$). The tracing algorithm is a two phased process.

Briefly the tracing is implemented as follows:

1. First, we run the Index.Trace algorithm as defined in Figure 3.8, on the index space $[n_{\text{ind}}]$. If the key associated with index $i \in [n_{\text{ind}}]$ is used in constructing the decoder box, then the bit b in the output of the Index.Trace algorithm is 1. We maintain a set T_{index} to store all such indices on which $b = 1$.
2. Next, the ID.Trace algorithm, defined in Figure 3.9.1, takes the set T_{index} as input and uses the decoder D to compute the identity associated with each index.

Specifically, the algorithm runs as follows:

- Set $T_{\text{index}} := \phi$. For $i = 1$ to n_{ind} ,
 - Compute $(b, p, q) \leftarrow \text{Index.Trace}(\text{tk}, y, m_0, m_1, L, i)$.
 - If $b = 1$, set $T_{\text{index}} := T_{\text{index}} \cup (i, p, q)$.
- Set $T = \phi$. For $(i, p, q) \in T_{\text{index}}$,
 - Compute $\text{id} \leftarrow \text{ID.Trace}(\text{tk}, y, m_0, m_1, L, (i, p, q))$.
 - Set $T := T \cup \{\text{id}\}$.
- Output T .

Correctness. We prove that the above construction of indexed (secret/public tracing)-

Algorithm Index.Trace(tk, y, m₀, m₁, L, i)

Inputs: Tracing key tk, a parameter y, messages m₀, m₁, revocation list L and an index i.

Output : (b, p, q), where b ∈ {0, 1} and p, q ∈ [0, 1] ∪ {⊥}.

Let ε = ⌊1/y⌋. It sets N = λ · n_{ind}/ε and temp₁ = temp₂ = 0. For j = 1 to N, it computes the following:

1. It samples b_j ← {0, 1} and computes ct_{j,1} ← RPE.Enc(tk, f[i, ⊥, 0], m_{b_j}, L) and sends ct_{j,1} to D.

If D outputs b_j, set temp₁ = temp₁ + 1, else set temp₁ = temp₁ - 1.

2. It samples c_j ← {0, 1} and computes ct_{j,2} ← RPE.Enc(tk, f[i + 1, ⊥, 0], m_{c_j}, L) and sends ct_{j,2} to D.

If D outputs c_j, set temp₂ = temp₂ + 1, else set temp₂ = temp₂ - 1.

If $\frac{\text{temp}_1 - \text{temp}_2}{N} > \frac{\epsilon}{4n_{\text{ind}}}$, output $\left(1, \frac{\text{temp}_1}{N}, \frac{\text{temp}_2}{N}\right)$ else output (0, ⊥, ⊥).

Figure 3.8: Index Tracing

Algorithm ID.Trace(tk, y, m₀, m₁, L, (i, p, q))

Inputs: Tracing key tk, a parameter y, messages m₀, m₁, revocation list L, index i, and probabilities p, q.

Output : id ∈ {0, 1}^κ

Let ε = ⌊1/y⌋. It sets N = λ · n_{ind}/ε and temp_ℓ = 0 for ℓ ∈ [κ]. For ℓ = 1 to κ, it does as follows:

1. For j = 1 to N

It samples b_j ← {0, 1} and computes ct_j ← RPE.Enc(tk, f[i, ℓ, 0], m_{b_j}, L) and sends ct_j to D.

If D outputs b_j, set temp_ℓ = temp_ℓ + 1, else set temp_ℓ = temp_ℓ - 1.

Let id be an empty string. For ℓ = 1 to κ, do the following:

1. If $\frac{p+q}{2} > \frac{\text{temp}_\ell}{N}$, set id_ℓ = 0, else set id_ℓ = 1.

Output id = id₁ || ... || id_κ.

Figure 3.9: Identity Tracing

EITR scheme satisfies correctness (Def. 3.22) via the following theorem.

Theorem 3.37. *If RPE is a correct (secret/public key)-RPE scheme, then the above construction of indexed (secret/public tracing)-EITR scheme is correct.*

Proof. For the secret trace setting, the correctness follows directly from the broadcast correctness of the underlying (secret-key) RPE scheme.

For the public trace setting, firstly we observe that for any $(\text{id}, i) \in \mathcal{ID} \times [n_{\text{ind}}]$, $f[1, \perp, 0](\text{id}, i) = 1$. Hence, correctness follows directly from the encryption correctness of the underlying (public-key) RPE scheme.

■

Efficiency. We can instantiate the above construction using public key and secret key RPE. The size of each parameter is directly inherited from that of the underlying RPE. We have $|x| = |\text{id}| + \log n_{\text{ind}} \leq |\text{id}| + \lambda$ and $|f| := |f[j, \ell, b]| = \log n_{\text{ind}} + \log \kappa + 1 = O(\lambda)$. We note that here, $|f[j, \ell, b]|$ refers to the description size of the function $|f[j, \ell, b]|$, not the size of the circuit implementing it. In particular, the latter can depend on $|\text{id}|$ while the former is independent from $|\text{id}|$.

Secret Tracing Setting. If we instantiate the scheme with our secret key RPE in Sec. 3.7, we have $|\text{mpk}|, |\text{ct}|, |\text{sk}| = \text{poly}(\lambda, |f|, |x|, |\text{lb}|) = \text{poly}(\lambda, |\text{id}|, |\text{lb}|)$.

Public Tracing Setting. If we instantiate the scheme with our public key RPE in Sec. 3.5, we have $|\text{mpk}|, |\text{ct}| = \text{poly}(\lambda, |f|, |\text{lb}|) = \text{poly}(\lambda, |\text{lb}|)$, $|\text{sk}| = \text{poly}(\lambda, |f|, |x|, |\text{lb}|) = \text{poly}(\lambda, |\text{id}|, |\text{lb}|)$.

3.9.2 Security

In this section we show that our construction of the indexed (secret/public tracing)-EITR scheme is secure.

IND-CPA security.

Theorem 3.38. *If (secret/public key)-RPE scheme satisfies (0-query- very selective/adaptive) message hiding property (Def. 3.14/Def. 3.13) then our construction of indexed (secret/public tracing)-EITR scheme is (very selective/adaptive) ind-cpa secure (Def. 3.23/Def. 3.24).*

Proof. Recall that for ind-cpa security, we need

$$\text{Enc}(\text{mpk}, m_0, L) \approx_c \text{Enc}(\text{mpk}, m_1, L).$$

For indexed *secret tracing* EITR scheme, this is equivalent to

$$\text{RPE.Broadcast}(\text{RPE.mpk}, m_0, L) \approx_c \text{RPE.Broadcast}(\text{RPE.mpk}, m_1, L). \quad (3.5)$$

To prove this, we define the following hybrids:

Hyb₀ : This is the real world with the challenge bit $b = 0$. In particular, the challenger returns the ciphertext $\text{ct} = \text{RPE.Broadcast}(\text{RPE.mpk}, m_0, L)$.

Hyb₁ : In this hybrid, the challenger returns the ciphertext $\text{ct} = \text{RPE.Enc}(\text{RPE.mpk}, f[1, \perp, 0], m_0, L)$. Indistinguishability from **Hyb₀** follows from the broadcast security of secret-key RPE.

Hyb₂ : In this hybrid, the challenger returns the ciphertext $\text{ct} = \text{RPE.Enc}(\text{RPE.mpk}, f[1, \perp, 0], m_1, L)$. Indistinguishability from **Hyb₁** follows from the message hiding property of secret-key RPE.

Hyb₃ : In this hybrid, the challenger returns the ciphertext $\text{ct} = \text{RPE.Broadcast}(\text{RPE.mpk}, m_1, L)$. This is the real world with $b = 1$. Indistinguishability from **Hyb₂** follows from the broadcast security of secret-key RPE.

For indexed *public trace* EITR scheme, ind-cpa security is equivalent to

$$\text{RPE.Enc}(\text{RPE.mpk}, f[1, \perp, 0], m_0, L) \approx_c \text{RPE.Enc}(\text{RPE.mpk}, f[1, \perp, 0], m_1, L). \quad (3.6)$$

The indistinguishability here follows directly from the message hiding property of the underlying RPE scheme. \blacksquare

Secure Tracing Analysis. We now show that our construction of indexed (secret/public tracing)-EITR scheme achieves the correct and false trace guarantees. The following analysis is mostly taken from ([113], Section 5.2.2) with appropriate modifications to incorporate the revocation list.

False Trace Guarantee. The false trace guarantee of a trace and revoke scheme ensures that the tracing algorithm does not falsely accuse any user. We prove that there does not exist a PPT adversary who can output a decoder D such that the tracing algorithm, when executed using this decoder D , outputs an identity for which key was not queried by the adversary or the corresponding label is in the revocation list.

Theorem 3.39. *For every stateful PPT adversary \mathcal{A} and non-negligible function $\epsilon(\cdot)$, there exists negligible functions $\text{negl}(\cdot)$, such that for all $\lambda \in \mathbb{N}$*

$$\text{Pr-Fal-Tr}_{\mathcal{A}, \epsilon}(\lambda) \leq \text{negl}(\lambda)$$

where the probability $\text{Pr-Fal-Tr}_{\mathcal{A}, \epsilon}$ is as defined in Def. 3.25.

Proof. Let $S \subseteq \mathcal{L} \times \mathcal{ID} \times [n_{\text{ind}}]$ be the set of label-identity-index pairs on which \mathcal{A} issues key queries. That is, $S = \{(\text{lb}, \text{id}, i) : \mathcal{A} \text{ issues a key query on } (\text{lb}, \text{id}, i)\}$. Let us first recall the assumption that \mathcal{A} can issue at most one key query for any index $i \in [n_{\text{ind}}]$. Next, we set up some notations. For $\text{lb} \in \mathcal{L}$, $i \in [n_{\text{ind}}]$, $\text{id} \in \mathcal{ID}$ and a revocation list L ,

let

$$\begin{aligned}
S_{\text{index}} &= \{i : (\text{lb}, \text{id}, i) \in S \text{ for some } (\text{lb}, \text{id}) \in \mathcal{L} \times \mathcal{ID}\}, \\
L_{\text{index}} &= \{i : (\text{lb}, \text{id}, i) \in S \text{ and } \text{lb} \in L\}, \\
B &= \{(\text{id}, i) : (\text{lb}, \text{id}, i) \in S \text{ and } \text{lb} \notin L\}, \\
B_{\text{index}} &= \{i : (\text{id}, i) \in B\}.
\end{aligned}$$

For any decoder box D , messages m_0, m_1 , revocation list L , for any $i \in [n_{\text{ind}} + 1], \ell \in [\kappa]$, $\text{lb} \in \mathcal{L}$ we define

$$p_{i,\perp}^D = \Pr[D(\text{ct}) = b : b \leftarrow \{0, 1\}, \text{ct} \leftarrow \text{RPE.Enc}(\text{tk}, f[i, \perp, 0], m_b, L)]$$

$$p_{i,\ell}^D = \Pr[D(\text{ct}) = b : b \leftarrow \{0, 1\}, \text{ct} \leftarrow \text{RPE.Enc}(\text{tk}, f[i, \ell, 0], m_b, L)]$$

where the probability is taken over the random coins to decoder D and the randomness used

during the encryption. We show that $\text{Pr-Fal-Tr}_{\mathcal{A}, \epsilon(\lambda)} \leq \text{negl}(\lambda)$. For $i \in [n_{\text{ind}}], \ell \in [\kappa]$,

we also define the following events:

$$\text{Diff-Adv}_i^D : p_{i,\perp}^D - p_{i+1,\perp}^D > \epsilon/8n_{\text{ind}}$$

$$\text{Diff-Adv}_{i,\ell,\text{lwr}}^D : p_{i,\perp}^D - p_{i,\ell}^D > \epsilon/16n_{\text{ind}}$$

$$\text{Diff-Adv}_{i,\ell,\text{upr}}^D : p_{i,\ell}^D - p_{i+1,\perp}^D > \epsilon/16n_{\text{ind}}$$

$$\text{Diff-Adv}^D : \bigvee_{i \in [n_{\text{ind}}] \setminus (S_{\text{index}} \setminus L_{\text{index}})} \text{Diff-Adv}_i^D \bigvee_{\substack{(\text{id}, i) \in B, \ell \in [\kappa], \\ \text{s.t. id}_\ell = 1}} \text{Diff-Adv}_{i,\ell,\text{lwr}}^D \bigvee_{\substack{(\text{id}, i) \in B, \ell \in [\kappa] \\ \text{s.t. id}_\ell = 0}} \text{Diff-Adv}_{i,\ell,\text{upr}}^D$$

We will drop the dependence on D for the ease of notation. We have

$$\begin{aligned}
\text{Pr}[\text{Fal-Tr}] &= \text{Pr}[\text{Fal-Tr} \mid \overline{\text{Diff-Adv}}] \text{Pr}[\overline{\text{Diff-Adv}}] + \text{Pr}[\text{Fal-Tr} \mid \text{Diff-Adv}] \text{Pr}[\text{Diff-Adv}] \\
&\leq \text{Pr}[\text{Fal-Tr} \mid \overline{\text{Diff-Adv}}] + \text{Pr}[\text{Diff-Adv}]
\end{aligned}$$

$$\begin{aligned}
&= \Pr[\text{Fal-Tr} \mid \overline{\text{Diff-Adv}}] + \sum_{i \in [n_{\text{ind}}]} \Pr[i \notin S_{\text{index}} \setminus L_{\text{index}} \wedge \text{Diff-Adv}_i] \\
&\quad + \sum_{(i, \ell) \in [n_{\text{ind}}] \times [\kappa]} \Pr \left[\exists \text{id} \in \{0, 1\}^\kappa \text{ s.t. } (\text{id}, i) \in B \right. \\
&\quad \left. \wedge ([\text{Diff-Adv}_{i, \ell, \text{low}} \wedge \text{id}_\ell = 1] \vee [\text{Diff-Adv}_{i, \ell, \text{upr}} \wedge \text{id}_\ell = 0]) \right]
\end{aligned}$$

Now, we argue that each of the terms on the RHS is bounded by a negligible function.

Lemma 3.40. *For every stateful PPT adversary \mathcal{A} , there exists a negligible function $\text{negl}_1(\cdot)$, such that $\forall \lambda \in \mathbb{N}$, we have*

$$\Pr[\text{Fal-Tr} \mid \overline{\text{Diff-Adv}}] \leq \text{negl}_1(\lambda).$$

Proof. Note that the event Fal-Tr occurs if and only if the trace algorithm outputs an identity that was not queried by the adversary or outputs an identity which was queried but also revoked. Since, the tracing scheme is two phased, the false tracing can happen in any of the two stages. First, during the index-tracing procedure, it can happen that $\exists (i, p, q) \in T_{\text{index}}$, such that, $i \notin S_{\text{index}} \setminus L_{\text{index}}$ and secondly, during the identity-tracing procedure, the ID.Trace algorithm outputs an incorrect identity corresponding to some $i \in S_{\text{index}} \setminus L_{\text{index}}$. So, we have

$$\begin{aligned}
&\Pr[\text{Fal-Tr} \mid \overline{\text{Diff-Adv}}] \\
&\leq \sum_{i \in [n_{\text{ind}}]} \Pr[\text{Fal-Tr} \wedge i \notin S_{\text{index}} \setminus L_{\text{index}} \wedge (\exists p, q : (i, p, q) \in T_{\text{index}}) \mid \overline{\text{Diff-Adv}}] \\
&\quad + \sum_{(i, \ell) \in [n_{\text{ind}}] \times [\kappa]} \Pr[\text{Fal-Tr} \wedge \exists \text{id}, \hat{\text{id}} : (\text{id}, i) \in B \wedge (\exists p, q : (i, p, q) \in T_{\text{index}}) \\
&\quad \wedge \hat{\text{id}} \leftarrow \text{ID.Trace}(\text{tk}, 1^\gamma, m_0, m_1, (i, p, q)) \wedge \text{id}_\ell \neq \hat{\text{id}}_\ell \mid \overline{\text{Diff-Adv}}].
\end{aligned}$$

Consider the first term on RHS. If $\overline{\text{Diff-Adv}}$ happens then $\forall i \notin S_{\text{index}} \setminus L_{\text{index}}$ event $\overline{\text{Diff-Adv}_i}$ happens.

We know that $\overline{\text{Diff-Adv}_i}$ implies $p_{i, \perp} - p_{i+1, \perp} \leq \epsilon/8n_{\text{ind}}$. Also the event $(\exists p, q : (i, p, q) \in T_{\text{index}})$ implies $\hat{p}_{i, \perp} - \hat{p}_{i+1, \perp} > \epsilon/4n_{\text{ind}}$, where $\hat{p}_{j, \perp}$ is the estimated probability

for $p_{j,\perp}$.

By applying Chernoff bound, we have, for every $i \in [n_{\text{ind}}]$

$$\Pr[i \notin S_{\text{index}} \setminus L_{\text{index}} \wedge (\exists p, q : (i, p, q) \in T_{\text{index}}) \mid \overline{\text{Diff-Adv}}] \leq e^{-\lambda/24} \leq 2^{-O(\lambda)}.$$

Now, in the second term, for a fixed (i, ℓ) the probability term corresponds to the event when the ID.Trace outputs an identity $\hat{\text{id}}$ corresponding to index i , such that $\hat{\text{id}}_\ell \neq \text{id}_\ell$, where id is such that the adversary made a key query for (id, i) . So, if the event $\overline{\text{Diff-Adv}}$ happens then $\forall (\text{id}, i) \in B, \ell \in [\kappa]$, event $\overline{\text{Diff-Adv}_{i,\ell,X}}$ happens where $X = \text{lwr}$ if $\text{id}_\ell = 1$ else $X = \text{upr}$.

Thus, when we have $\text{id}_\ell = 1$, then $\overline{\text{Diff-Adv}_{i,\ell,\text{lwr}}}$ implies $p_{i,\perp} - p_{i,\ell} \leq \epsilon/16n_{\text{ind}}$.

Also, the event $(\exists p, q : (i, p, q) \in T_{\text{index}})$ implies that $\hat{p}_{i,\perp} - \hat{p}_{i+1,\perp} \geq \epsilon/4n_{\text{ind}}$ and the event $\hat{\text{id}} \leftarrow \text{ID.Trace}(\text{tk}, 1^\gamma, m_0, m_1, (i, p, q)) \wedge \hat{\text{id}}_\ell = 0$ implies $(\hat{p}_{i,\perp} + \hat{p}_{i+1,\perp} > 2\hat{p}_{i,\ell})$.

These together imply that $\hat{p}_{i,\perp} - \hat{p}_{i,\ell} > \epsilon/8n_{\text{ind}}$. Similarly, when $\text{id}_\ell = 0$ and $\hat{\text{id}}_\ell = 1$, $\overline{\text{Diff-Adv}_{i,\ell,\text{upr}}}$ implies $p_{i,\ell} - p_{i+1,\perp} \leq \epsilon/16n_{\text{ind}}$ and following the similar reasoning as above, we get $\hat{p}_{i,\perp} - \hat{p}_{i+1,\perp} \geq \epsilon/4n_{\text{ind}}$ and $\hat{p}_{i,\perp} + \hat{p}_{i+1,\perp} \leq 2\hat{p}_{i,\ell}$ which implies $\hat{p}_{i,\ell} - \hat{p}_{i+1,\perp} \geq \epsilon/8n_{\text{ind}}$.

Hence, using Chernoff bound, we get that

$$\Pr[\exists \text{id}, \hat{\text{id}} : (\text{id}, i) \in S \wedge (\exists p, q : (i, p, q) \in T_{\text{index}}) \wedge \hat{\text{id}} \leftarrow \text{ID.Trace}(\text{tk}, 1^\gamma, m_0, m_1, (i, p, q)) \wedge \hat{\text{id}} \in T \wedge \text{id}_\ell \neq \hat{\text{id}}_\ell \mid \overline{\text{Diff-Adv}}] \leq 2^{-O(\lambda)}.$$

Combining the probability of both the RHS terms, we get

$$\Pr[\text{Fal-Tr} \mid \overline{\text{Diff-Adv}}] \leq n_{\text{ind}} \cdot 2^{-O(\lambda)} + n_{\text{ind}} \cdot \kappa \cdot 2^{-O(\lambda)} = \text{negl}_1(\lambda).$$

■

Lemma 3.41. *If the underlying (secret/public key)-RPE satisfies (1-query-selective/adaptive) function hiding security property (Def. 3.17/Def. 3.16), then for every stateful PPT adversary \mathcal{A} , there exists a negligible function $\text{negl}_2(\cdot)$, such*

that $\forall \lambda \in \mathbb{N}$ and $i \in [n_{\text{ind}}]$, we have

$$\Pr[i \notin S_{\text{index}} \setminus L_{\text{index}} \wedge \text{Diff-Adv}_i] \leq \text{negl}_2(\lambda).$$

Proof. We give a proof by contradiction. Suppose there exists a PPT adversary \mathcal{A} that outputs a good decoder D along with messages m_0, m_1 and a revocation list L such that there exists $i^* \in [n_{\text{ind}}]$ for which $\Pr[i^* \notin S_{\text{index}} \setminus L_{\text{index}} \wedge \text{Diff-Adv}_{i^*}] \geq \delta$ where δ is some non-negligible function in the security parameter λ . We use this adversary \mathcal{A} to build a reduction \mathcal{B} that can break the function hiding security of the underlying RPE scheme. The reduction is as follows:

1. In the beginning, \mathcal{A} outputs $1^{n_{\text{ind}}}, 1^\kappa, L$ (\mathcal{A} will output L adaptively in the public trace setting).
2. The RPE challenger samples $(\text{RPE.mpk}, \text{RPE.msk}) \leftarrow \text{RPE.Setup}(1^\lambda)$ and sends RPE.mpk to \mathcal{B} . \mathcal{B} sets $\text{mpk} = \text{RPE.mpk}$ and forwards it to \mathcal{A} .
3. When \mathcal{A} issues a secret key query $(\text{lb}, \text{id}, j)$ (recall that \mathcal{A} can make at most one query for each index j), \mathcal{B} sends a key query to the RPE challenger on $(\text{lb}, x = (\text{id}, j))$. The challenger returns $\text{sk}_{\text{lb}, \text{id}, j}$ which \mathcal{B} forwards to \mathcal{A} .
4. In the end, \mathcal{A} outputs a decoder D and messages m_0, m_1 (and a revocation list L in the case of public tracing EITR scheme) to \mathcal{B} . \mathcal{B} then does the following:
 - a) Samples $i \leftarrow [n_{\text{ind}}] \setminus (S_{\text{index}} \setminus L_{\text{index}})$ and sets $f_0 = f[i, \perp, 0]$ and $f_1 = f[i + 1, \perp, 0]$.
 - b) Samples $b \leftarrow \{0, 1\}$ and sends (f_0, f_1, m_b) to the RPE challenger. The challenger samples $\alpha \leftarrow \{0, 1\}$, computes $\text{ct}_1 \leftarrow \text{RPE.Enc}(\text{RPE.msk}, f_\alpha, m_b, L)$ and sends ct_1 to \mathcal{B} . (In public trace setting, $\text{ct}_1 \leftarrow \text{RPE.Enc}(\text{RPE.mpk}, f_\alpha, m_b, L)$.)
 - c) \mathcal{B} samples $\beta \leftarrow \{0, 1\}$ and sends a encryption query on (f_β, m_b, L) and sets the ciphertext returned by the RPE challenger as ct_2 . (In public trace setting, \mathcal{B} computes ct_2 itself as $\text{ct}_2 \leftarrow \text{RPE.Enc}(\text{RPE.mpk}, f_\beta, m_b, L)$.)
 - d) \mathcal{B} samples $\beta \leftarrow \{0, 1\}$, computes $\text{ct}_2 \leftarrow \text{RPE.Enc}(\text{RPE.mpk}, f_\beta, m_b, L)$.
 - e) If $D(\text{ct}_1) = D(\text{ct}_2)$, sets $\alpha' = \beta$, else $\alpha' = 1 - \beta$.
 - f) \mathcal{B} returns α' to the RPE challenger.

\mathcal{B} wins the game if $\alpha = \alpha'$. Note that since $i \leftarrow [n_{\text{ind}}] \setminus (S_{\text{index}} \setminus L_{\text{index}})$, for every key query $(\text{lb}, \text{id}, j)$ that \mathcal{B} issues to the RPE challenger, either $f_0(\text{id}, j) = f_1(\text{id}, j)$ (when $i \notin S_{\text{index}}$) or $\text{lb} \in L$ (when $i \in S_{\text{index}} \cap L_{\text{index}}$). This establishes the admissibility of \mathcal{B} . Now, let us analyse the probability of \mathcal{B} winning against the RPE challenger.

Let $q_{j,b} = \Pr[D(\text{ct}) = b : b \leftarrow \{0, 1\}, \text{ct} \leftarrow \text{RPE.Enc}(\text{tk}, f[j, \perp, 0], m_b, L)]$.

Let E^D be the event that \mathcal{B} wins and E_b^D be the event that \mathcal{B} wins when it samples b as the message bit in step 4b. So, we have

$$\begin{aligned} \Pr[E_b^D] &= \frac{1}{4} \left(\Pr[E_b^D \mid \alpha = 0, \beta = 0] + \Pr[E_b^D \mid \alpha = 0, \beta = 1] \right. \\ &\quad \left. + \Pr[E_b^D \mid \alpha = 1, \beta = 0] + \Pr[E_b^D \mid \alpha = 1, \beta = 1] \right) \\ &= \frac{1}{4} \left(q_{i,b}^2 + (1 - q_{i,b})^2 + 2(q_{i,b}(1 - q_{i+1,b}) + (1 - q_{i,b})q_{i+1,b}) + q_{i+1,b}^2 + (1 - q_{i+1,b})^2 \right) \\ &= \frac{1}{2} + \frac{(q_{i,b} - q_{i+1,b})^2}{2}. \end{aligned}$$

We have assumed $\exists i^* \in [n_{\text{ind}}]$ such that $\Pr[i^* \notin S_{\text{index}} \setminus L_{\text{index}} \wedge \text{Diff-Adv}_{i^*}] \geq \delta$.

$\Pr[i = i^*] = 1/n_{\text{ind}}$. So we get $\Pr[i = i^* \wedge i^* \notin S_{\text{index}} \setminus L_{\text{index}} \wedge \text{Diff-Adv}_{i^*}] \geq \delta/n_{\text{ind}}$.

Let F be the event: $\exists i^* \in [n_{\text{ind}}]$ such that $i = i^* \wedge i^* \notin S_{\text{index}} \setminus L_{\text{index}} \wedge \text{Diff-Adv}_{i^*}$. Then when F occurs, we have $p_{i,\perp} - p_{i+1,\perp} > \epsilon/8n_{\text{ind}} \Rightarrow \exists b' \in \{0, 1\}$ s.t $q_{i,b'} - q_{i+1,b'} > \epsilon/8n_{\text{ind}}$. Also, $\Pr[E_b^D] \geq 1/2$ for $b \in \{0, 1\}$, irrespective of occurrence of F . Now,

$$\begin{aligned} \Pr[E_{b'}^D] &= \Pr[E_{b'}^D \mid F] \Pr[F] + \Pr[E_{b'}^D \mid \bar{F}] \Pr[\bar{F}] \\ &\geq \left(1/2 + \frac{\epsilon^2}{128n_{\text{ind}}^2}\right) \times (\delta/n_{\text{ind}}) + 1/2 \times (1 - \delta/n_{\text{ind}}) \\ &= 1/2 + \frac{\epsilon^2 \delta}{128n_{\text{ind}}^3}. \end{aligned}$$

Again,

$$\Pr[E^D] = \frac{\Pr[E_{b'}^D]}{2} + \frac{\Pr[E_{\bar{b}'}^D]}{2} \geq \frac{1}{2} \left(\frac{1}{2} + \frac{\epsilon^2 \delta}{128n_{\text{ind}}^3} \right) + \frac{1}{4} = \frac{1}{2} + \eta, \text{ where } \eta = \frac{\epsilon^2 \delta}{128n_{\text{ind}}^3}.$$

Thus, \mathcal{B} wins against the function-hiding security of the underlying RPE scheme with advantage $\geq \eta$, which is non-negligible for non-negligible ϵ and δ , a contradiction.

Hence the lemma follows.

■

Lemma 3.42. *If the underlying (secret/public key)-RPE scheme satisfies (1-query-selective/adaptive) function hiding security property (Def. 3.17/Def. 3.16), then for every stateful PPT adversary \mathcal{A} , there exists a negligible function $\text{negl}_3(\cdot)$, such that $\forall \lambda \in \mathbb{N}$ and $i \in [n_{\text{ind}}], \ell \in [\kappa]$, we have*

$$\Pr \left[\exists \text{id} \in \{0, 1\}^\kappa \text{ s.t. } (\text{id}, i) \in B \wedge ([\text{Diff-Adv}_{i, \ell, \text{lwr}} \wedge \text{id}_\ell = 1] \vee [\text{Diff-Adv}_{i, \ell, \text{upr}} \wedge \text{id}_\ell = 0]) \right] \leq \text{negl}_3(\lambda).$$

Proof. The proof is similar to the proof of Lemma 3.41. We show that if there exists a PPT adversary \mathcal{A} who outputs a good decoder D along with messages m_0, m_1 and a revocation list L for which there exists $i^* \in [n_{\text{ind}}], \ell^* \in [\kappa]$ such that $\Pr[\exists \text{id} \in \{0, 1\}^\kappa \text{ s.t. } (i^*, \text{id}) \in B \wedge ([\text{Diff-Adv}_{i^*, \ell^*, \text{lwr}} \wedge \text{id}_{\ell^*} = 1] \vee [\text{Diff-Adv}_{i^*, \ell^*, \text{upr}} \wedge \text{id}_{\ell^*} = 0])] \geq \delta$.

We use this adversary \mathcal{A} to build a reduction \mathcal{B} that can break the function hiding security of the underlying RPE scheme. \mathcal{B} is defined in the same way as in the proof of Lemma 3.41, except Step 4a, which is now described as follows:

\mathcal{B} samples $i \leftarrow B_{\text{index}}, \ell \in [\kappa]$ and a bit $g \leftarrow \{0, 1\}$ and sets $f_0 = f[i, \perp, 0]$ and $f_1 = f[i, \ell, 0]$, if $g = 0$, else sets $f_0 = f[i, \ell, 0]$ and $f_1 = f[i + 1, \perp, 0]$.

Analysis of \mathcal{B} 's advantage:

Let $q_{j,k,b} = \Pr[D(\text{ct}) = b : b \leftarrow \{0, 1\}, \text{ct} \leftarrow \text{RPE.Enc}(\text{tk}, f[j, k, 0], m_b, L)]$.

Let E^D be the event that \mathcal{B} wins and E_b^D be the event that \mathcal{B} wins when \mathcal{B} samples b as the message bit in step 4b. So, we have

$$\begin{aligned} \Pr[E_b^D] &= \frac{1}{8} \left(\Pr[E_b^D \mid \alpha = 0, \beta = 0, g = 0] + \Pr[E_b^D \mid \alpha = 0, \beta = 1, g = 0] \right. \\ &\quad \left. + \Pr[E_b^D \mid \alpha = 1, \beta = 0, g = 0] + \Pr[E_b^D \mid \alpha = 1, \beta = 1, g = 0] \right. \\ &\quad \left. + \Pr[E_b^D \mid \alpha = 0, \beta = 0, g = 1] + \Pr[E_b^D \mid \alpha = 0, \beta = 1, g = 1] \right) \end{aligned}$$

$$\begin{aligned}
& +\Pr[E_b^D \mid \alpha = 1, \beta = 0, g = 1] + \Pr[E_b^D \mid \alpha = 1, \beta = 1, g = 1]) \\
& = \frac{1}{8} \left(q_{i,\perp,b}^2 + (1 - q_{i,\perp,b})^2 + 2(q_{i,\perp,b}(1 - q_{i,\ell,b}) + (1 - q_{i,\perp,b})q_{i,\ell,b}) + q_{i,\ell,b}^2 + \right. \\
& \quad (1 - q_{i,\ell,b})^2 + q_{i+1,\perp,b}^2 + (1 - q_{i+1,\perp,b})^2 + 2(q_{i,\ell,b}(1 - q_{i+1,\perp,b}) + \\
& \quad \left. (1 - q_{i,\ell,b})q_{i+1,\perp,b}) + q_{i+1,\perp,b}^2 + (1 - q_{i+1,\perp,b})^2 \right) \\
& = \frac{1}{2} + \frac{(q_{i,\perp,b} - q_{i,\ell,b})^2}{4} + \frac{(q_{i,\ell,b} - q_{i+1,\perp,b})^2}{4}.
\end{aligned}$$

We have assumed that $\exists i^* \in [n_{\text{ind}}], \ell^* \in [\kappa]$ such that $\Pr[\exists \text{id} \in \{0, 1\}^\kappa \text{ s.t. } (i^*, \text{id}) \in B \wedge ([\text{Diff-Adv}_{i^*, \ell^*, \text{low}} \wedge \text{id}_{\ell^*} = 1] \vee [\text{Diff-Adv}_{i^*, \ell^*, \text{upr}} \wedge \text{id}_{\ell^*} = 0])] \geq \delta$.

Using $\Pr[i = i^* \wedge \ell = \ell^*] = 1/\kappa n_{\text{ind}}$, we get $\Pr[(i = i^* \wedge \ell = \ell^*) \wedge \exists \text{id} \in \{0, 1\}^\kappa \text{ s.t. } ((i^*, \text{id}) \in B \wedge ([\text{Diff-Adv}_{i^*, \ell^*, \text{low}} \wedge \text{id}_{\ell^*} = 1] \vee [\text{Diff-Adv}_{i^*, \ell^*, \text{upr}} \wedge \text{id}_{\ell^*} = 0]))] \geq \frac{\delta}{\kappa n_{\text{ind}}}$

Let F be the event: $\exists i^* \in [n_{\text{ind}}], \ell^* \in [\kappa]$ such that $(\exists \text{id} \in \{0, 1\}^\kappa \text{ s.t. } (i^*, \text{id}) \in B \wedge ([\text{Diff-Adv}_{i^*, \ell^*, \text{low}} \wedge \text{id}_{\ell^*} = 1] \vee [\text{Diff-Adv}_{i^*, \ell^*, \text{upr}} \wedge \text{id}_{\ell^*} = 0]))$. Then when F occurs, we have $(p_{i,\perp} - p_{i,\ell} > \epsilon/8n_{\text{ind}}) \vee (p_{i,\ell} - p_{i+1,\perp} > \epsilon/8n_{\text{ind}})$. This implies that there exists $b' \in \{0, 1\}$ s.t. $(q_{i,\perp,b'} - q_{i,\ell,b'} > \epsilon/8n_{\text{ind}}) \vee ((q_{i,\ell,b'} - q_{i+1,\perp,b'} > \epsilon/8n_{\text{ind}}))$. Also, $\Pr[E_b^D] \geq 1/2$ for $b \in \{0, 1\}$, irrespective of occurrence of F . Now,

$$\begin{aligned}
\Pr[E_{b'}^D] &= \Pr[E_{b'}^D \mid F] \Pr[F] + \Pr[E_{b'}^D \mid \bar{F}] \Pr[\bar{F}] \\
&\geq \left(1/2 + \frac{\epsilon^2}{256n_{\text{ind}}^2}\right) \times (\delta/\kappa n_{\text{ind}}) + 1/2 \times (1 - \delta/\kappa n_{\text{ind}}) \\
&= 1/2 + \frac{\epsilon^2 \delta}{256\kappa n_{\text{ind}}^3}.
\end{aligned}$$

Again,

$$\Pr[E^D] = \frac{\Pr[E_{b'}^D]}{2} + \frac{\Pr[E_{\bar{b}'}^D]}{2} \geq \frac{1}{2} \left(\frac{1}{2} + \frac{\epsilon^2 \delta}{256\kappa n_{\text{ind}}^3} \right) + \frac{1}{4} = \frac{1}{2} + \eta, \text{ where } \eta = \frac{\epsilon^2 \delta}{512n_{\text{ind}}^3 \kappa}.$$

■

Combining the result of Lemmas 3.40, 3.41 and 3.42, we get

$$\Pr\text{-Fal-Tr}_{\mathcal{A},\epsilon}(\lambda) \leq \text{negl}_1(\lambda) + n_{\text{ind}} \cdot \text{negl}_2(\lambda) + n_{\text{ind}} \cdot \kappa \cdot \text{negl}_3(\lambda) = \text{negl}(\lambda).$$

■

Correct trace guarantee. We prove that whenever an adversary outputs a good decoder, the tracing algorithm will output, with all but negligible probability, at least one valid user identity which was queried by the adversary.

Theorem 3.43. *If the underlying (secret/public key)-RPE scheme in the indexed (secret/public tracing)-EITR scheme construction satisfies (1-query-selective broadcast and very selective message/adaptive message) hiding property, then for every stateful PPT adversary \mathcal{A} for the (very selective/adaptive)-tracing game (Def. 3.26/Def. 3.25) and non-negligible function ϵ , there exists a negligible function $\text{negl}(\cdot)$ such that for all $\lambda \in \mathbb{N}$, the following holds*

$$\Pr[\text{Cor-Tr}] \geq \Pr[\text{Good-Decoder}] - \text{negl}(\lambda)$$

Proof. Let us define $p_{\text{Broadcast}}^D = \Pr[D(\text{ct}) = b : b \leftarrow \{0, 1\}, \text{ct} \leftarrow \text{RPE.Broadcast}(\text{RPE.mpk}, m_b, L)]$. Then in the secret trace setting, if the event $\text{Good-Decoder}_{\mathcal{A},\epsilon}$ occurs, then this implies $p_{\text{Broadcast}}^D \geq 1/2 + \epsilon$. Further, from the broadcast security of (secret-key) RPE, we get $p_{1,\perp}^D \geq p_{\text{Broadcast}}^D - \text{negl}_1(\lambda)$, which implies

$$p_{1,\perp}^D \geq 1/2 + \epsilon - \text{negl}_1(\lambda). \quad (3.7)$$

In public trace setting, the event $\text{Good-Decoder}_{\mathcal{A},\epsilon}$ directly implies $p_{1,\perp}^D \geq 1/2 + \epsilon$ by definition. Also, by message hiding property of the underlying (secret/public key) RPE scheme, we have

$$p_{n_{\text{ind}}+1,\perp}^D \leq 1/2 + \text{negl}_2(\lambda) \quad (3.8)$$

for some negligible function $\text{negl}_2(\cdot)$ with overwhelming probability. This is so, because

$f[n_{\text{ind}}, \perp, 0](\text{id}, i) = 0$ for all $(\text{id}, i) \in \mathcal{ID} \times [n_{\text{ind}}]$.

Combining equations (3.7) and (3.8), we get $p_{1,\perp}^D - p_{n_{\text{ind}}+1,\perp}^D > \epsilon/2$.

Let $S^{\text{ind}} = \{i \in [n_{\text{ind}}] \mid p_{i,\perp}^D - p_{i+1,\perp}^D > \epsilon/2n_{\text{ind}}\}$. Then if the event **Good-Decoder** occurs, $S^{\text{ind}} \neq \emptyset$.

By Chernoff bound, we have

$$\forall i \in S^{\text{index}}, \quad \Pr [\hat{p}_{i,\perp}^D - \hat{p}_{i+1,\perp}^D \leq \epsilon/4n_{\text{ind}}] \leq 2^{-O(\lambda)} = \text{negl}_3(\lambda)$$

for some negligible function $\text{negl}_3(\cdot)$. Here, \hat{p} denotes the estimate for p computed by tracing algorithm.

So, with all but negligible probability,

$$T_{\text{index}} \neq \emptyset \text{ and } \forall (i, p, q) \in T_{\text{index}}, p - q > \epsilon/4n_{\text{ind}}$$

where T_{index} and p, q are as defined in the **Trace** algorithm 3.9.1.

Note that the **ID.Trace** algorithm (Figure 3.9.1) takes as input $(i, p, q) \in T_{\text{index}}$ and outputs a corresponding id , where $\text{id}_\ell = 1$ if $\hat{p}_{i,\ell}^D > (p + q)/2$ else $\text{id}_\ell = 0$ for $\ell \in [\kappa]$. Then, for every $(i, p, q) \in T_{\text{index}}$, the tracing algorithm outputs an identity. Hence, if $T_{\text{index}} \neq \emptyset \Rightarrow T \neq \emptyset$, where T is defined as in the tracing algorithm. So,

$$\begin{aligned} \Pr[T \neq \emptyset] &\geq \Pr[T \neq \emptyset \wedge \text{Good-Decoder}] \\ &\geq (1 - \text{negl}_2(\lambda))\Pr[\text{Good-Decoder}] \\ &\geq \Pr[\text{Good-Decoder}] - \text{negl}(\lambda) \end{aligned}$$

for some negligible function $\text{negl}(\cdot)$. Combining this with the false trace guarantee, we have

$$\Pr[\text{Cor-Tr}] \geq \Pr[\text{Good-Decoder}] - \text{negl}(\lambda).$$

■

3.10 BOUNDED TRACE AND REVOKE WITH EMBEDDED IDENTITY

In this section we show how to construct a bounded (secret/public tracing)-EITR scheme from an indexed (secret/public tracing)-EITR scheme. The construction and security analysis in this section is an adaptation of ([113], Section 9) with appropriate modifications to incorporate the revocation list L . We present the construction and proofs for the secret trace setting primarily and also outline the differences in the public trace setting simultaneously.

3.10.1 Construction

Let $\text{Ind-TR} = (\text{Ind.Setup}, \text{Ind.KeyGen}, \text{Ind.Enc}, \text{Ind.Dec}, \text{Ind.Trace})$ be an indexed (secret/public tracing)-EITR system. Let $\sigma = (\sigma.\text{KeyGen}, \sigma.\text{Sign}, \sigma.\text{Verify})$ be a signature scheme that satisfies unforgeability with signature space $\{0, 1\}^{\ell_s}$. We let n_{bd} denote the bound on the number of key queries that the adversary can make. For our purpose, we can assume $n_{\text{bd}} \leq 2^\lambda$. We construct a bounded (secret/public tracing)-EITR scheme as follows:

$\text{Setup}(1^\lambda, 1^\kappa, n_{\text{bd}}) \rightarrow (\text{mpk}, \text{msk})$. The setup algorithm does the following:

1. Set $n_{\text{index}} = 2n_{\text{bd}}^2$.
2. For $j = 1$ to λ , sample $(\text{Ind.mpk}_j, \text{Ind.msk}_j) \leftarrow \text{Ind.Setup}(1^\lambda, 1^{\kappa'}, n_{\text{index}})$, where κ' is $\kappa + \ell_s$.
3. Sample $(\text{sig.sk}, \text{sig.vk}) \leftarrow \sigma.\text{KeyGen}(1^\lambda)$.
4. Output $\text{mpk} = (\text{sig.vk}, \{\text{Ind.mpk}_j\}_{j \in [\lambda]})$ and $\text{msk} = (\text{sig.vk}, \text{sig.sk}, \{\text{Ind.mpk}_j, \text{Ind.msk}_j\}_{j \in [\lambda]})$.

$\text{KeyGen}(\text{msk}, \text{lb}, \text{id}) \rightarrow \text{sk}_{\text{lb}, \text{id}}$. The key generation algorithm does the following:

1. Parse msk as $(\text{sig.vk}, \text{sig.sk}, \{\text{Ind.mpk}_j, \text{Ind.msk}_j\}_{j \in [\lambda]})$.
2. Compute $\sigma = \sigma.\text{Sign}(\text{sig.sk}, \text{id})$ and let $\text{id}' = (\text{id}, \sigma)$.
3. For $j = 1$ to λ , do the following:
 - a) Sample $i_j \leftarrow [2n_{\text{bd}}^2]$.

b) $\text{Ind.sk}_{\text{lb},\text{id},j} \leftarrow \text{Ind.KeyGen}(\text{Ind.msk}_j, \text{lb}, \text{id}', i_j)$.

4. Return $\text{sk}_{\text{lb},\text{id}} = \{\text{Ind.sk}_{\text{lb},\text{id},j}\}_{j \in [\lambda]}$.

$\text{Enc}(\text{mpk}, m, L) \rightarrow \text{ct}$. The encryption algorithm does the following:

1. Parse mpk as $(\text{sig.vk}, \{\text{Ind.mpk}_j\}_{j \in [\lambda]})$.
2. For $j = 1$ to $\lambda - 1$, randomly sample $r_j \leftarrow \mathcal{M}$ and set $r_\lambda = m \oplus r_1 \oplus \dots \oplus r_{\lambda-1}$.
3. For $j = 1$ to λ , compute $\text{Ind.ct}_j \leftarrow \text{Ind.Enc}(\text{Ind.mpk}_j, r_j, L)$.
4. Return $\text{ct} = \{\text{Ind.ct}_j\}_{j \in [\lambda]}$.

$\text{Dec}(\text{sk}_{\text{lb},\text{id}}, \text{ct}, L) \rightarrow m'$. The decryption algorithm does the following:

1. Parse $\text{sk}_{\text{lb},\text{id}} = \{\text{Ind.sk}_{\text{lb},\text{id},j}\}_{j \in [\lambda]}$ and $\text{ct} = \{\text{Ind.ct}_j\}_{j \in [\lambda]}$.
2. For $j = 1$ to λ , compute $r'_j = \text{Ind.Dec}(\text{Ind.sk}_{\text{lb},\text{id},j}, \text{Ind.ct}_j, L)$.
3. If any of the decryption fails then output \perp , else output $m' = r'_1 \oplus \dots \oplus r'_\lambda$.

$\text{Trace}^D(\text{tk}, y, m_0, m_1, L) \rightarrow T^{\text{final}}$. The trace algorithm uses two algorithms **Bnd-isGoodDecoder** and **Bnd-Subtrace** defined in Figures 3.10 and 3.11, respectively as subroutines and is defined as follows:

1. Parse tk as $\text{msk} = (\text{sig.vk}, \text{sig.sk}, \{\text{Ind.mpk}_j, \text{Ind.msk}_j\}_{j \in [\lambda]})$ and let $\text{Ind.tk}_j = \text{Ind.msk}_j$ for $j \in [\lambda]$. (In the public trace setting, $\text{tk} = \text{mpk}$ and $\text{Ind.tk}_j = \text{Ind.mpk}_j$).
2. Set $j = 1$.
3. Set $\text{flag} = 0$. For $\text{itr} = 1$ to $\lambda \cdot y$, do the following
 - a) Choose a random message $r \leftarrow \mathcal{M}$.
 - b) Run **Bnd-isGoodDecoder** as
 $\text{flag} \leftarrow \text{Bnd-isGoodDecoder}^D(\{\text{Ind.mpk}_j\}_{j \in [\lambda]}, 1^y, m_0, m_1, r, L, j)$.
 - c) If $\text{flag} = 1$, break. Else, continue.
4. If $\text{flag} = 1$, run **Bnd-Subtrace** as

Algorithm Bnd-isGoodDecoder^D(key, 1^y, m₀, m₁, r, L, i)

Inputs: keys key = {Ind.mpk_j}_{j∈[λ]}, parameter y, messages m₀, m₁, r, revocation list L and position-index i ∈ [λ].

Output : 0/1.

1. Set count = 0. Let $\epsilon = 1/y$.
2. For $j = 1$ to $\lambda \cdot y$:
 - Split r in $\lambda - 1$ shares. That is, randomly sample $\lambda - 1$ messages, $r_1, \dots, r_{i-1}, r_{i+1}, \dots, r_\lambda$ such that $\bigoplus_{k \in [\lambda] \setminus \{i\}} r_k = r$.
 - Sample $b \leftarrow \{0, 1\}$. For $k \in [\lambda] \setminus \{i\}$, compute $\text{Ind.ct}_k \leftarrow \text{Ind.Enc}(\text{Ind.mpk}_k, r_k, L)$. Compute $\text{Ind.ct}_i \leftarrow \text{Ind.Enc}(\text{Ind.mpk}_i, r \oplus m_b, L)$.
 - Query D on $\text{ct} = (\text{Ind.ct}_1, \dots, \text{Ind.ct}_\lambda)$. Let b' be D 's response.
 - If $b' = b$, set count = count + 1.
3. If $\text{count}/(\lambda \cdot y) \geq 1/2 + \epsilon/3$, then output 1, else output 0.

Figure 3.10: Algorithm Bnd-isGoodDecoder

$T \leftarrow \text{Bnd-Subtrace}^D(\{\text{Ind.mpk}_j, \text{Ind.tk}_j\}_{j \in [\lambda]}, 1^y, m_0, m_1, r, L, j)$. Else, set $T = \phi$.

5. If $T = \phi$ and $j < \lambda$, set $j = j + 1$ and go to step 3. Otherwise do the following:

- Set $T^{\text{temp}} = \phi$. For each $\text{id}' = (\text{id}, \sigma) \in T$, if $\sigma.\text{Verify}(\text{sig.vk}, \text{id}, \sigma) = 1$, add id to T^{temp} . Concretely,

$$T^{\text{temp}} = \{\text{id} : \exists \sigma \text{ s.t. } (\text{id}, \sigma) \in T \text{ and } \sigma.\text{Verify}(\text{sig.vk}, \text{id}, \sigma) = 1\}.$$

- Recall the function $\text{map} : \mathcal{ID} \rightarrow \mathcal{L}$, that maps a given identity id to its corresponding label lb (Remark 7). For each $\text{id} \in T^{\text{temp}}$, if $\text{map}(\text{id}) \in L$, then set $T^{\text{temp}} = T^{\text{temp}} \setminus \{\text{id}\}$.
- Set $T^{\text{final}} = T^{\text{temp}}$.
- If $T^{\text{final}} = \phi$ and $j < \lambda$, set $j = j + 1$ and go to step 3. Otherwise exit and return T^{final} .

Algorithm $\text{Bnd-Subtrace}^D(\text{key}, 1^y, m_0, m_1, r, L, i)$

Inputs: keys $\text{key} = \{\text{Ind.mpk}_j, \text{Ind.tk}_j\}_{j \in [\lambda]}$, parameter y , messages m_0, m_1, r , revocation list L and index-position $i \in [\lambda]$.

Output : $T \subseteq \{0, 1\}^\kappa$

1. Define oracle $\tilde{D}[\{\text{Ind.mpk}_j\}_{j \in [\lambda]}, r, L, i]$ as in Figure 3.12.
2. Output $T \leftarrow \text{Ind.Trace}^{\tilde{D}}(\text{Ind.tk}_i, 4y, m_0 \oplus r, m_1 \oplus r, L)$.

Figure 3.11: Algorithm Bnd-Subtrace

Algorithm $\tilde{D}^D[\text{key}, r, L, i]$

Hardwired values: keys $\text{key} = \{\text{Ind.mpk}_j\}_{j \in [\lambda]}$, message r , revocation list L and index-position $i \in [\lambda]$.

Inputs: Ind.ct .

Output : 0/1

Upon input Ind.ct , the \tilde{D} oracle does the following:

- Shares r in $\lambda - 1$ components as follows: it chooses $\lambda - 1$ random messages r_k for $k \in [\lambda] \setminus \{i\}$, such that $\oplus_{k \in [\lambda] \setminus \{i\}} r_k = r$.
- For $k \in [\lambda] \setminus \{i\}$, computes $\text{Ind.ct}_k = \text{Ind.Enc}(\text{Ind.mpk}_k, r_k, L)$.
- Sets $\text{ct}_{\text{bd}} = (\text{Ind.ct}_1, \dots, \text{Ind.ct}_{i-1}, \text{Ind.ct}, \text{Ind.ct}_{i+1}, \dots, \text{Ind.ct}_\lambda)$.
- Queries oracle D as $b' \leftarrow D(\text{ct}_{\text{bd}})$.
- Outputs b' .

Figure 3.12: Oracle \tilde{D}

Correctness. We prove that the above construction of bounded (secret/public tracing)-EITR scheme satisfies correctness (Def. 3.22) via the following theorem.

Theorem 3.44. *Assume Ind-TR is a correct indexed (secret/public tracing)-EITR scheme then the above construction of bounded (secret/public tracing)-EITR scheme is correct.*

Proof. We have, as per the construction, that any message m is split in λ components r_1, \dots, r_λ such that $r_1 \oplus \dots \oplus r_\lambda = m$. Then, from the correctness of Ind-TR, we get $r'_k = r_k$ for all $k \in [\lambda]$, where $r'_k \leftarrow \text{Ind.Dec}(\text{Ind.sk}_{\text{lb}, \text{id}, k}, \text{Ind.ct}_k, L)$, as long as $\text{lb} \notin L$. Hence, the decryption algorithm correctly outputs m as $r_1 \oplus \dots \oplus r_\lambda$. ■

Efficiency. We can instantiate the above construction by the indexed public/secret tracing-EITR scheme in Sec. 3.9.1. The above construction is basically λ times repetition of the underlying indexed EITR scheme. Additional overhead is induced due to the usage of the signature scheme, where identity becomes longer by $\ell_s = \text{poly}(\lambda)$ bit and the master public key is longer by $|\text{sig.vk}| = \text{poly}(\lambda)$ bit.²⁶ These changes do not alter the dependency on $|\text{id}|$ and $|\text{lb}|$ of the parameter size. Therefore, the parameter size is as follows.

Secret Tracing Setting. In the secret tracing setting, we have $|\text{mpk}|, |\text{ct}|, |\text{sk}| = \text{poly}(\lambda, |\text{id}|, |\text{lb}|)$.

Public Tracing Setting. In the public tracing setting, we have $|\text{mpk}|, |\text{ct}| = \text{poly}(\lambda, |\text{lb}|), |\text{sk}| = \text{poly}(\lambda, |\text{id}|, |\text{lb}|)$.

3.10.2 Security

In this section, we prove that our construction of bounded (secret/public tracing)-EITR scheme is secure.

²⁶Parameter sizes of most signature schemes depend on the length of the message space, which is $|\text{id}|$ in our case. However, this dependency can be removed by hashing the message before signing using the collision resistant hash functions.

IND-CPA security.

Theorem 3.45. *If Ind-TR is a (very selective/adaptive)-ind-cpa secure indexed (secret/public tracing)-EITR scheme, then the above construction of bounded (secret/public tracing)-EITR scheme is (very selective/adaptive)-ind-cpa secure.*

Proof. We show that if there exists a PPT adversary that breaks the ind-cpa security of the bounded EITR scheme, then we can use \mathcal{A} to build a PPT algorithm \mathcal{B} that breaks ind-cpa security of the underlying Ind-TR scheme. The reduction is as follows:

1. \mathcal{B} gets $1^\kappa, 1^{n_{\text{bd}}}, L$ and key queries $\{(\text{lb}_i, \text{id}_i)\}_{i \in [Q]}$, where Q is the number of key queries issued, from the adversary \mathcal{A} . (In the public trace setting, \mathcal{A} can make adaptive key queries and outputs L adaptively along with the challenge messages in Step 6).
2. \mathcal{B} generates $(\text{sig.sk}, \text{sig.vk}) \leftarrow \sigma.\text{KeyGen}(1^\lambda)$.
3. For the i -th key query $(\text{lb}_i, \text{id}_i)$, $i \in [Q]$, \mathcal{B} samples $i_1 \leftarrow [n_{\text{index}}]$. It sets $n_{\text{index}} = 2n_{\text{bd}}^2$ and sends $1^{\kappa+\ell_s}, 1^{n_{\text{index}}}, L$ and $\{(\text{lb}_i, \text{id}'_i, i_1)\}_{i \in [Q]}$, where id'_i is computed as in the construction, to the Ind-TR challenger. The Ind-TR challenger returns Ind.mpk and $\{\text{Ind.sk}_{\text{lb}_i, \text{id}_i}\}_{i \in [Q]}$ where $\text{Ind.sk}_{\text{lb}_i, \text{id}_i} \leftarrow \text{Ind.KeyGen}(\text{Ind.msk}, \text{lb}_i, \text{id}'_i, i_1)$.
4. \mathcal{B} sets $\text{Ind.mpk}_1 = \text{Ind.mpk}$, generates $(\text{Ind.mpk}_j, \text{Ind.msk}_j) \leftarrow \text{Ind.Setup}(1^\lambda, 1^{\kappa+\ell_s}, n_{\text{index}})$ for $j \in \{2, \dots, \lambda\}$. It sends $\text{mpk} = (\text{sig.vk}, \{\text{Ind.mpk}_j\}_{j \in [\lambda]})$ to \mathcal{A} .
5. For each key query $(\text{lb}_i, \text{id}_i)$, \mathcal{B} sets $\text{Ind.sk}_{\text{lb}_i, \text{id}_i, i_1} = \text{Ind.sk}_{\text{lb}_i, \text{id}_i}$ and computes $\text{Ind.sk}_{\text{lb}_i, \text{id}_i, i_j}$ for $j \in \{2, \dots, \lambda\}$ itself and returns $\text{sk}_{\text{lb}, \text{id}} = \{\text{Ind.sk}_{\text{lb}_i, \text{id}_i, i_j}\}_{j \in [\lambda]}$ to \mathcal{A} .
6. When \mathcal{A} sends the challenge messages (m_0, m_1) for the challenge query, \mathcal{B} samples $r_j \leftarrow \mathcal{M}$ for $j \in \{2, \dots, \lambda\}$, sets $m'_b = \bigoplus_{j>1} r_j \oplus m_b$, sends (m'_0, m'_1) for the challenge query to the Ind-TR challenger and gets back Ind.ct . \mathcal{B} sets $\text{ct}_1 = \text{Ind.ct}$, computes $\text{ct}_j \leftarrow \text{Ind.Enc}(\text{Ind.mpk}_j, r_j, L)$ for $j \in \{2, \dots, \lambda\}$ and sends $\text{ct} = \{\text{ct}_j\}_{j \in [\lambda]}$ to \mathcal{A} .
7. \mathcal{A} outputs a bit b' , \mathcal{B} returns b' to the challenger.

Observe that \mathcal{B} issues key queries to the Ind-TR challenger only when there is a key query from \mathcal{A} . From the admissibility of \mathcal{A} , for any key query of the form $(\text{lb}_i, \text{id}_i, i_1)$ that \mathcal{B} issues $\text{lb} \in L$. So, \mathcal{A} wins the ind-cpa security game of the BD-TR scheme

with advantage ϵ , then \mathcal{B} also wins the ind-cpa security game of the underlying Ind-TR scheme with the same advantage. \blacksquare

Secure Tracing Analysis. Now we show that our construction satisfies false trace and correct trace guarantees.

False Trace Guarantee. First, we show that the probability of false trace in our scheme is negligible in the security parameter via the following theorem.

Theorem 3.46. *Assume that σ is unforgeable, then our construction of bounded (secret/public tracing)-EITR scheme satisfies the (very selective/adaptive) false trace guarantee (Def. 3.27/Def. 3.28), even if the adversary makes unbounded polynomial number of key queries.*

Proof. Let us first setup some notations. For $\text{lb} \in \mathcal{L}$, $\text{id} \in \mathcal{ID}$ and a revocation list L , let

$$\begin{aligned} S_{\mathcal{ID}} &= \{\text{id} : (\text{lb}, \text{id}) \in S\}, \\ L_{\mathcal{ID}} &= \{\text{id} : (\text{lb}, \text{id}) \in S \wedge \text{lb} \in L\}, \\ T_{\text{lb}}^{\text{final}} &= \{\text{map}(\text{id}) : \text{id} \in T^{\text{final}}\} \end{aligned}$$

where map is as defined in Remark 7 and T^{final} is as in the construction.

False trace happens when $T^{\text{final}} \not\subseteq S_{\mathcal{ID}}$ or $T_{\text{lb}}^{\text{final}} \cap L \neq \emptyset$. Observe that $T_{\text{lb}}^{\text{final}} \cap L = \emptyset$ by definition. So, all we need to argue is that $T^{\text{final}} \subseteq S_{\mathcal{ID}}$.

Recall that the tracing algorithm uses Bnd-Subtrace algorithm to find a set $T = \{(\text{id}_k, \sigma_k)\}_k$. Identity id_k is added to the set T^{temp} and then to set T^{final} only if $\sigma.\text{Verify}(\text{sig.vk}, \text{id}_k, \sigma_k) = 1$. Next we show that if there is an adversary \mathcal{A} who outputs a decoder D along with messages m_0, m_1 and revocation list L such that the tracing algorithm outputs T^{final} where $T^{\text{final}} \not\subseteq S_{\mathcal{ID}}$, then there exists a reduction \mathcal{B} against

unforgeability of signature scheme σ . The reduction is defined as follows:

Upon receiving sig.vk from the σ challenger, \mathcal{B} does the following:

1. Run \mathcal{A} on input 1^λ to obtain $1^\kappa, 1^{n_{\text{bd}}}, L$ and key queries $\{(\text{lb}_i, \text{id}_i)\}_{i \in [Q]}$, where Q is the number of key queries issued. (In the public trace setting, \mathcal{A} can make adaptive key queries and output L adaptively along with the challenge messages).
2. Samples $(\text{Ind.mpk}_j, \text{Ind.msk}_j) \leftarrow \text{Ind.Setup}(1^\lambda, 1^{\kappa'}, n_{\text{index}})$ for $j \in [\lambda]$ and sends $\text{mpk} = (\text{sig.vk}, \{\text{mpk}_j\}_{j \in [\lambda]})$ to \mathcal{A} .
3. For each key query $(\text{lb}_i, \text{id}_i)$, $i \in [Q]$, \mathcal{B} sends id_i to the σ challenger for signature. The σ challenger returns $\sigma_i = \sigma(\text{sig.sk}, \text{id}_i)$. \mathcal{B} generates the secret key $\text{sk}_{\text{lb}_i, \text{id}_i}$ using $\text{id}'_i = (\text{id}_i, \sigma_i)$ as in the construction and sends $\text{sk}_{\text{lb}_i, \text{id}_i}$ to \mathcal{A} .
4. In the end, \mathcal{A} outputs a decoder D , messages m_0, m_1 .
5. \mathcal{B} runs the trace algorithm with the help of the decoder D and gets a set T^{final} .
6. If there exists an $\text{id}^* \in T^{\text{final}}$ such that $\text{id}^* \notin S_{\text{ID}}$, then \mathcal{B} returns a forgery for id^* to σ challenger as follows: $\text{id}^* \in T^{\text{final}}$ implies that there exists σ^* , such that $(\text{id}^*, \sigma^*) \in T$, and $\sigma.\text{Verify}(\text{sig.vk}, \text{id}^*, \sigma^*) = 1$. \mathcal{B} returns (id^*, σ^*) as a forgery to the σ challenger.

Note that since $\text{id}^* \notin S_{\text{ID}}$, \mathcal{B} must not have queried a signature on id^* to the σ challenger and hence (id^*, σ^*) is a valid forgery. If the false tracing happens with non-negligible probability ϵ , then \mathcal{B} also wins the unforgeability game with probability ϵ .

■

Correct Trace Guarantee. Recall that in the experiment for correct tracing, the adversary first sends $(1^\kappa, 1^{n_{\text{bd}}})$, a revocation list L (in the public trace setting, it outputs L adaptively), and at most n_{bd} key queries. Let $S = \{(\text{lb}, \text{id})\}$ be the set of label-identity pairs for which the adversary issues key queries. Next, the challenger sends the public key to the adversary. At the end, the adversary outputs a decoder box D along with two messages m_0 and m_1 . If D is a good decoder then for correct tracing we want that the tracing algorithm outputs non empty set of traitors $T^{\text{final}} \subseteq S_{\text{ID}} \setminus L_{\text{ID}}$. We prove the correct trace guarantee of our scheme via the following theorem.

Theorem 3.47. *Assume that the underlying (secret/public tracing)-Ind-TR scheme satisfies (very selective/adaptive) correct trace guarantee (Def. 3.26 /Def. 3.25) and let n_{bd} be the bound on the number of key queries for an admissible adversary, then our bounded (secret/public tracing)-EITR scheme also satisfies the (very selective/adaptive) correct trace guarantee (Def. 3.28 /Def. 3.27).*

Proof. Similar to [113] we begin with defining some events of interest. We modify the definition of some events to take into account the constraint that $T^{\text{final}} \subseteq S_{\mathcal{ID}} \setminus L_{\mathcal{ID}}$. We drop the subscripts \mathcal{A} , ϵ and security parameter λ in the following to keep the notations simple.

- **Event Admissible-Adversary:** It is defined as the event that the adversary \mathcal{A} makes at most n_{bd} key queries.
- **Event Tr (Tracing without correctness):** Similar to Corr-Tr, except that we don't need $T^{\text{final}} \subseteq S_{\mathcal{ID}} \setminus L_{\mathcal{ID}}$, i.e., Tr is the event that $T^{\text{final}} \neq \emptyset$ occurs. Denote $\text{Pr-Tr} := \text{Pr}[\text{Tr}]$.
- **Event Dist-Indx (Position with distinct indices for each key):** Defined as the event that there exists $i \in [\lambda]$ such that the i -th index of each key is distinct.
- **Event Dist-Indx _{i} :** It is defined as the event that $i \in [\lambda]$ is the first position such that the i -th index of each key is distinct. By definition, Dist-Indx _{i} are disjoint events for all $i \in [\lambda]$ and $\cup_{i \in [\lambda]} \text{Dist-Indx}_i = \text{Dist-Indx}$.
- **Event Tr _{i} (Tracing without correctness in i -th iteration):** Let T_i denote the set of (identity, signature) pairs traced in the i -th iteration. The event Tr _{i} happens if T_i is non-empty.

- Event Corr-Tr-Sig_i (Tracing with same signature as that received in key): The event that T_i is not empty and for all $(\text{id}, \sigma) \in T_i$, (lb, id) was queried for a key and $\text{lb} \notin L$, i.e. $\text{id} \in S_{\mathcal{ID}} \setminus L_{\mathcal{ID}}$ and that the key generation oracle output $\text{Ind.sk}_{\text{lb}, \text{id}, i} \leftarrow \text{Ind.KeyGen}(\text{Ind.msk}_i, \text{lb}, (\text{id}, \sigma), j)$ for some j .
- Event $\text{Found-Good-}r_i$: This event occurs if flag is set to 1 in the i -th iteration of the tracing algorithm.
- Event $\text{Good-}\tilde{D}_i$ (Good decoder \tilde{D} during the Bnd-Subtrace routine execution in i -th iteration): It is defined as the event that in the i -th iteration, the execution reaches step 3 (that is, it found a ‘good’ r in the i -th iteration), and the decoder \tilde{D} constructed is an $\epsilon/4$ good decoder for distinguishing messages $m_0 \oplus r$ and $m_1 \oplus r$.

Note that if no good r is found in step 3, $\text{Good-}\tilde{D}_i$ is said to not have happened. With the above events, the correctness of tracing is argued via the following series of inequalities:

$$\text{Pr-Corr-Tr}(\lambda) \geq \text{Pr-Tr} - \text{negl} \quad (3.9)$$

$$\geq \text{Pr}[\text{Tr} \wedge \text{Dist-Indx}] - \text{negl} \quad (3.10)$$

$$= \sum_{i \in [\lambda]} \text{Pr}[\text{Tr} \wedge \text{Dist-Indx}_i] - \text{negl} \quad (3.11)$$

$$\geq \sum_{i \in [\lambda]} \text{Pr}[\text{Corr-Tr-Sig}_i \wedge \text{Dist-Indx}_i] - \text{negl} \quad (3.12)$$

$$\geq \sum_{i \in [\lambda]} \text{Pr}[\text{Good-}\tilde{D}_i \wedge \text{Admissible-Adversary} \wedge \text{Dist-Indx}_i] - \text{negl} \quad (3.13)$$

$$\geq \sum_{i \in [\lambda]} \text{Pr}[\text{Good-}\tilde{D}_i \wedge \text{Found-Good-}r_i \wedge \text{Good-Decoder} \wedge \text{Admissible-Adversary} \wedge \text{Dist-Indx}_i] - \text{negl} \quad (3.14)$$

$$\geq \sum_{i \in [\lambda]} \text{Pr}[\text{Found-Good-}r_i \wedge \text{Good-Decoder} \wedge \text{Admissible-Adversary}]$$

$$\wedge \text{Dist-Indx}_i] - \text{negl} \quad (3.15)$$

$$\geq \sum_{i \in [\lambda]} \Pr[\text{Good-Decoder} \wedge \text{Admissible-Adversary} \wedge \text{Dist-Indx}_i] - \text{negl} \quad (3.16)$$

$$= \Pr[\text{Good-Decoder} \wedge \text{Admissible-Adversary} \wedge \text{Dist-Indx}] - \text{negl} \quad (3.17)$$

$$\geq \Pr[\text{Good-Decoder} \wedge \text{Admissible-Adversary}] - \text{negl} \quad (3.18)$$

Explanation for each of the inequalities is the same as in [113] and is omitted. Here, we argue the transitions from equations (3.11) to (3.12) and (3.12) to (3.13) only, since our definition of event Corr-Tr-Sig_i is slightly modified.

- Transition from (3.11) to (3.12) follows from the observation that the event Corr-Tr-Sig_i implies the event Tr because by definition, if Corr-Tr-Sig_i happens then T_i is non empty and for each (id, σ) pair in T_i , $\sigma.\text{Verify}(\text{sig.vk}, \text{id}, \sigma)$ verifies and $\text{id} \in S_{\mathcal{ID}} \setminus L_{\mathcal{ID}}$. Hence, the set of traitors T^{final} obtained from T_i will also be non empty.
- Transition from (3.12) to (3.13) is argued via the following claim:

Claim 3.48. Assume that the underlying Ind-TR scheme satisfies correct trace guarantee.

Then for all $i \in [\lambda]$

$$\Pr[\text{Corr-Tr-Sig}_i \wedge \text{Dist-Indx}_i] \geq \Pr[\text{Good-}\tilde{\text{D}}_i \wedge \text{Admissible-Adversary} \wedge \text{Dist-Indx}_i] - \text{negl}.$$

Proof. We prove the claim by contradiction. We assume that there exists an adversary \mathcal{A} who outputs a good decoder D along with messages m_0, m_1 and a revocation list L such that $\Pr[\text{Good-}\tilde{\text{D}}_i \wedge \text{Admissible-Adversary} \wedge \text{Dist-Indx}_i] - \Pr[\text{Corr-Tr-Sig}_i \wedge \text{Dist-Indx}_i]$ is non-negligible. Then we use \mathcal{A} to build an adversary \mathcal{B} against correct trace guarantee of Ind-TR, defined as follows:

1. Run \mathcal{A} on input 1^λ to obtain $1^\kappa, 1^{n_{\text{bd}}}, L$ and key queries $\{(\text{lb}_k, \text{id}_k)\}_{k \in [Q]}$, where Q is the number of key queries issued. (In the public trace setting, \mathcal{A} can make adaptive key queries and output L adaptively along with the challenge messages). It sets $n_{\text{index}} = 2n_{\text{bd}}^2$.

2. Samples $(\text{sig.sk}, \text{sig.vk}) \leftarrow \sigma.\text{KeyGen}(1^\lambda)$, computes $\sigma_k \leftarrow \sigma.\text{Sign}(\text{sig.sk}, \text{id}_k)$ and sets $\text{id}'_k = (\text{id}_k, \sigma_k)$ for $k \in [Q]$.
3. For each $k \in [Q]$ and $j \in [\lambda]$, it randomly samples $k_j \leftarrow n_{\text{index}}$. It sends $1^{\kappa+\ell_s}, 1^{n_{\text{index}}}, L$ and key queries $\{(\text{lb}_k, \text{id}_k, k_i)\}_{k \in [Q]}$ to the Ind-TR challenger and gets back Ind.mpk and $\{\text{Ind.sk}_{\text{lb}_k, \text{id}_k, k_i}\}_{k \in [Q]}$.
4. \mathcal{B} sets $\text{Ind.mpk}_i = \text{Ind.mpk}$ and does the following:
 - For $j \in [\lambda] \setminus \{i\}$, $(\text{Ind.mpk}_k, \text{Ind.msk}_k) \leftarrow \text{Ind.Setup}(1^{\kappa+\ell_s}, 1^{n_{\text{index}}})$.
 - Set $\text{mpk} = (\text{sig.vk}, \{\text{Ind.mpk}_k\}_{k \in [\lambda]})$ and sends mpk to \mathcal{A} .
5. For each key query $(\text{lb}_k, \text{id}_k)$ by \mathcal{A} , \mathcal{B} does the following:
 - For $j \in [\lambda] \setminus \{i\}$, $\text{Ind.sk}_{\text{lb}_k, \text{id}_k, k_j} \leftarrow \text{Ind.KeyGen}(\text{Ind.msk}_k, \text{lb}_k, \text{id}'_k, k_j)$.
 - Sends $\text{sk}_{\text{lb}, \text{id}} = (\text{Ind.sk}_{\text{lb}_k, \text{id}_k, k_1}, \dots, \text{Ind.sk}_{\text{lb}_k, \text{id}_k, k_\lambda})$ to \mathcal{A} .
6. If i is not the first position for which the i -th position indices (k_i 's) are distinct for all the key queries, then \mathcal{B} outputs a random decoder and quits the game.
7. In the end, \mathcal{A} outputs a decoder D along with messages m_0, m_1 .
8. \mathcal{B} runs $\text{Bnd-isGoodDecoder}(\{\text{Ind.mpk}_j\}_{j \in [\lambda]}, 1^y, m_0, m_1, L, r)$ for uniformly and independently sampled r until it finds a r for which Bnd-isGoodDecoder algorithm outputs 1. If Bnd-isGoodDecoder algorithm does not output 1 even after $\lambda \cdot y$ attempts, then \mathcal{B} outputs a random decoder and quits.
9. \mathcal{B} constructs decoder \tilde{D} as defined in Figure 3.12 and sets $\tilde{m}_b = m_b \oplus r$ for $b \in \{0, 1\}$ and sends $(\tilde{D}, \tilde{m}_0, \tilde{m}_1, L)$ to the Ind-TR challenger.

Now let us analyze the probability that \mathcal{B} outputs a $1/4y$ good decoder box. This happens if (i) i is the first position such that the i -th position indices are different for all the key queries (ii) Bnd-isGoodDecoder outputs 1 for some r and \tilde{D} is $\epsilon/4 = 1/4y$ good decoder for distinguishing $m_0 \oplus r, m_1 \oplus r$. First event is same as Dist-Indx_i and second event is same as $\text{Good-}\tilde{D}_i$. Hence, the probability that \mathcal{B} outputs a $1/4y$ good decoder box is $\Pr[\text{Good-}\tilde{D}_i \wedge \text{Dist-Indx}_i]$. Then, by correct trace guarantee of Ind-TR,

$$\begin{aligned} \Pr[\text{Corr-Tr-Sig}_i \wedge \text{Dist-Indx}_i] &\geq \Pr[\text{Good-}\tilde{D}_i \wedge \text{Dist-Indx}_i] - \text{negl} \\ &\geq \Pr[\text{Good-}\tilde{D}_i \wedge \text{Admissible-Adversary} \wedge \text{Dist-Indx}_i] - \text{negl}. \end{aligned}$$

This contradicts our assumption that $\Pr[\text{Good-}\tilde{D}_i \wedge \text{Admissible-Adversary} \wedge \text{Dist-Indx}_i] - \Pr[\text{Corr-Tr-Sig}_i \wedge \text{Dist-Indx}_i]$ is non-negligible, hence the proof. ■

■

3.11 UNBOUNDED TRACE AND REVOKE WITH EMBEDDED IDENTITIES

In this section we show how to construct unbounded (secret/public tracing)-EITR scheme from a bounded (secret/public tracing)-EITR scheme. The transformation technique and the security analysis in this section is adapted from [113] with modifications to incorporate the revocation list. We present the construction and proofs for *secret-key* trace setting primarily and also outline the differences in the public-key trace setting simultaneously.

3.11.1 Construction

Let $\text{BD-TR} = (\text{BD.Setup}, \text{BD.KeyGen}, \text{BD.Enc}, \text{BD.Dec}, \text{BD.Trace})$ be a bounded (secret/public tracing)-EITR scheme for identity space $\mathcal{ID} = \{0, 1\}^\kappa$. We construct an unbounded (secret/public tracing)-EITR scheme as follows.

$\text{Setup}(1^\lambda, 1^\kappa) \rightarrow (\text{mpk}, \text{msk})$. The setup algorithm does the following:

1. For $j = 1$ to λ , sample $(\text{BD.mpk}_j, \text{BD.msk}_j) \leftarrow \text{BD.Setup}(1^\lambda, 1^\kappa, n_{\text{bd}} = 2^j)$.
2. Output $\text{mpk} = \{\text{BD.mpk}_j\}_{j \in [\lambda]}$ and $\text{msk} = \{\text{BD.mpk}_j, \text{BD.msk}_j\}_{j \in [\lambda]}$.

$\text{KeyGen}(\text{msk}, \text{lb}, \text{id}) \rightarrow \text{sk}_{\text{lb}, \text{id}}$. The KeyGen algorithm does the following:

1. Parse $\text{msk} = \{\text{BD.mpk}_j, \text{BD.msk}_j\}_{j \in [\lambda]}$.
2. For $j = 1$ to λ , it computes $\text{BD.sk}_j \leftarrow \text{BD.KeyGen}(\text{BD.msk}_j, \text{lb}, \text{id})$.
3. Returns $\text{sk}_{\text{lb}, \text{id}} = \{\text{BD.sk}_j\}_{j \in [\lambda]}$.

$\text{Enc}(\text{mpk}, m, L) \rightarrow \text{ct}$. The encryption algorithm does the following:

1. Parse $\text{mpk} = \{\text{BD.mpk}_j\}_{j \in [\lambda]}$.
2. Secret share m in λ shares as follows. For $j = 1$ to $\lambda - 1$, randomly sample $r_j \leftarrow \mathcal{M}$ and set $r_\lambda = m \oplus r_1 \oplus \dots \oplus r_{\lambda-1}$.
3. For $j = 1$ to λ , compute $\text{BD.ct}_j = \text{BD.Enc}(\text{BD.mpk}_j, r_j, L)$.
4. Output $\text{ct} = \{\text{BD.ct}_j\}_{j \in [\lambda]}$.

$\text{Dec}(\text{sk}_{\text{lb}, \text{id}}, \text{ct}, L) \rightarrow m'$. The decryption algorithm does the following:

1. Parse $\text{sk}_{\text{lb}, \text{id}} = \{\text{BD.sk}_j\}_{j \in [\lambda]}$ and $\text{ct} = \{\text{BD.ct}_j\}_{j \in [\lambda]}$.
2. For $j = 1$ to λ , compute $r'_j = \text{BD.Dec}(\text{BD.sk}_j, \text{BD.ct}_j, L)$.
3. If any of the decryption fails then output $m' = \perp$, else output $m' = \bigoplus_{j \in [\lambda]} r'_j$.

$\text{Trace}^D(\text{tk}, y, \text{Q}_{\text{bd}}, m_0, m_1, L) \rightarrow T$. The trace algorithm uses two algorithms `isGoodDecoder` and `SubTrace` defined in Figures 3.13 and 3.14, respectively as subroutines. The tracing algorithm is as follows.

1. Parse tk as $\text{msk} = \{\text{BD.mpk}_j, \text{BD.msk}_j\}_{j \in [\lambda]}$ and let $\text{BD.tk}_j = \text{BD.msk}_j$ for $j \in [\lambda]$. (For the public trace setting $\text{tk} = \text{mpk}$ and $\text{BD.tk}_j = \text{BD.mpk}_j$)
2. Set $j = \lceil \log \text{Q}_{\text{bd}} \rceil$.
3. Set $\text{flag} = 0$. For $\text{itr} = 1$ to $\lambda \cdot y$, do the following
 - a) Choose a random message $r \leftarrow \mathcal{M}$.
 - b) Run `isGoodDecoder` as
 $\text{flag} \leftarrow \text{isGoodDecoder}^D(\{\text{BD.mpk}_j\}_{j \in [\lambda]}, 1^y, m_0, m_1, r, L, j)$
 - c) If $\text{flag} = 1$, break. Else, continue.
4. If $\text{flag} = 1$, run $T \leftarrow \text{SubTrace}^D(\{\text{BD.mpk}_j, \text{BD.tk}_j\}_{j \in [\lambda]}, 1^y, m_0, m_1, r, L, j)$. Else, set $T = \phi$.
5. Output T .

Correctness. We show that the above construction of bounded (secret/public tracing)-EITR satisfies correctness (Def. 3.22) via the following theorem.

Theorem 3.49. *Assume BD-TR is a correct bounded (secret/public tracing)-EITR*

Algorithm $\text{isGoodDecoder}^D(\text{key}, 1^y, m_0, m_1, r, L, i)$

Inputs: keys $\text{key} = \{\text{BD.mpk}_j\}_{j \in [\lambda]}$, parameter y , messages m_0, m_1, r , revocation list L and a position-index $i \in [\lambda]$.

Output : 0/1.

1. Set $\text{count} = 0$. Let $\epsilon = 1/y$.
2. For $j = 1$ to $\lambda \cdot y$:
 - Sample $\lambda - 1$ messages, $r_1, \dots, r_{i-1}, r_{i+1}, \dots, r_\lambda$ randomly such that $\bigoplus_{k \in [\lambda] \setminus \{i\}} r_k = r$.
 - Sample $b \leftarrow \{0, 1\}$, and compute ciphertexts $\text{BD.ct}_k = \text{BD.Enc}(\text{BD.mpk}_k, r_k, L)$ for $k \in [\lambda] \setminus \{i\}$ and $\text{BD.ct}_i = \text{BD.Enc}(\text{BD.mpk}_i, r \oplus m_b, L)$.
 - Query D on $\text{ct} = (\text{BD.ct}_1, \dots, \text{BD.ct}_\lambda)$. Let b' be the response of D .
 - If $b' = b$, set $\text{count} = \text{count} + 1$.
3. If $\text{count}/(\lambda \cdot y) \geq 1/2 + \epsilon/3$, then output 1, else output 0.

Figure 3.13: Algorithm isGoodDecoder for Unbounded EITR

Algorithm $\text{SubTrace}^D(\text{key}, 1^y, m_0, m_1, r, L, i)$

Inputs: keys $\text{key} = \{\text{BD.mpk}_j, \text{BD.tk}_j\}_{j \in [\lambda]}$, parameter y , messages m_0, m_1, r , revocation list L and a position-index $i \in [\lambda]$.

Output : $T \subseteq \{0, 1\}^\kappa$.

1. Define oracle $\tilde{D}[\{\text{BD.mpk}_j\}_{j \in [\lambda]}, r, L, i]$ as in Figure 3.15.
2. Output $T \leftarrow \text{BD.Trace}^{\tilde{D}}(\text{BD.tk}_i, 4y, m_0 \oplus r, m_1 \oplus r, L)$.

Figure 3.14: Algorithm: SubTrace for Unbounded EITR

Algorithm $\tilde{D}^D[\text{key}, r, L, i]$

Hardwired values: keys $\text{key} = \{\text{BD.mpk}_j\}_{j \in [\lambda]}$, message r , revocation list L and a position-index $i \in [\lambda]$.

Inputs: BD.ct .

Output : 0/1

On input BD.ct , the \tilde{D} oracle does the following:

- It first shares r in $\lambda - 1$ components as follows: it chooses $\lambda - 1$ random messages r_k for $k \in [\lambda] \setminus \{i\}$, such that $\oplus_{k \in [\lambda] \setminus \{i\}} r_k = r$.
- It computes $\text{BD.ct}_k = \text{BD.Enc}(\text{BD.mpk}_k, r_k, L)$ for $k \in [\lambda] \setminus \{i\}$.
- Sets $\text{ct} = (\text{BD.ct}_1, \dots, \text{BD.ct}_{i-1}, \text{BD.ct}, \text{BD.ct}_{i+1}, \dots, \text{BD.ct}_\lambda)$.
- Outputs $b' \leftarrow D(\text{ct})$.

Figure 3.15: Oracle \tilde{D} for Unbounded EITR

scheme then the above construction of unbounded (secret/public tracing)-EITR scheme is correct.

Proof. For all $k \in [\lambda]$, if $\text{BD.ct}_k \leftarrow \text{BDEnc}(\text{BD.mpk}_k, r_k, L)$ and $\text{BD.sk}_k \leftarrow \text{BDKeyGen}(\text{BD.msk}_k, (\text{lb}, \text{id}))$, we have $r_k \leftarrow \text{BD.Dec}(\text{BD.sk}_k, \text{BD.ct}_k, L)$, as long as $\text{id} \notin L$, from the correctness of the underlying BD-TR scheme. Hence, the decryption of $\text{ct} = (\text{BD.ct}_1, \dots, \text{BD.ct}_\lambda)$ correctly outputs m as $r_1 \oplus \dots \oplus r_\lambda$. ■

Efficiency. We can instantiate the above construction by the bounded public/secret tracing-EITR scheme in Sec. 3.10. Since the above construction is simple λ times repetition of the underlying bounded EITR scheme, the parameter size of the scheme is as follows.

Secret Tracing Setting. In the secret tracing setting, we have $|\text{mpk}|, |\text{ct}|, |\text{sk}| = \text{poly}(\lambda, |\text{id}|, |\text{lb}|)$.

Public Tracing Setting. In the public tracing setting, we have

$$|\text{mpk}|, |\text{ct}| = \text{poly}(\lambda, |\text{lb}|), |\text{sk}| = \text{poly}(\lambda, |\text{id}|, |\text{lb}|).$$

3.11.2 Security

In this section, we prove that our construction of unbounded (secret/public tracing)-EITR scheme is secure.

IND-CPA security.

Theorem 3.50. *If the underlying BD-TR scheme is (very selective/adaptive) ind-cpa secure bounded (secret/public tracing)-EITR scheme, then the above construction of unbounded (secret/public tracing)-EITR is (very selective/adaptive) ind-cpa secure.*

Proof. We show that if there exists a PPT adversary that breaks the ind-cpa security of unbounded EITR scheme, then we can use \mathcal{A} to build a PPT algorithm \mathcal{B} that breaks ind-cpa security of the underlying BD-TR scheme. ■

The reduction is defined as follows:

1. \mathcal{B} first gets $1^\kappa, L$ and key queries $\{(\text{lb}_k, \text{id}_k)\}_{k \in [Q]}$, where Q is the number of key queries issued, from the adversary \mathcal{A} . (In the public trace setting, \mathcal{A} can make adaptive key queries and output L adaptively along with the challenge messages).
2. It sets $n_{\text{bd}} = 2$ and sends $1^\kappa, 1^{n_{\text{bd}}}, L$ and key queries $\{(\text{lb}_k, \text{id}_k)\}_{k \in [Q]}$ to the BD-TR challenger. The BD-TR challenger returns BD.mpk and $\{\text{BD.sk}_k\}_{k \in [Q]}$.
3. \mathcal{B} sets $\text{BD.mpk}_1 = \text{BD.mpk}$, generates $((\text{BD.mpk}_j, \text{BD.msk}_j) \leftarrow \text{BD.Setup}(1^\lambda, 1^\kappa, 2^j))$ for $j \in \{2, \dots, \lambda\}$ and sends $\text{mpk} = \{\text{BD.mpk}_j\}_{j \in [\lambda]}$ to \mathcal{A} .
4. For each key query $(\text{lb}_k, \text{id}_k)$, \mathcal{B} sets $\text{BD.sk}_{k,1} = \text{BD.sk}_k$, computes $\text{BD.sk}_{k,j} \leftarrow \text{BD.KeyGen}(\text{BD.msk}_j, \text{lb}_k, \text{id}_k)$ for $j \in \{2, \dots, \lambda\}$ and sends $\text{sk}_{\text{lb}_k, \text{id}_k} = \{\text{BD.sk}_{k,j}\}_{j \in [\lambda]}$ to \mathcal{A} .
5. When \mathcal{A} sends the challenge query (m_0, m_1) , \mathcal{B} samples $r_j \leftarrow \mathcal{M}$ for $j \in \{2, \dots, \lambda\}$, sets $m'_b = \bigoplus_{j>1} r_j \oplus m_b$ for $b \in \{0, 1\}$ and sends (m'_0, m'_1) as challenge query to the BD challenger and sets the returned ciphertext as BD.ct_1 . \mathcal{B} then computes $\text{BD.ct}_j \leftarrow \text{BD.Enc}(\text{BD.mpk}_j, r_j, L)$ for $j \in \{2, \dots, \lambda\}$ and sends $\text{ct} = \{\text{BD.ct}_j\}_{j \in [\lambda]}$ to \mathcal{A} .

6. \mathcal{A} outputs a bit b' , \mathcal{B} sends b' to the BD-TR challenger.

We observe that \mathcal{B} issues a key query (lb, id) to the BD-TR challenger only when \mathcal{A} makes a key query on (lb, id) . So, by the admissibility of \mathcal{A} , we have $\text{lb} \in L$ and thus \mathcal{B} is admissible in the ind-cpa game of BD-TR. Also, if \mathcal{A} has an advantage ϵ in distinguishing the encryptions of m_0, m_1 , then clearly \mathcal{B} has the same advantage in distinguishing the encryptions of m'_0, m'_1 and thus breaking the ind-cpa security of BD-TR.

Secure Tracing Analysis. Now we show that our construction satisfies false trace and correct trace guarantees.

False Trace Guarantee.

Theorem 3.51. *Assume that the underlying (secret/public tracing)-BD-TR scheme satisfies (selective/adaptive) false trace guarantee, then our construction of unbounded (secret/public tracing)-EITR satisfies (selective/adaptive) false trace guarantee as defined in Def. 3.30/Def. 3.29.*

Proof. Suppose there exists a PPT adversary \mathcal{A} , polynomial $p(\lambda)$ and non-negligible functions $\epsilon(\cdot), \delta(\cdot)$ such that $\Pr[\text{Fal-Tr}]_{\mathcal{A}, \epsilon, p} \geq \delta(\lambda)$, then we can build a PPT reduction \mathcal{B} that can break the false trace guarantee of BD-TR. The reduction is as follows:

1. \mathcal{B} first gets $1^\kappa, L$ and key queries $\{(\text{lb}_k, \text{id}_k)\}_{k \in [Q]}$, where Q is the number of key queries issued, from the adversary \mathcal{A} . (In the public trace setting, \mathcal{A} can make adaptive key queries and output L adaptively along with the challenge messages).
2. It sets $i = \lceil \log p(\lambda) \rceil$, sends $1^\kappa, 1^{2^i}, L$ and key queries $\{(\text{lb}_k, \text{id}_k)\}_{k \in [Q]}$ to the BD-TR challenger. The BD-TR challenger returns BD.mpk and $\{\text{BD.sk}_k\}_{k \in [Q]}$.
3. \mathcal{B} sets $\text{BD.mpk}_i = \text{BD.mpk}$, generates $((\text{BD.mpk}_j, \text{BD.msk}_j) \leftarrow \text{BD.Setup}(1^\lambda, 1^\kappa, 2^j))$ for $j \in [\lambda] \setminus \{i\}$ and sends $\text{mpk} = \{\text{BD.mpk}_j\}_{j \in [\lambda]}$ to \mathcal{A} .
4. For each key query $(\text{lb}_k, \text{id}_k)$, \mathcal{B} sets $\text{BD.sk}_{k,i} = \text{BD.sk}_k$, computes

$\text{BD.sk}_{k,j} \leftarrow \text{BD.KeyGen}(\text{BD.msk}_j, \text{lb}_k, \text{id}_k)$ for $j \in [\lambda] \setminus \{i\}$ and sends $\text{sk}_{\text{lb}_k, \text{id}_k} = \{\text{BD.sk}_{k,j}\}_{j \in \lambda}$ to \mathcal{A} .

5. Finally, when \mathcal{A} outputs a decoder box D and messages m_0, m_1 , \mathcal{B} runs $\text{isGoodDecoder}((\text{BD.mpk}_j)_{j \in [\lambda]}, 1^{1/\epsilon}, m_0, m_1, r, L, i)$ as defined in Figure 3.13, for $\lambda \cdot y$ many choices of r , until it finds an r s.t isGoodDecoder outputs 1.
6. \mathcal{B} constructs a decoding box \tilde{D} as defined in Figure 3.14, sends $(\tilde{D}, m_0 \oplus r, m_1 \oplus r)$ to the BD-TR challenger.

We observe that \tilde{D} uses the decoder D as a subroutine and it returns the response of D as its output. So, if \mathcal{A} outputs (D, m_0, m_1, L) such that the false trace guarantee does not hold with non-negligible probability, then \mathcal{B} breaks the false trace guarantee of the underlying BD-TR scheme. ■

Correct Trace Guarantee.

Theorem 3.52. *Assume that the underlying (secret/public tracing)-BD-TR scheme satisfies (very selective/adaptive) correct trace guarantee, then our construction of unbounded (secret/public tracing)-EITR satisfies (very selective/adaptive) correct trace guarantee as defined in Def. 3.30/Def. 3.29.*

Proof. Let $i = \lceil \log p(\lambda) \rceil$ and let $S, S_{\mathcal{ID}}, T_{\text{lb}}$ be as defined in Def. 3.29. Consider the following events

Event Cor-Tr_i : is defined as the event that the **SubTrace** algorithm, when run for position i outputs a correct traitor set T , i.e. $|T| > 0, (T \subseteq S_{\mathcal{ID}}) \wedge (T_{\text{lb}} \cap L = \emptyset)$.

Event $\text{Good-}\tilde{D}_i$: is defined as the event that the flag is set to 1 in step 3 and the decoder \tilde{D} defined in Fig 3.15 is $\epsilon/4$ good decoder.

Event Found-Good-r_i : is defined as the event that the **isGoodDecoder** algorithm, when run for position i outputs 1.

With these definitions, the correctness is argued via following series of inequalities:

$$\Pr\text{-Corr-Tr}(\lambda) = \Pr[\text{Corr-Tr}_i] \quad (3.19)$$

$$\geq \Pr[\text{Corr-Tr}_i \wedge \text{Found-Good-}r_i] \quad (3.20)$$

$$\geq \Pr[\text{Good-}\tilde{D}_i \wedge \text{Found-Good-}r_i \wedge p(\lambda) \geq |S_{ID}|] - \text{negl}(\lambda) \quad (3.21)$$

$$\geq \Pr[\text{Good-}\tilde{D}_i \wedge \text{Found-Good-}r_i \wedge \text{Good-Decoder} \wedge p(\lambda) \geq |S_{ID}|] - \text{negl}(\lambda) \quad (3.22)$$

$$\geq \Pr[\text{Found-Good-}r_i \wedge \text{Good-Decoder} \wedge p(\lambda) \geq |S_{ID}|] - \text{negl}(\lambda) \quad (3.23)$$

$$\geq \Pr[\text{Good-Decoder} \wedge p(\lambda) \geq |S_{ID}|] - \text{negl}(\lambda) \quad (3.24)$$

Equation (3.19) and (3.20) follow directly from the definition of the events involved.

Equation (3.21) follows from the following claim:

Claim 3.53. If BD-TR guarantees correct tracing then

$$\begin{aligned} & \Pr [\text{Good-}\tilde{D}_i \wedge \text{Found-Good-}r_i \wedge p(\lambda) \geq |S_{ID}|] \\ & \quad - \Pr[\text{Corr-Tr}_i \wedge \text{Found-Good-}r_i] \leq \text{negl}(\lambda). \end{aligned}$$

Proof. We show that if there exists an adversary \mathcal{A} who outputs a good decoder along with messages m_0, m_1 and a revocation list L such that $\Pr[\text{Good-}\tilde{D}_i \wedge \text{Found-Good-}r_i \wedge p(\lambda) \geq |S_{ID}|] - \Pr[\text{Cor-Tr}_i]$ is non negligible then we can use \mathcal{A} to construct an adversary \mathcal{B} against correct trace guarantee of the underlying BD-TR. The reduction is defined as follows:

It sets $i = \lceil \log p(\lambda) \rceil$,

1. \mathcal{B} first gets $1^\kappa, L$ and key queries $\{(\text{lb}_k, \text{id}_k)\}_{k \in [Q]}$, where Q is the number of key queries issued, from the adversary \mathcal{A} . (In the public trace setting, \mathcal{A} can make adaptive key queries and output L adaptively along with the challenge messages).
2. It sends $1^\kappa, 1^{2^i}, L$ and key queries $\{(\text{lb}_k, \text{id}_k)\}_{k \in [Q]}$ to the BD-TR challenger. The BD-TR challenger returns BD.mpk and $\{\text{BD.sk}_k\}_{k \in [Q]}$.

3. \mathcal{B} sets $\text{BD.mpk}_i = \text{BD.mpk}$, generates $((\text{BD.mpk}_j, \text{BD.msk}_j) \leftarrow \text{BD.Setup}(1^\lambda, 1^\kappa, 2^j))$ for $j \in [\lambda] \setminus \{i\}$ and sends $\text{mpk} = \{\text{BD.mpk}_j\}_{j \in [\lambda]}$ to \mathcal{A} .
4. For each key query $(\text{lb}_k, \text{id}_k)$, \mathcal{B} sets $\text{BD.sk}_{k,i} = \text{BD.sk}_k$, computes $\text{BD.sk}_{k,j} \leftarrow \text{BD.KeyGen}(\text{BD.msk}_j, \text{lb}_k, \text{id}_k)$ for $j \in [\lambda] \setminus \{i\}$ and sends $\text{sk}_{\text{lb}_k, \text{id}_k} = \{\text{BD.sk}_{k,j}\}_{j \in [\lambda]}$ to \mathcal{A} .
5. In the end, \mathcal{A} outputs a decoder D along with messages m_0, m_1 .
6. \mathcal{B} does the following:
 - If $Q > p(\lambda)$, then \mathcal{B} outputs a random decoder and quits.
 - Else, \mathcal{B} runs $\text{isGoodDecoder}(\{\text{BD.mpk}_j\}_{j \in [\lambda]}, 1^y, m_0, m_1, r, L, i)$ for uniformly and independently sampled r until it finds a r for which isGoodDecoder algorithm outputs 1. If isGoodDecoder algorithm does not output 1 even after $\lambda \cdot y$ attempts, then \mathcal{B} outputs a random decoder and quits.
 - Else, \mathcal{B} constructs decoder \tilde{D} as defined in Figure 3.15 and sets $\tilde{m}_b = m_b \oplus r$ for $b \in \{0, 1\}$.
7. \mathcal{B} sends $(\tilde{D}, \tilde{m}_0, \tilde{m}_1)$ to the BD-TR challenger.

Now we analyze the probability that \mathcal{B} outputs a good decoder. Observe that \mathcal{B} does not abort and outputs a genuine decoder when both Found-Good- r_i and $|S_{ID}| \leq p(\lambda)$ happens. Since the decoder returned by \mathcal{B} is the same decoder as defined in the SubTrace algorithm, we get that the probability that \mathcal{B} outputs a good decoder is $\Pr[\text{Good-}\tilde{D}_i \wedge \text{Found-Good-}r_i \wedge p(\lambda) \geq |S_{ID}|]$. Furthermore, the probability that the BD.Trace algorithm outputs correct set of traitors T (using the decoder returned by \mathcal{B}) is same as $\Pr[\text{Corr-Tr}_i \wedge \text{Found-Good-}r_i]$. Hence, if $\Pr[\text{Good-}\tilde{D}_i \wedge \text{Found-Good-}r_i \wedge p(\lambda) \geq |S_{ID}|] - \Pr[\text{Corr-Tr}_i \wedge \text{Found-Good-}r_i]$ is non negligible, then it breaks the security of correct trace guarantee of the underlying BD-TR scheme. ■

Equation (3.22) again follows directly from the definition. Argument for transition from equation (3.22) to (3.23) and (3.23) to (3.24) is same as that of transition from (3.16) to (3.17) and (3.17) to (3.18), respectively. This completes the proof. ■

3.12 EXTENSION TO SUPER-POLYNOMIAL SIZE REVOCATION LIST

While our main focus in the paper is on the case where the revocation list L is of polynomial size, we can consider an extension where it is of super-polynomial size. In particular, we consider the setting where L has efficient representation by a circuit C_L of polynomial size. Namely, we have $C_L(\text{lb}) = 0$ for revoked label lb and $C_L(\text{lb}) = 1$ for non-revoked label lb . We assume that the depth of C_L is bounded by some polynomial \bar{d} . Namely, $C_L \in \mathcal{C}_{|\text{lb}|, \bar{d}}$. This assumption is necessary because we will use **kpABE** (resp., **cpABE**) with the same restriction to generate a secret key (resp., a ciphertext) associated with C_L .

EITR scheme for this setting can be obtained both in the secret and the public tracing settings very similarly to the case where the revocation list is of polynomial size. Namely, we first construct (secret key/public key) RPE with super-polynomial revocation list and then apply the chain of conversions (Sec. 3.9, 3.10, and 3.11). The conversions work almost without change with the natural adaptation of replacing L with C_L and the membership check $\text{lb} \stackrel{?}{\in} L$ with $C(\text{lb}) \stackrel{?}{=} 0$. The constructions of (public key/secret key) RPE are also almost the same as those in (Sec. 3.5/Sec. 3.6 and 3.7) with the natural adaptation of generating **kpABE** secret key for C_L ²⁷ in Sec. 3.5 and 3.6 and replacing the condition $\text{lb} \notin L$ in Eq. (3.4) defining $C_{L, \text{RMFE.ct}}$ with $C_L(\text{lb}) = 1$. The main difference is that we have to assume sub-exponential LWE assumption instead of (polynomial) LWE assumption for both secret key and public key settings here, because we need adaptive security for the underlying **kpABE**. We give further details in the following.

- In the public key setting, indistinguishability of $\text{Hyb}_{6,a}$ and Hyb_7 shown in Claim 3.17, where post-challenge key queries are dealt with, cannot be proven any more if we only assume selective security for **kpABE**. The reason why the original proof does not go through is that we have to deal with the **kpABE** queries in the order of key first and ciphertext later. With polynomial size L , this does not pose a problem because when the adversary chooses L , all the labels for which we use **kpABE** security are in L and we can perform a hybrid argument over these labels. However, this is not possible for super-polynomial size L . To deal with the queries

²⁷Originally, we start from the revocation list L and then construct the circuit C_L that hardwires L into it. Here, we directly use the circuit C_L that efficiently represents the super-polynomial set L .

in this order, we assume adaptive security (as per Definition 2.9) for kpABE. Then, the indistinguishability of the games can easily be shown by changing the post-challenge kpABE ciphertexts associated with lb with $C_L(\text{lb}) = 0$ one by one.

- In the secret key setting, we also encounter similar difficulty. In particular, indistinguishability of Hyb_1 and Hyb_2 shown in Claim 3.30 cannot be proven any more if we only assume selective security for kpABE by exactly the same reason. We can overcome the problem by assuming adaptive security for kpABE.

Finally, we briefly discuss the parameter size of the resulting EITR scheme. The parameter size of the resulting scheme is the same as the case of polynomial size revocation list except that they rely on \bar{d} . Notably, they are independent from the size of the circuits being supported inheriting the succinctness properties of the underlying kpABE and cpABE.

CHAPTER 4

ATTRIBUTE BASED ENCRYPTION FOR TURING MACHINES FROM LATTICES

4.1 INTRODUCTION

Attribute based encryption (ABE) [157, 116] is a fundamental primitive in cryptography that enables fine-grained access control on encrypted data. In an ABE scheme, the ciphertext encodes a secret message m and a public *attribute* \mathbf{x} , the secret key encodes a (public) function f , and decryption outputs m if and only if $f(\mathbf{x}) = 1$. Security posits that an adversary cannot distinguish between an encryption of (m_0, \mathbf{x}) and (m_1, \mathbf{x}) given secret keys that do not decrypt the challenge. ABE has two variants – “key-policy” where the function f (typically representing an access control policy) is embedded in the key, or “ciphertext-policy” where it is embedded in the ciphertext, as the names suggest. These are denoted by KP-ABE and CP-ABE respectively.

Traditionally, the function f is represented by a circuit and while there has been fantastic progress in building ABE supporting general circuits [116, 107, 43, 26, 169, 122], there are inherent limitations to the circuit model. Circuits force the size of the input to be fixed ahead of time and also incur *worst case* running time on every input – this is dissatisfying from both the theoretical and practical perspective. To overcome these limitations, a line of works [165, 102, 32, 15, 136, 103, 104, 17, 18, 142] has studied ABE for uniform models of computation but success has been much more limited. Without relying on heavy hammers such as indistinguishability obfuscation or compact functional encryption, the state of art ABE in this regime¹ supports non-deterministic log space Turing machines from pairings [142]. In the post-quantum regime, the situation is even less satisfactory – the state of the art construction supports non-deterministic finite state

¹In this work, we only consider ABE with unbounded collusion resistance.

automata from learning with errors (LWE) but only in the *symmetric* key setting [17]. In the public key setting, even ABE for DFA – supporting unbounded lengths and with provable security – is not known, to the best of our knowledge. Thus, an outstanding open question is:

Can we extend ABE for uniform computation beyond NL?

4.1.1 Our Results

In this work, we take a leap forward and provide the first ABE scheme for Turing machines supporting unbounded collusions from lattice assumptions. In more detail, the encryptor encodes an attribute \mathbf{x} together with a bound t on the machine running time and a message m into the ciphertext, the key generator embeds a Turing machine M into the secret key and decryption returns m if and only if $M(\mathbf{x}) = 1$. Crucially, the input \mathbf{x} and machine M can be of unbounded size, the time bound t can be chosen dynamically for each input and decryption runs in input specific time.

In more detail, our results are:

1. We construct the first ABE for NL from the LWE, evasive LWE [169, 163] and Tensor LWE [169] assumptions. This yields the first (conjectured) post-quantum ABE for NL.
2. Relying on LWE, evasive LWE and a new assumption called *circular tensor* LWE, we construct ABE for all Turing machines. At a high level, the circular tensor LWE assumption incorporates circularity into the tensor LWE (Wee, Eurocrypt 2022) assumption.

Towards our ABE for Turing machines, we obtain the first CP-ABE for circuits of unbounded depth and size from the same assumptions – this may be of independent interest.

Our Assumptions. Below, we describe the evasive and circular tensor LWE assumptions. Below, we adopt the convention by Wee [169] and let the underline denote

that noise is added to the term, whose exact value is not important.

Evasive LWE. The evasive LWE assumption [169] states that if

$$\begin{aligned} \text{if } \mathbf{B}, \mathbf{A}, \mathbf{P}, \underline{\mathbf{s}^\top \mathbf{B}}, \underline{\mathbf{s}^\top \mathbf{A}}, \underline{\mathbf{s}^\top \mathbf{P}}, \text{aux} &\approx \mathbf{B}, \mathbf{A}, \mathbf{P}, \$, \$, \$, \text{aux} \\ \text{then } \mathbf{B}, \mathbf{A}, \mathbf{P}, \underline{\mathbf{s}^\top \mathbf{B}}, \underline{\mathbf{s}^\top \mathbf{A}}, \mathbf{B}^{-1}(\mathbf{P}), \text{aux} &\approx \mathbf{B}, \mathbf{A}, \mathbf{P}, \$, \$, \mathbf{B}^{-1}(\mathbf{P}), \text{aux} \end{aligned}$$

Above $\mathbf{B}^{-1}(\mathbf{P})$ is a low norm Gaussian matrix such that $\mathbf{B} \mathbf{B}^{-1}(\mathbf{P}) = \mathbf{P}$, and \mathbf{A}, \mathbf{P} are sampled in a correlated way. In the public coin version of this assumption, the auxiliary information aux above contains random coins using during the sampling. We rely on the private coin version of this assumption, similarly to prior work [22, 164].

Tensor Circular LWE. Similar to how circularity was incorporated into the evasive LWE assumption by [122], we will need to incorporate it in the tensor LWE assumption introduced by Wee [169]. The tensor LWE assumption states that for all $\mathbf{x}_1, \dots, \mathbf{x}_Q \in \{0, 1\}^\ell$, we have

$$\textcircled{1}: \mathbf{A}, \left\{ \underline{(\mathbf{s}(\mathbf{I} \otimes \mathbf{r}_i))^\top (\mathbf{A} - \mathbf{x}_i \otimes \mathbf{G}), \mathbf{r}_i} \right\}_{i \in [Q]} \approx_c \textcircled{2}: \mathbf{A}, \{\$, \mathbf{r}_i\}_{i \in [Q]}$$

In the original formulation \mathbf{s} is uniform and \mathbf{r} is sampled from a discrete Gaussian. We will require \mathbf{s} to be small, i.e. also sampled from a discrete Gaussian.

Our *tensor circular LWE assumption* basically incorporates the circular terms into the assumption. For notational brevity, we denote $\mathbf{s}(\mathbf{I} \otimes \mathbf{r}_i)$ by $\mathbf{s}_{\mathbf{r}_i}$. In more detail, for all $\mathbf{x}_1, \dots, \mathbf{x}_Q \in \{0, 1\}^L$,

$$\mathbf{A}_{\text{circ}}, \left\{ \textcircled{1}, \mathbf{A}_{\mathbf{s}_{\mathbf{r}_i}}, \mathbf{S}_{\mathbf{s}_{\mathbf{r}_i}}, \underline{\mathbf{s}_{\mathbf{r}_i}^\top (\mathbf{A}_{\text{circ}} - \mathbf{S}_{\mathbf{s}_{\mathbf{r}_i}) \otimes \mathbf{G}}, \mathbf{r}_i} \right\}_{i \in [Q]} \approx_c \mathbf{A}_{\text{circ}}, \left\{ \textcircled{2}, \$, \$, \$, \mathbf{r}_i \right\}_{i \in [Q]} \quad (4.1)$$

Above $\mathbf{A}_{\mathbf{s}_{\mathbf{r}_i}}$ is the FHE public key corresponding to secret $\mathbf{s}_{\mathbf{r}_i}$, and $\mathbf{S}_{\mathbf{s}_{\mathbf{r}_i}} = \text{hct}_{\mathbf{s}_{\mathbf{r}_i}}(\mathbf{s}_{\mathbf{r}_i})$ is an FHE ciphertext.

The rationale of security for tensor LWE does not change with the addition of the circular terms, to the best of our knowledge. While we do not see any attack on this strengthening of tensor LWE, we note that the usage of tensor LWE in our construction is inherited from the construction by Wee [169]. Any improvements to Wee’s construction are likely to lead to improvements in our setting as well.

4.1.2 Technical Overview

Next, we outline the main technical ideas developed in our work.

Using KP-ABE and CP-ABE for circuits to build ABE for TM. The starting point of our work is the compiler by Agrawal et al. [16] that shows how to construct public key functional encryption (FE) for Turing machines using two circuit FE schemes in tandem, one ciphertext-policy and one key-policy. In more detail, their construction, building upon techniques developed in [17], relies on two restricted FE schemes: one that supports decryption in the case $|(\mathbf{x}, 1^t)| > |M|$ and one that supports the case where $|(\mathbf{x}, 1^t)| \leq |M|$. Here, \mathbf{x} is the input chosen by the encryptor and is of unbounded length, M is the Turing machine chosen by the key generator and t is an upper bound on the runtime of the Turing machine on a given input \mathbf{x} . We note that t can be chosen by the encryptor dynamically for each x and is not a-priori bounded. These restricted schemes are then run in parallel so that the first scheme can be used to decrypt the ciphertext when $|(\mathbf{x}, 1^t)| > |M|$ and the second scheme can be used otherwise.

It was shown by [16] that the first sub-scheme can be instantiated using a CP-FE that supports unbounded size and unbounded depth circuits while the second can be instantiated using a KP-FE for unbounded size but bounded depth circuits. If the first scheme only supports bounded depth (but still supports unbounded size), this still yields an FE for a smaller class of computation NL. While the techniques of [16] were developed for the setting of bounded key FE, they can be adapted to the setting of unbounded key ABE as well. We skip the details here and refer the reader to Section 4.5 for the detailed

compiler. Since KP-ABE for unbounded size but bounded depth circuits has been known since 2013 [107, 43], it remains to find a CP-ABE for unbounded size, unbounded depth circuits to instantiate the compiler.

CP-ABE for Unbounded Size Circuits. Traditionally ciphertext-policy schemes have been much harder to construct than their key-policy cousins – for instance, while KP-ABE for circuits has been known for over a decade from the standard LWE assumption, it was only very recently that the first CP-ABE for circuits was discovered [169], and that too by relying on new lattice assumptions called the *evasive* and *tensor* LWE assumptions. Moreover, of these, the evasive LWE is actually a fairly strong, non-falsifiable assumption. Yet, these assumptions have generated significant excitement in the community since they enabled progress on problems that had remained “stuck” for a while.

Wee’s construction supports circuits of unbounded size but only bounded depth. While this will turn out to be a significant barrier, this construction is all that we have to begin with. Let us recap this construction below and then try to adapt it to suit our needs.

Overview of Wee’s CP-ABE. Wee’s construction can be seen as broadly following a two step outline: (i) start with a CP-ABE which satisfies only single key security, (ii) add collusion resistance via randomization of keys. Indeed, this two-step recipe for CP-ABE has been used before [27, 64], where the first step is common to all works (barring minor syntactic differences) and can be based on LWE, while the second step has been implemented using pairings in [27] and using heuristic lattice methods in [64]. Wee’s construction achieves randomization of keys via *tensors*.

In more detail, from prior work on lattice based KP-ABE [43], it is known that given an input $\mathbf{x} \in \{0, 1\}^\ell$, and a matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times \ell m}$, one can homomorphically evaluate a circuit $f : \{0, 1\}^\ell \rightarrow \{0, 1\}$ on an “input encoding” matrix of form $\mathbf{A} - \mathbf{x} \otimes \mathbf{G}$ by multiplying

on the right by a low norm matrix $\mathbf{H}_{\mathbf{A},f,\mathbf{x}}$ to obtain the term $\mathbf{A}_f - f(x)\mathbf{G}$. Here, \mathbf{G} is a special gadget matrix defined as follows. Let $\mathbf{g} = [1, 2, 2^2, \dots, 2^{\log q}]^\top$ and $\mathbf{G} = \mathbf{I} \otimes \mathbf{g}^\top$.

Towards CP-ABE, let the public key be the matrix \mathbf{A} and the ciphertext for function f and message μ be of the form $\underline{\mathbf{s}^\top \mathbf{A}_f + \mu \lceil q/2 \rceil}$ for some randomly sampled LWE secret \mathbf{s} . To decrypt correctly, we would ideally want the secret key encoding attribute \mathbf{x} to be $\underline{\mathbf{s}^\top (\mathbf{A} - \mathbf{x} \otimes \mathbf{G})}$ – this would let us recover the mask $\mathbf{s}^\top \mathbf{A}_f$ via homomorphic computation and subtract it from the ciphertext to recover the message (after rounding out a suitably bounded noise).

Of course, the key generator cannot know the encryption randomness \mathbf{s} , and indeed this randomness will change for every ciphertext (among an unbounded number) so the above does not work ². At this stage, to enable key generation, the evasive LWE assumption described above comes to our rescue. Loosely speaking, the evasive LWE assumption can be seen as a kind of “secret sharing” mechanism to generate our desired term $\underline{\mathbf{s}^\top (\mathbf{A} - \mathbf{x} \otimes \mathbf{G})}$ when the encryptor holds randomness \mathbf{s} and the key generator holds attribute \mathbf{x} . The encryptor will now additionally provide an extra LWE sample $\underline{\mathbf{s}^\top \mathbf{B}}$ under a public matrix \mathbf{B} and the key generator will provide a low norm Gaussian matrix \mathbf{R} such that $\mathbf{B} \mathbf{R} = (\mathbf{A} - \mathbf{x} \otimes \mathbf{G})$. For readability, \mathbf{R} is referred to as $\mathbf{B}^{-1}(\mathbf{A} - \mathbf{x} \otimes \mathbf{G})$ in the literature. Now, the decryptor can compute:

$$(\underline{\mathbf{s}^\top \mathbf{B}}) \mathbf{B}^{-1}(\mathbf{A} - \mathbf{x} \otimes \mathbf{G}) = \underline{\mathbf{s}^\top (\mathbf{A} - \mathbf{x} \otimes \mathbf{G})}$$

and decryption can proceed as desired. However, while we achieve correctness, the above scheme is not secure. It is easy to see that an attacker, given keys for any \mathbf{x} and $\bar{\mathbf{x}}$ can trivially break security. To achieve collusion resistance, it is necessary to randomize keys so that mix and match attacks are thwarted.

²A simple twist to this scheme yields a construction satisfying single-key security though by [156]

Randomizing keys. Wee [169] observed that the correctness of the evaluation algorithm is preserved even under tensoring with random Gaussian vectors \mathbf{r} – in more detail, one can evaluate f even on the randomized $(\mathbf{A} - \mathbf{x} \otimes \mathbf{G}) \otimes \mathbf{r}$ to recover $(\mathbf{A}_f - f(x)\mathbf{G}) \otimes \mathbf{r}$ by simply augmenting the low norm multiplication matrix to $\mathbf{H}_{\mathbf{A},f,\mathbf{x}} \otimes \mathbf{I}$. This is just a direct application of the mixed product property on tensors.

Armed with this technique, a modification of the above construction that is not immediately broken is:

$$\begin{aligned} \text{ct}_f &= \underline{\mathbf{s}^\top(\mathbf{A}_f \otimes \mathbf{I})} + \mu \cdot \mathbf{g}, \quad \underline{\mathbf{s}^\top \mathbf{B}}, \\ \text{sk}_x &= \mathbf{B}^{-1}((\mathbf{A} - \mathbf{x} \otimes \mathbf{G}) \otimes \mathbf{r}), \quad \mathbf{r} \end{aligned}$$

To decrypt, we compute

$$\begin{aligned} (\underline{\mathbf{s}^\top(\mathbf{A}_f \otimes \mathbf{I})} + \mu \cdot \mathbf{g})(\mathbf{I} \otimes \mathbf{r}) &\approx \mathbf{s}^\top(\mathbf{A}_f \otimes \mathbf{r}) + \mu \mathbf{g} \cdot (\mathbf{I} \otimes \mathbf{r}) \\ \underline{\mathbf{s}^\top \mathbf{B}}(\mathbf{B}^{-1}((\mathbf{A} - \mathbf{x} \otimes \mathbf{G}) \otimes \mathbf{r})) \cdot (\mathbf{H}_{\mathbf{A},f,\mathbf{x}} \otimes \mathbf{I}) &\approx \mathbf{s}^\top(\mathbf{A}_f \otimes \mathbf{r}) \end{aligned}$$

and subtract the second from the first to recover μ . Now, for different keys, the random vectors \mathbf{r}_i are different and mix and match attacks no longer apply.

The above scheme needs to be refined further to admit a proof, but we skip the details here for ease of exposition. We only remark that to support arbitrary (bounded) depth, Wee needs to make an additional assumption called *tensor* LWE as discussed above.

Overcoming bounded depth. Wee’s construction only supports circuits of bounded depth, primarily because the evaluation algorithms used on the ciphertext and public key only support bounded depth. In this context, the recent exciting work by Hsieh et al. [122] provides hope since they adapted the public key and ciphertext evaluation algorithms to support circuits of unbounded depth, and provided the first KP-ABE for circuits from a generalization of the evasive LWE assumption, called the *circular small*

secret evasive LWE assumption.

One may hope that plugging in the algorithms of [122] (henceforth HLL) into the CP-ABE construction of [169] can give a CP-ABE of unbounded depth, but as we shall see this approach runs into multiple difficulties. While this approach is what we will finally make work, it is via several new ideas, new assumptions and an involved security proof. We proceed to outline these next.

Unbounded Evaluation by HLL. The work of HLL supports unbounded homomorphism via two steps: (i) noise removal and (ii) bootstrapping. Suppose we evaluate on the lattice encoding $\mathbf{s}^\top(\mathbf{A} - \mathbf{x} \otimes \mathbf{G})$ as described previously and obtain $\mathbf{s}^\top(\mathbf{A}_f - f(\mathbf{x})\mathbf{G}) + \mathbf{e}_{\text{large}}$ where $\mathbf{e}_{\text{large}}$ is large noise, as the name suggests. We would like to reduce this noise, so that we can continue homomorphic evaluation. HLL provide a noise removal procedure inspired by the modulus reduction technique developed in the context of levelled FHE [61, 55], which allows to perform public operations and recover a *noiseless* encoding of the desired function $f(\mathbf{x})$. Unfortunately, the recovered encoding does not have the right structure and disallows further homomorphic computation. In more detail, their noise removal procedure allows to obtain a term

$$\text{RndPad}_{\mathbf{A}_f}(\mathbf{s}) - f(\mathbf{x}) \mathbf{s}^\top \mathbf{G}$$

where the second summand is what we want but the first is an ill-formed mask that is not amenable to homomorphism. The exact details of what this term looks like and how one obtains it are not too important here, and we will defer these to the technical sections. However, having reached this point, the obvious question is whether and how one can proceed to obtain a well formed encoding of $f(\mathbf{x})$ which has noise within the desired bound and permits further computation.

Bootstrapping or Structure Restoration. HLL provide a clever way out of the predicament by using circular security of LWE based FHE. At a high level, circular security posits that an FHE ciphertext remains pseudorandom even if it encrypts its own secret key. HLL first observe that if there was a way to compute $\underline{s^\top \mathbf{A}'_f - \text{RndPad}_{\mathbf{A}_f}(\mathbf{s})}$ then this could be combined with the term obtained above to recover $\underline{(s^\top \mathbf{A}'_f - f(\mathbf{x}) s^\top \mathbf{G})}$ which has the right shape and allows to continue homomorphic evaluation. Thus, it suffices to compute the term $\underline{s^\top \mathbf{A}'_f - \text{RndPad}_{\mathbf{A}_f}(\mathbf{s})}$.

To do so, they leverage the algebraic structure of the GSW FHE scheme [99] together with an elegant “automatic decryption” technique by Brakerski et al. [60]. In more detail, recall that in the GSW FHE scheme, the secret key is \mathbf{s}^\top , a ciphertext for message \mathbf{y}^\top is a matrix \mathbf{C} and decryption computes $\mathbf{s}^\top \mathbf{C}$ to recover $\underline{\mathbf{y}^\top}$. The work of [60] proposes to use the same secret \mathbf{s} for encoding the attribute, and set the attribute to be an FHE encryption under \mathbf{s} . As we will see, this causes “automatic decryption” which will be useful to recover the term we are after. In more detail, HLL additionally provide the following elements in their ciphertext:

$$\mathbf{S} = \text{hct}_{\mathbf{s}}(\mathbf{s}), \quad \underline{s^\top (\mathbf{A}_{\text{circ}} - \mathbf{S} \otimes \mathbf{G})} \quad (4.2)$$

Given this, one can evaluate the circuit $\text{RndPad}_f(\cdot)$ homomorphically on the second term to obtain $\underline{s^\top \mathbf{A}'_f - s^\top \text{hct}(\text{RndPad}_{\mathbf{A}_f}(\mathbf{s}))}$ for some matrix \mathbf{A}'_f . Next, given the structure of GSW ciphertext, $s^\top \text{hct}(\text{RndPad}_{\mathbf{A}_f}(\mathbf{s}))$ decrypts *automatically* to yield $\text{RndPad}_f(\mathbf{s})$ giving the overall term $\underline{s^\top \mathbf{A}'_f - \text{RndPad}_{\mathbf{A}_f}(\mathbf{s})}$ as desired.

Randomizing Unbounded Homomorphic Evaluation. Let us now try to plug in the unbounded evaluation procedures of HLL into Wee’s construction. As described above, decryption in Wee’s CP-ABE recovers the terms

$$\underline{s^\top (\mathbf{A}_f \otimes \mathbf{r}) + \mu \mathbf{g} \cdot (\mathbf{I} \otimes \mathbf{r})}, \quad \underline{s^\top ((\mathbf{A} - \mathbf{x} \otimes \mathbf{G}) \otimes \mathbf{r})}$$

By the mixed product property of tensors, this can be written as

$$\underline{s^\top(\mathbf{I} \otimes \mathbf{r})\mathbf{A}_f + \mu \mathbf{g} \cdot (\mathbf{I} \otimes \mathbf{r})}, \quad \underline{s^\top(\mathbf{I} \otimes \mathbf{r})(\mathbf{A} - \mathbf{x} \otimes \mathbf{G})}$$

Denoting $s^\top(\mathbf{I} \otimes \mathbf{r})$ as \mathbf{s}_r^\top , we obtain well formed ABE encodings

$$\underline{\mathbf{s}_r^\top \mathbf{A}_f + \mu \mathbf{g} \cdot (\mathbf{I} \otimes \mathbf{r})}, \quad \underline{\mathbf{s}_r^\top (\mathbf{A} - \mathbf{x} \otimes \mathbf{G})}$$

In particular, we can perform the bounded depth evaluation of [43] as done by Wee [169] to obtain $\underline{\mathbf{s}_r^\top (\mathbf{A}_f - f(\mathbf{x})\mathbf{G})}$ and plug this into the noise removal procedure of HLL. Taking appropriate care, this works out to yield

$$\underline{\mathbf{s}_r^\top (\text{RndPad}_{\mathbf{A}_f}(\mathbf{s}_r) - f(\mathbf{x})\mathbf{G})}$$

which appears suitable for further processing. To continue further, HLL requires the terms

$$\mathbf{S}_r = \text{hct}_{\mathbf{s}_r}(\mathbf{s}_r), \quad \mathbf{E}_r = \underline{\mathbf{s}_r^\top (\mathbf{A}_{\text{circ}} - \mathbf{S}_r \otimes \mathbf{G})} \quad (4.3)$$

Above, the subscript in hct denotes the secret key under which the ciphertext can be decrypted. Thus, by HLL, it suffices to compute the above terms so as to bootstrap and continue.

Randomized Bootstrapping. Unfortunately, careful examination shows how demanding the above requirement is. Note that \mathbf{r} is chosen by the key generator while \mathbf{s} (and hence $\mathbf{S} = \text{hct}_{\mathbf{s}}(\mathbf{s})$) by the encryptor. To randomize \mathbf{S} to $\text{hct}_{\mathbf{s}_r}(\mathbf{s}_r)$ might seem intuitively possible at first – after all, this is an FHE ciphertext and can be evaluated upon using knowledge of \mathbf{r} to obtain $\text{hct}_{\mathbf{s}}(\mathbf{s}_r)$. However, while modifying the underlying message is indeed easy, modifying the underlying *secret key* is anything but! In its full generality, modifying the secret key of an FHE scheme should not even be possible since it violates semantic security. Yet, there are amazing ways already known to modify the underlying secret key in FHE schemes – key shrinking in single key FHE [61, 55] as well as key *expansion* in multi-key FHE [76, 149]. Sadly, neither approach (nor anything

related) works for us since they require providing some “advice” terms that cannot be computed in our setting. Note that our situation is additionally complicated by the fact that we are also required to provide a circular encoding of \mathbf{S}_r under the secret \mathbf{s}_r .

Another difficulty is that if we wish to randomize the circular encoding $\mathbf{s}^\top (\mathbf{A}_{\text{circ}} - \mathbf{S} \otimes \mathbf{G})$ using tensors as discussed earlier, our key generator would be required to give a term that looks like $\mathbf{B}^{-1}((\mathbf{A}_{\text{circ}} - \mathbf{S} \otimes \mathbf{G}) \otimes \mathbf{r})$, which it cannot because unlike before, the inner quantity is not fixed and public but depends on \mathbf{S} which is only known to the encryptor.

Our Approach. We overcome both hurdles together by relying on the following obvious fact, which previously seemed like a barrier – FHE ciphertexts are good for computing on the underlying messages, not on the underlying keys. We observe that nothing forces us to provide an FHE ciphertext of \mathbf{s} under the secret key \mathbf{s} itself! We can have our encryptor sample a new FHE scheme with unrelated secret, \mathbf{t} (say) and provide an FHE encryption of \mathbf{s} under the public key corresponding to \mathbf{t} . Now, given knowledge of \mathbf{r} , an evaluator can easily compute any complicated circuit, including those necessary to compute \mathbf{S}_r and \mathbf{E}_r described in Eq. 1. We can also provide an ABE encoding using the same trick of reusing \mathbf{t} as the randomness to cause automatic decryption as was done in HLL.

In more detail, the encryptor can provide

$$\mathbf{T} = \text{hct}_{\mathbf{t}}(\mathbf{s}, \text{sd}), \quad \mathbf{D} = \underline{\mathbf{t}^\top (\mathbf{A}_1 - (1, \text{bits}(\mathbf{T})) \otimes \mathbf{G})}$$

where sd is a PRF seed, and \mathbf{A}_1 is a public matrix of appropriate dimensions. Now, one can homomorphically evaluate on the encoding \mathbf{D} in *bounded* depth, using knowledge of \mathbf{T} , to obtain

$$\underline{\mathbf{t}^\top \mathbf{A}'_r} + \underline{\mathbf{t}^\top \text{hct}_{\mathbf{t}}(\mathbf{S}_r, \mathbf{E}_r)} = \underline{\mathbf{t}^\top \mathbf{A}'_r} + (\mathbf{S}_r, \mathbf{E}_r)$$

where \mathbf{A}'_r is some \mathbf{r} dependent matrix and the equality follows by automatic decryption.

Recall that:

$$\mathbf{S}_r = \text{hct}(\mathbf{A}_{s_r}, \mathbf{s}_r), \quad \mathbf{E}_r = \underline{\mathbf{s}_r^\top (\mathbf{A}_{\text{circ}} - (1, \text{bits}(\mathbf{S}_r)) \otimes \mathbf{G})}$$

which we require to plug into the HLL bootstrapping procedure.

Now, our desired output is masked by $\mathbf{t}^\top \mathbf{A}'_r$ so it is unclear we have achieved anything. However, note that we now have the happy situation that \mathbf{t} is known to the encryptor and \mathbf{A}'_r can be computed by the key generator! So we can plug back evasive LWE so that the encryptor gives $\underline{\mathbf{t}^\top \mathbf{C}}$ for some fixed matrix \mathbf{C} and the key generator gives $\mathbf{C}^{-1}(\mathbf{A}'_r)$ so that we can recover the term $\underline{\mathbf{t}^\top \mathbf{A}'_r}$ and cancel it out. Please see Section 4.3 for the detailed description.

Onward to Unbounded CP-ABE. The core idea for randomizing the bootstrapping discussed above applied carefully to Wee's construction allows us to obtain a correct scheme supporting unbounded depth circuits. However, proving security must still contend with several hurdles. While the high level structure of our proof resembles the proof of Wee's CP-ABE, our distributions are significantly more complex and need more work to analyze. To argue indistinguishability of hybrids, we must rely on the new assumptions described previously. We refer the reader to Section 4.4 for a detailed description of the assumptions, the scheme and the proof.

4.2 PRELIMINARIES

4.2.1 Garbled Circuits

Our definition of garbled circuits is based upon the one considered in [110]. For $\lambda \in \mathbb{N}$, let C_{inp} denote a family of circuits with inp bit inputs and $C = \{C_{\text{inp}(\lambda)}\}_{\lambda \in \mathbb{N}}$. A garbling scheme for C consists of two algorithms $\text{GC} = (\text{Garble}, \text{Eval})$ with the following syntax.

$\text{Garble}(1^\lambda, C) \rightarrow \{\text{lab}_{i,b}\}_{i \in [\text{inp}], b \in \{0,1\}}$. The garbling algorithm takes as input the security parameter λ and a circuit $C \in C_{\text{inp}(\lambda)}$. It outputs a set of labels

$$\{\text{lab}_{i,b}\}_{i \in [\text{inp}], b \in \{0,1\}}.$$

$\text{Eval}(\{\text{lab}_i\}_{i \in [\text{inp}]}) \rightarrow y$. The evaluation algorithm takes as input an inp labels $\{\text{lab}_i\}_{i \in [\text{inp}]}$ and outputs y .

Definition 4.1 (Correctness). A garbling scheme GC for circuit family $\{C_{\text{inp}(\lambda)}\}_{\lambda \in \mathbb{N}}$ is correct if for all $C \in C_{\text{inp}(\lambda)}$ and all $x \in \{0, 1\}^{\text{inp}(\lambda)}$, we have

$$\Pr [\text{Eval}(\{\text{lab}_{i,x_i}\}_{i \in [\text{inp}]}) \neq C(x) : (\{\text{lab}_{i,b}\}_{i \in [\text{inp}], b \in \{0,1\}}) \leftarrow \text{Garble}(1^\lambda, C)] = \text{negl}(\lambda),$$

where the probability is taken over the random coins of Garble .

Definition 4.2 (Security). A garbling scheme $\text{GC} = (\text{Garble}, \text{Eval})$ for $C = \{C_{\text{inp}(\lambda)}\}_{\lambda \in \mathbb{N}}$ is said to be a secure garbling scheme if there exists a PPT simulator Sim such that for all $\lambda \in \mathbb{N}$, inp , $C \in C_{\text{inp}(\lambda)}$ and $x \in \{0, 1\}^{\text{inp}(\lambda)}$, we have

$$\begin{aligned} & \left\{ \{\text{lab}_i\}_{i \in [\text{inp}]} : \{\text{lab}_i\}_{i \in [\text{inp}]} \leftarrow \text{Sim}(1^\lambda, 1^{\text{inp}}, 1^{|C|}, C(x)) \right\} \\ & \approx_c \left\{ \{\text{lab}_{i,x_i}\}_{i \in [\text{inp}]} : \{\text{lab}_{i,b}\}_{i \in [\text{inp}], b \in \{0,1\}} \leftarrow \text{Garble}(1^\lambda, C) \right\} \end{aligned}$$

Garbled circuits are known to exist from one-way functions [173, 36]. Though the above definition is not as general as one in [36], it suffices for our constructions in this paper.

Remark 11. Note that in the standard syntax of garbling scheme [145, 35], the output of the garbling algorithm consists of labels along with a garbled circuit. On the other hand, the output only consists of labels in our syntax. However, the former can easily be converted into the latter by including the garbled circuit into a label for example. This change in syntax is made for simplifying our constructions and notations. We also note that the same primitive is called decomposable randomized encoding in [106].

Remark 12 (Multi-instance security definition). The above definition allows to argue security of a single instance of GC . In the multi-instance variant, the adversary can adaptively make polynomial number of garbling queries to the challenger. All queries are answered by honestly generated labels in the real world whereas they are all simulated in the ideal world. If both worlds are computationally indistinguishable, we say that GC

satisfies multi-instance security. Note that such *multi-instance* security follows from the standard single instance security (from Definition 4.2 above) by a set of simple hybrid arguments.

4.2.2 Identity-Based Encryption

We define identity-based encryption (IBE) [137, 42, 77] in this section. An IBE scheme IBE with identity space $\mathcal{I} = \{\mathcal{I}_\lambda\}_{\lambda \in \mathbb{N}}$ consists of the following algorithms.

$\text{Setup}(1^\lambda) \rightarrow (\text{mpk}, \text{msk})$. The setup algorithm takes as input the unary representation of security parameter and outputs the master public and secret keys mpk and msk .

$\text{KeyGen}(\text{msk}, \text{id}) \rightarrow \text{sk}_{\text{id}}$. The key generation algorithm takes as input the master secret key msk and an identity $\text{id} \in \mathcal{I}_\lambda$. It outputs a corresponding secret key sk_{id} .

$\text{Enc}(\text{mpk}, \text{id}, m) \rightarrow \text{ct}$. The encryption algorithm takes as input the master public parameter mpk , an identity $\text{id} \in \mathcal{I}_\lambda$ and a message $m \in \{0, 1\}^*$. It outputs a ciphertext ct .

$\text{Dec}(\text{sk}_{\text{id}}, \text{ct}) \rightarrow m' / \perp$. The decryption algorithm takes as input the secret key sk_{id} and a ciphertext ct , and outputs either a message $m' \in \{0, 1\}^*$ or \perp .

Definition 4.3 (Correctness). For correctness, we require that for all $\lambda \in \mathbb{N}, m \in \{0, 1\}^*, \text{id} \in \mathcal{I}_\lambda$,

$$\Pr \left[\begin{array}{l} \text{Dec}(\text{sk}_{\text{id}}, \text{ct}_{\text{id}}) = m \mid (\text{mpk}, \text{msk}) \leftarrow \text{Setup}(1^\lambda, 1^L), \\ \text{sk}_{\text{id}} \leftarrow \text{KeyGen}(\text{msk}, \text{id}), \text{ct}_{\text{id}} \leftarrow \text{Enc}(\text{mpk}, \text{id}, m) \end{array} \right] = 1,$$

where the probability is taken over the random coins of Setup , KeyGen and Enc .

Remark 13 (About the identity space.). In this paper, we set the identity space \mathcal{I} to be $\{0, 1\}^*$ by default. This is without loss of generality assuming collision resistant hash function and an IBE scheme with sufficiently large identity space (e.g., $\mathcal{I}_\lambda = \{0, 1\}^{2\lambda}$),

since we can hash any string into a string of fixed length using collision resistant hash function.

Remark 14 (About the message length.). In the above definition, we assume that the message space is $\{0, 1\}^*$. This is without loss of generality if we consider IND-CPA security for IBE, since it is straightforward to extend the message space by encrypting each bit (or chunk) of the message in parallel.

Definition 4.4 (IND-CPA Security). An IBE scheme IBE for an identity space \mathcal{I}_λ and message space $\{0, 1\}^*$ is said to satisfy indistinguishability based security if for any stateful PPT adversary \mathcal{A} , there exists a negligible function $\text{negl}(\cdot)$ such that

$$\text{Adv}_{\text{IBE}, \mathcal{A}}(1^\lambda) = \left| \Pr[\text{Exp}_{\text{IBE}, \mathcal{A}}^{(0)}(1^\lambda) = 1] - \Pr[\text{Exp}_{\text{IBE}, \mathcal{A}}^{(1)}(1^\lambda) = 1] \right| \leq \text{negl}(\lambda),$$

where for each $b \in \{0, 1\}$ and $\lambda \in \mathbb{N}$, the experiment $\text{Exp}_{\text{IBE}, \mathcal{A}}^{(b)}$, modelled as a game between adversary and challenger, is defined as follows:

1. **Setup Phase:** On input 1^λ from the adversary \mathcal{A} , the challenger samples $(\text{mpk}, \text{msk}) \leftarrow \text{Setup}(1^\lambda)$ and replies to \mathcal{A} with mpk .
2. **Query phase:** During the game, \mathcal{A} adaptively makes the following queries in any arbitrary order and unbounded many times.
 - a) **Key Queries:** \mathcal{A} chooses an identity $\text{id} \in \mathcal{I}_\lambda$ and sends it to the challenger. For each such query, the challenger replies with $\text{sk}_{\text{id}} \leftarrow \text{KeyGen}(\text{msk}, \text{id})$.
 - b) **Encryption Queries:** \mathcal{A} submits to the challenger, an identity $\text{id}^* \in \mathcal{I}_\lambda$ and a pair of equal length messages (m_0, m_1) . The challenger replies with $\text{ct}_{\text{id}^*} \leftarrow \text{Enc}(\text{mpk}, \text{id}^*, m_b)$.
3. **Output phase:** \mathcal{A} outputs its guess b' as output of the experiment.

We say an adversary \mathcal{A} is legitimate if during the challenge phase, it is restricted to query for ciphertexts corresponding to identities $\text{id}^* \in \mathcal{I}_\lambda$ whose secret keys are not queried during any key query.

Remark 15. Since we are in the public key setting, we can simplify the above experiment so that the encryption query is allowed for only once. This simplified definition is shown to equivalent to the definition above by a simple hybrid argument. We adopt the above

multi-challenge security definition since it is convenient for our purpose in this paper.

4.2.3 Turing Machines

Here, we recall the definition of a Turing machine (TM) following [142]. The definition considers Turing machines with two tapes, namely, input tape and working tape.

Definition 4.5 (Turing Machine). A (deterministic) TM M is represented by the tuple $M = (Q, \delta, F)$ where Q is the number of states (we use $[Q]$ as the set of states and 1 as the initial state), $F \subset [Q]$ is the set of accepting state and

$$\begin{aligned} \delta : [Q] \times \{0, 1\} \times \{0, 1\} &\rightarrow [Q] \times \{0, 1\} \times \{0, \pm 1\} \times \{0, \pm 1\} \\ (q, b_1, b_2) &\mapsto (q', b'_2, \Delta i, \Delta j) \end{aligned}$$

is the state transition function, which, given the current state q , the symbol b_1 on the input tape under scan, and the symbol b_2 on the work tape under scan, specifies the new state q' , the symbol b'_2 , overwriting b_2 , the direction Δi to which the input tape pointers moves, and the direction Δj to which the work tape pointer moves. The machine is required to hang (instead of halting) once it reaches an accepting state, i.e., for all $q \in [Q]$ such that $q \in F$ and $b_1, b_2 \in \{0, 1\}$, it holds that $\delta(q, b_1, b_2) = (q, b_2, 0, 0)$.

For input length $n \geq 1$ and space complexity bound $s \geq 1$, the set of internal configurations of M is

$$\mathcal{Q}_{M,n,s} = [n] \times [s] \times \{0, 1\}^s \times [Q]$$

where $(i, j, W, q) \in \mathcal{Q}_{M,n,s}$ specifies the input tape pointer $i \in [n]$, the work tape pointer $j \in [s]$, the content of the work tape $W \in \{0, 1\}^s$ and the machine state $q \in [Q]$.

For any bit-string $x \in \{0, 1\}^n$ for $n \geq 1$ and time/space complexity bounds $t, s \geq 1$, the machine M accepts x within time t and space s if there exists a sequence of internal configurations (computation path of t steps) $c_0, \dots, c_t \in \mathcal{Q}_{M,n,s}$ with $c_k = (i_k, j_k, W_k, q_k)$

such that $(i_0, j_0, W_0, q_0) = (1, 1, 0^s, 1)$ (initial configuration),

$$\text{for all } 0 \leq k < t: \begin{cases} \delta(q_k, x[i_k], W_k[j_k]) = (q_{k+1}, W_{k+1}[j_k], i_{k+1} - i_k, j_{k+1} - j_k) \\ W_{k+1}[j] = W_k[j] \quad \text{for all } j \neq j_k \end{cases} \quad (\text{valid transitions});$$

where $x[i]$ is the i the bit of the string x and $W_k[j]$ is the j the bit of the string W_k , and $q_t \in F$ (accepting). We also say M accepts x within time t (without the space bound) if M accepts x within time t and space $s = t$.

Next, we define time/space bounded computation with *non-deterministic* Turing machines.

The definition is the same as Definition 4.5, except with the following changes:

- The transition criterion δ can be any relation between (i.e., any subset of the Cartesian product of) $[Q] \times \{0, 1\}^2$ and $[Q] \times \{0, 1\} \times \{0, \pm q\}^2$, where $((q, b_1, b_2), (b'_2, b'_2, \Delta i, \Delta j)) \in \delta$ means that if the current state is q , the input tape symbol under scan is b_1 and the work tape symbol under scan is b_2 , then it is valid to transit into state q' , overwrite b_2 with b'_2 , and move the input and work tape pointer by offsets Δi and Δj respectively.
- The definition of hanging in accepting states is that for all $q \in [Q]$ such that $q \in F$ and all $b_1, b_2 \in \{0, 1\}$,

$$\delta \cap \left(\{(q, b_1, b_2)\} \times ([Q] \times \{0, 1\} \times \{0, \pm 1\}^2) \right) = \{(q, b_1, b_2), (q, b_2, 0, 0)\}.$$

- In the definition of acceptance

$$\delta(q_k, x[i_k], W_k[j_k]) = (q_{k+1}, W_{k+1}[j_k], i_{k+1} - i_k, j_{k+1} - j_k)$$

is changed to $((q_k, x[i_k], W_k[j_k]), (q_{k+1}, W_{k+1}[j_k], i_{k+1} - i_k, j_{k+1} - j_k)) \in \delta$.

The following lemma can be obtained by a simple argument on emulating a Turing machine on Boolean circuits. While we have many other clever methods of simulating Turing machines on circuits (e.g., [152]), we use the following simple one because it evaluates the depth of the circuit as a function on the size of Turing machine, which is usually regarded as a constant and ignored.

Lemma 4.1 (Emulating a Turing Machine on Circuit). *Consider a circuit that takes as input a description of a (deterministic) Turing machine $M = (Q, \delta, F)$, input x to M , a configuration $(i, j, W, q) \in Q_{M, |x|, |W|}$ and outputs the next configuration (i', j', W', q') .*

We can implement such a circuit with depth $\text{poly}(\log |x|, \log |W|, \log |M|)$ and size $\text{poly}(|x|, |W|, |M|)$.

Proof. The circuit is implemented as follows. We focus on the depth of the circuits, since the bound on the size will be clear from the description. Given the input, it first retrieves the i -th bit $x[i] \in \{0, 1\}$ of the input tape. This can be done by a circuit with depth $O(\log |x|)$, which checks whether $i = \nu$ or not for each position ν of the string x in parallel and returns $x[\nu]$ for ν such that $\nu = i$. Similarly, it can retrieve $W[j]$ with depth $O(\log |W|)$. Given $x[i]$ and $W[j]$, it then retrieves $\delta(q, x[i], W[j])$ from δ , which can be done with depth $O(\log |Q|)$ similarly to the above. Given $\delta(q, x[i], W[j]) = (q', b', \Delta i, \Delta j)$, the update of i and j can be done in depth $O(\log |x|)$ and $O(\log |W|)$, respectively. Writing back the new value b' can also be done in depth $O(\log |W|)$ by finding the right place to write in the tape and change the value there. From the above discussion, the total depth of the circuit is $\text{poly}(\log |x|, \log |W|, \log |M|)$ as desired. ■

We also need the following lemma, which can be obtained by a simple observation.

Lemma 4.2 (Checking Transition for Non-deterministic Turing Machine). *Consider a circuit that takes as input a description of a non-deterministic Turing machine $M = (Q, \delta, F)$, input x to M , two configurations $(i, j, W, q) \in Q_{M, |x|, |W|}$ and $(i', j', W', q') \in Q_{M, |x|, |W|}$ and outputs whether $((i, j, W, q), (i', j', W', q')) \in \delta$ or not. We can implement such a circuit with depth $\text{poly}(\log |x|, \log |W|, \log |M|)$ and size $\text{poly}(|x|, |W|, |M|)$.*

Proof. The circuit is implemented as follows. We focus on the depth of the circuits, since the bound on the size will be clear from the description. Given the input, it first checks whether $i' - i \in \{0, \pm 1\}$, $j' - j \in \{0, \pm 1\}$. Clearly, this can be done in depth $\text{poly}(\log |x|, \log |W|)$. It then checks whether $W'[k] = W[k]$ for all $k \in [|W|] \setminus \{k\}$, which can be done in depth $\text{poly}(\log |W|)$. It then checks whether $((q, x[i], W[j]), (q', W'[j], i' - i, j' - j)) \in \delta$ or not. This can be checked in depth

$\text{poly}(\log |Q|)$. Therefore, the total depth of the circuit is $\text{poly}(\log |x|, \log |W|, \log |M|)$ as desired. \blacksquare

4.2.4 Attribute Based Encryption

In the following, we recall definitions of various ABEs by specifying the relation. We refer to Section 2.2 for the definition of an ABE scheme.

Key-policy Attribute Based encryption (kpABE) for circuits. To define kpABE for circuits, we set $\mathcal{X} = \{0, 1\}^*$ and \mathcal{Y} as the set of all circuits and define $R(\mathbf{x}, C) = C(\mathbf{x})$. In this paper, we consider circuit class $C_{\text{inp}, \text{dep}}$ that consists of circuits with input length $\text{inp} := \text{inp}(\lambda)$ and depth $\text{dep} := \text{dep}(\lambda)$. To do so, we set $\text{prm} = (1^{\text{inp}}, 1^{\text{dep}})$, $\mathcal{X}_{\text{prm}} = \{0, 1\}^{\text{inp}}$, and $\mathcal{Y}_{\text{prm}} = C_{\text{inp}, \text{dep}}$.

Ciphertext-policy Attribute Based encryption (cpABE) for circuits. To define cpABE for circuits, we set $\mathcal{Y} = \{0, 1\}^*$ and \mathcal{X} as the set of all circuits and define $R(C, \mathbf{x}) = C(\mathbf{x})$. In this paper, we consider circuit class $C_{\text{inp}, \text{dep}}$ that consists of circuits with input length $\text{inp} := \text{inp}(\lambda)$ and depth $\text{dep} := \text{dep}(\lambda)$. To do so, we set $\text{prm} = (1^{\text{inp}}, 1^{\text{dep}})$, $\mathcal{X}_{\text{prm}} = C_{\text{inp}, \text{dep}}$, and $\mathcal{Y}_{\text{prm}} = \{0, 1\}^{\text{inp}}$.

ABE for Turing Machines. To define ABE for Turing machines, we set $\mathcal{X} = \{0, 1\}^*$, \mathcal{Y} to be set of all Turing machine, and define $R : \mathcal{X} \times \mathcal{Y} \rightarrow \{0, 1\} \cup \{\perp\}$ as

$$R((\mathbf{x}, 1^t), M) = \begin{cases} 0 & \text{if } M \text{ accepts } \mathbf{x} \text{ in } t \text{ steps} \\ 1 & \text{otherwise.} \end{cases}$$

ABE for NL. To define ABE for NL, we set $\mathcal{X} = \{0, 1\}^*$, \mathcal{Y} to be set of all non-deterministic Turing machines with two tapes, one of which encodes the input and can only be read, whereas the other tape can be read as well as written. When we measure

the space complexity of the computation, we consider the space being used for the latter tape. We define $R : \mathcal{X} \times \mathcal{Y} \rightarrow \{0, 1\} \cup \{\perp\}$ as

$$R((\mathbf{x}, 1^t, 1^{2^s}), M) = \begin{cases} 0 & \text{if } M \text{ accepts } \mathbf{x} \text{ within } t \text{ steps and space } s \\ 1 & \text{otherwise.} \end{cases}$$

Note that here, s is in the exponent to reflect the idea that the space for the computation is logarithmically bounded.

We will use the **kpABE** scheme given by [43] and the **cpABE** scheme by [169] for our constructions in this chapter. We refer to Theorem 3.3 and Theorem 3.4 for the properties of [43] and [169] schemes, respectively.

4.2.5 Tensors

In this work, similarly to [169], we use the tensor product techniques. Let $\mathbf{A} = (a_{i,j}) \in \mathbb{Z}_q^{m \times n}$ and $\mathbf{B} \in \mathbb{Z}_q^{s \times t}$. The tensor product is defined as:

$$\mathbf{A} \otimes \mathbf{B} \stackrel{\text{def}}{=} \begin{pmatrix} a_{1,1}\mathbf{B} & \cdots & a_{1,n}\mathbf{B} \\ \vdots & & \vdots \\ a_{m,1}\mathbf{B} & \cdots & a_{m,n}\mathbf{B} \end{pmatrix} \in \mathbb{Z}_q^{ms \times nt}.$$

Throughout the paper, we will heavily use the mixed-product equality, stated as follows.

Let $\mathbf{A} \in \mathbb{Z}_q^{m \times n}$, $\mathbf{B} \in \mathbb{Z}_q^{s \times t}$, $\mathbf{C} \in \mathbb{Z}_q^{n \times u}$ and $\mathbf{D} \in \mathbb{Z}_q^{t \times v}$,

$$(\mathbf{A} \otimes \mathbf{B}) \cdot (\mathbf{C} \otimes \mathbf{D}) = (\mathbf{AC}) \otimes (\mathbf{BD}) \in \mathbb{Z}_q^{ms \times uv}.$$

The mixed-product can be naturally generalized following

$$(\mathbf{A}^1 \otimes \cdots \otimes \mathbf{A}^k) \cdot (\mathbf{B}^1 \otimes \cdots \otimes \mathbf{B}^k) = (\mathbf{A}^1 \mathbf{B}^1) \otimes \cdots \otimes (\mathbf{A}^k \mathbf{B}^k).$$

Note that we adopt the same convention as in [169] where matrix multiplication takes precedence over tensor products, i.e. $\mathbf{A} \otimes \mathbf{BC} = \mathbf{A} \otimes (\mathbf{BC})$.

4.2.6 Hardness Assumptions

Assumption 4.3 (Circular Small Secret LWE). [122] Let n, m, m', q, χ, χ' be functions of λ and

$$\begin{aligned} \bar{\mathbf{A}}_{\text{fhe}} &\leftarrow \mathbb{Z}_q^{n \times m}, \bar{\mathbf{A}}' \leftarrow \mathbb{Z}_q^{n \times m'}, \mathbf{r} \leftarrow \mathcal{D}_{\mathbb{Z}, \chi}^n, \mathbf{s} \leftarrow (\mathbf{r}^\top, -1)^\top, \mathbf{e}_{\text{fhe}} \leftarrow \mathcal{D}_{\mathbb{Z}, \chi}^m, \mathbf{e}' \leftarrow \mathcal{D}_{\mathbb{Z}, \chi'}^{m'}, \\ \mathbf{R} &\leftarrow \{0, 1\}^{m \times (n+1) \lceil \log_2 q \rceil m}, \delta_{\text{fhe}} \leftarrow \mathbb{Z}_q^m, \delta' \leftarrow \mathbb{Z}_q^{m'}, \Delta \leftarrow \mathbb{Z}_q^{(n+1) \times (n+1) \lceil \log_2 q \rceil m} \end{aligned}$$

The circular small-secret LWE assumption $\text{csLWE}_{n, m, m', q, \chi, \chi'}$ states that

$$\begin{aligned} &\left\{ \left(1^\lambda, \begin{pmatrix} \bar{\mathbf{A}}_{\text{fhe}} \\ \mathbf{r}^\top \bar{\mathbf{A}}_{\text{fhe}} + \mathbf{e}_{\text{fhe}}^\top \end{pmatrix}, \begin{pmatrix} \bar{\mathbf{A}}_{\text{fhe}} \\ \mathbf{r}^\top \bar{\mathbf{A}}_{\text{fhe}} + \mathbf{e}_{\text{fhe}}^\top \end{pmatrix} \mathbf{R} - \text{bits}(\mathbf{s}) \otimes \mathbf{G}, \bar{\mathbf{A}}', \mathbf{r}^\top \bar{\mathbf{A}}' + (\mathbf{e}')^\top \right) \right\}_{\lambda \in \mathbb{N}} \\ &\approx \left\{ \left(1^\lambda, \begin{pmatrix} \bar{\mathbf{A}}_{\text{fhe}} \\ \delta_{\text{fhe}}^\top \end{pmatrix}, \Delta, \bar{\mathbf{A}}', (\delta_i)^\top \right) \right\}_{\lambda \in \mathbb{N}} \end{aligned}$$

4.2.7 GSW Homomorphic Encryption and Evaluation

We recall the format (without the distribution) of the (leveled fully) homomorphic encryption [99] and the correctness property. We adapt the syntax from [122].

Lemma 4.4. *The leveled FHE scheme works as follows:*

- *The keys are*

$$(\text{public}) \quad \mathbf{A}_{\text{fhe}} = \begin{pmatrix} \bar{\mathbf{A}}_{\text{fhe}} \\ \bar{\mathbf{s}}^\top \bar{\mathbf{A}}_{\text{fhe}} + \mathbf{e}_{\text{fhe}}^\top \end{pmatrix} \in \mathbb{Z}_q^{(n+1) \times m}, \quad (\text{secret}) \quad \mathbf{s}^\top = (\bar{\mathbf{s}}^\top, -1),$$

where $\bar{\mathbf{s}} \in \mathbb{Z}^n, \bar{\mathbf{A}}_{\text{fhe}} \in \mathbb{Z}_q^{n \times m}$, and $\mathbf{e}_{\text{fhe}}^\top \in \mathbb{Z}^m$.

- *A ciphertext of $x \in \{0, 1\}$ is $\mathbf{X} = \mathbf{A}_{\text{fhe}} \mathbf{R} - x \mathbf{G} \in \mathbb{Z}_q^{(n+1) \times m}$, where $\mathbf{R} \in \mathbb{Z}_q^{m \times m}$ is the encryption randomness. The decryption equation is*

$$\mathbf{s}^\top \mathbf{X} = -\mathbf{e}_{\text{fhe}}^\top \mathbf{R} - x \mathbf{s}^\top \mathbf{G} \in \mathbb{Z}_q^m, \quad (4.4)$$

which can be used to extract x via multiplication by $\mathbf{G}^{-1}(\lfloor q/2 \rfloor \iota_{n+1})$.

- *There is an efficient algorithm*

$$\text{MakeHEvalCkt}(1^n, 1^m, q, C) = \text{HEval}_C$$

that takes as input n, m, q and a circuit $C : \{0, 1\}^L \rightarrow \{0, 1\}$ and outputs a circuit

$$\text{HEval}_C(\mathbf{X}_1, \dots, \mathbf{X}_L) = \mathbf{C}$$

taking L ciphertexts as input and outputting a new ciphertext \mathbf{C} .

- The depth of HEval_C is $dO(\log m \log \log q)$, where d is the depth of C .
- Suppose $\mathbf{X}_\ell = \mathbf{A}_{\text{fhe}} \mathbf{R}_\ell - \mathbf{x}[\ell] \mathbf{G}$ for $\ell \in [L]$ with $\mathbf{x} \in \{0, 1\}^L$, then

$$\mathbf{C} = \mathbf{A}_{\text{fhe}} \mathbf{R}_C - C(\mathbf{x}) \mathbf{G},$$

$$\text{where } \|\mathbf{R}_C^\top\| \leq (m+2)^d \max_{\ell \in [L]} \|\mathbf{R}_\ell^\top\|.$$

Additionally, in the circular version, ciphertexts of $\text{bits}(\mathbf{s})$ are published.

Lemma 4.5. (homomorphic evaluation for vector-valued functions [122]) There is an efficient algorithm

$$\text{MakeVEvalCkt}(1^n, 1^m, q, C) = \text{VEval}_C$$

that takes as input n, m, q and a vector-valued circuit $C : \{0, 1\}^L \rightarrow \mathbb{Z}_q^{1 \times m'}$ and outputs a circuit

$$\text{VEval}_C(\mathbf{X}_1, \dots, \mathbf{X}_L) = \mathbf{C},$$

taking L ciphertexts as input and outputting a new ciphertext \mathbf{C} of different format.

- The depth of VEval_C is $d \cdot O(\log m \log \log q) + O(\log^2 \log q)$ for C of depth d .
- Suppose $\mathbf{X}_\ell = \mathbf{A}_{\text{fhe}} \mathbf{R}_\ell - \mathbf{x}[\ell] \mathbf{G}$ for $\ell \in [L]$ with $\mathbf{x} \in \{0, 1\}^L$, then

$$\mathbf{C} = \mathbf{A}_{\text{fhe}} \mathbf{R}_C - \begin{pmatrix} \mathbf{0}_{n \times m'} \\ C(\mathbf{x}) \end{pmatrix} \in \mathbb{Z}_q^{(n+1) \times m'},$$

where $\|\mathbf{R}_C^\top\| \leq (m+2)^d \lceil \log q \rceil \max_{\ell \in [L]} \|\mathbf{R}_\ell^\top\|$. The new decryption equation is

$$\mathbf{s}^\top \mathbf{C} = -\mathbf{e}_{\text{fhe}}^\top \mathbf{R}_C + C(\mathbf{x}) \in \mathbb{Z}_q^{1 \times m'}.$$

4.2.8 BGG+ Homomorphic Evaluation Procedures

In this section we describe the properties of the attribute encoding and its homomorphic evaluation. We adapt the syntax from [122].

- For L -bit input, the public parameter is $\mathbf{A}_{\text{att}} \in \mathbb{Z}_q^{(n+1) \times (L+1)m}$.

- The encoding of $\mathbf{x} \in \{0, 1\}^L$ is

$$\mathbf{s}^\top (\mathbf{A}_{\text{att}} - (1, \mathbf{x}^\top) \otimes \mathbf{G}) + \mathbf{e}_{\text{att}}^\top,$$

where $\mathbf{s}^\top = (\bar{\mathbf{s}}^\top, -1)$ with $\bar{\mathbf{s}} \in \mathbb{Z}^n$ and $\mathbf{e}_{\text{att}}^\top \in \mathbb{Z}^{(L+1)m}$.

- There are efficient deterministic algorithms [43]

$$\text{EvalC}(\mathbf{A}_{\text{att}}, C) = \mathbf{H}_C \quad \text{and} \quad \text{EvalCX}(\mathbf{A}_{\text{att}}, C, \mathbf{x}) = \mathbf{H}_{C,\mathbf{x}}$$

that take as input \mathbf{A}_{att} , a circuit $C : \{0, 1\}^L \rightarrow \{0, 1\}$, and (for EvalCX) some $\mathbf{x} \in \{0, 1\}^L$, and output some matrix in $\mathbb{Z}^{(L+1)m \times m}$.

- Suppose C is of depth d , then $\|\mathbf{H}_C^\top\|, \|\mathbf{H}_{C,\mathbf{x}}^\top\| \leq (m+2)^d$.
- They satisfy encoding homomorphism, $(\mathbf{A}_{\text{att}} - (1, \mathbf{x}^\top) \otimes \mathbf{G})\mathbf{H}_{C,\mathbf{x}} = \mathbf{A}_{\text{att}}\mathbf{H}_C - C(\mathbf{x})\mathbf{G}$.

- There are efficient deterministic algorithms [60]

$$\text{MEvalC}(\mathbf{A}_{\text{att}}, C) = \mathbf{H}_C \quad \text{and} \quad \text{MEvalCX}(\mathbf{A}_{\text{att}}, C, \mathbf{x}) = \mathbf{H}_{C,\mathbf{x}}$$

that take as input \mathbf{A}_{att} , a matrix-valued circuit $C : \{0, 1\}^L \rightarrow \mathbb{Z}_q^{n+1 \times m'}$, and (for MEvalCX) some $\mathbf{x} \in \{0, 1\}^L$, and output some matrix in $\mathbb{Z}^{(L+1)m \times m'}$.

- Suppose C is of depth d , then $\|\mathbf{H}_C^\top\|, \|\mathbf{H}_{C,\mathbf{x}}^\top\| \leq (m+2)^d \lceil \log q \rceil$.
- The matrix encoding homomorphism is $(\mathbf{A}_{\text{att}} - (1, \mathbf{x}^\top) \otimes \mathbf{G})\mathbf{H}_{C,\mathbf{x}} = \mathbf{A}_{\text{att}}\mathbf{H}_C - C(\mathbf{x})$.

Dual-Use Technique and Extension. In [60], the attribute encoded with secret \mathbf{s}^\top is FHE ciphertexts under key \mathbf{s}^\top (the same, "dual-use") and the circuit being MEvalCX'ed is some HEval_C. This leads to automatic decryption. Let C be a circuit with Boolean output, \mathbf{x} an input, \mathbf{X} a bunch of FHE ciphertexts of $\text{bits}(\mathbf{x})$ under \mathbf{s}^\top , and $\mathbf{e}_{\text{att}}, \mathbf{e}', \mathbf{e}''$ some unspecified noises, then

$$\begin{aligned} & \mathbf{s}^\top (\mathbf{A}_{\text{att}} - (1, \text{bits}(\mathbf{X})) \otimes \mathbf{G}) + \mathbf{e}_{\text{att}}^\top \cdot \mathbf{H}_{\text{HEval}_C, \mathbf{X}} \\ &= \mathbf{s}^\top \mathbf{A}_{\text{att}} \mathbf{H}_{\text{HEval}_C} - \mathbf{s}^\top \text{HEval}_C(\mathbf{X}) + (\mathbf{e}')^\top \quad (\text{MEvalCX}) \end{aligned}$$

$$\begin{aligned}
&= \mathbf{s}^\top \mathbf{A}_{\text{att}} \mathbf{H}_{\text{HEval}_C} - \mathbf{s}^\top C(\mathbf{x}) \mathbf{G} + (\mathbf{e}'')^\top \quad (\text{HEval decryption}) \\
&= \mathbf{s}^\top (\mathbf{A}_{\text{att}} \mathbf{H}_{\text{HEval}_C} - C(\mathbf{x}) \mathbf{G}) + (\mathbf{e}'')^\top.
\end{aligned}$$

To extend the dual-use technique to vector-valued circuits, let the codomain of C be $\mathbb{Z}_q^{1 \times m'}$, then VEval_C is $\mathbb{Z}_q^{(n+1) \times m'}$ -valued and

$$\begin{aligned}
&\mathbf{s}^\top (\mathbf{A}_{\text{att}} - (1, \text{bits}(\mathbf{X})) \otimes \mathbf{G}) + \mathbf{e}_{\text{att}}^\top \cdot \mathbf{H}_{\text{VEval}_C, \mathbf{X}} \\
&= \mathbf{s}^\top \mathbf{A}_{\text{att}} \text{VEval}_C - \mathbf{s}^\top \text{VEval}_C(\mathbf{X}) + (\mathbf{e}')^\top \quad (\text{MEvalCX}) \\
&= \mathbf{s}^\top \mathbf{A}_{\text{att}} \text{VEval}_C - C(\mathbf{x}) + (\mathbf{e}'')^\top \quad (\text{VEval decryption}).
\end{aligned}$$

Extension to circular encryption means setting $\mathbf{x} = \mathbf{s}$, for which we say $\mathbf{S}, \mathbf{A}_{\text{circ}}$ in place of $\mathbf{X}, \mathbf{A}_{\text{att}}$.

4.3 BOOTSTRAPPING RANDOMIZED HOMOMORPHIC EVALUATION

In this section, we show how to achieve unbounded homomorphism for a randomized attribute encoding of the form

$$\mathbf{c}_{\text{att}}^\top = \mathbf{s}^\top (\mathbf{I}_{n+1} \otimes \mathbf{r}) (\mathbf{A}_{\text{att}} - (1, \mathbf{x}^\top) \otimes \mathbf{G}) + \mathbf{e}_{\text{att}}^\top = \mathbf{s}_r^\top (\mathbf{A}_{\text{att}} - (1, \mathbf{x}^\top) \otimes \mathbf{G}) + \mathbf{e}_{\text{att}}^\top$$

where $\mathbf{s} \in \mathbb{Z}^{m(n+1)}$, $\mathbf{r} \in \mathbb{Z}^m$, $\mathbf{A}_{\text{att}} \in \mathbb{Z}_q^{(n+1) \times m}$, $x \in \{0, 1\}^L$, $\mathbf{e}_{\text{att}}^\top \in \mathbb{Z}^m$. We let $\|\mathbf{e}_{\text{att}}\|$ and $\|(\mathbf{s}^\top (\mathbf{I}_{n+1} \otimes \mathbf{r}))^\top\|$ be bounded by B . Here B denotes the bound of removable noise. We achieve unbounded homomorphism for the above randomized attribute encoding using the following steps.

1. **Noise removal for randomized encoding.** First we transform $\mathbf{c}_{\text{att}}^\top = \mathbf{s}_r^\top (\mathbf{A}_{\text{att}} - (1, \mathbf{x}^\top) \otimes \mathbf{G}) + \mathbf{e}_{\text{att}}^\top$ into a noiseless encoding using the procedure $\text{RemoveNoise}(\cdot)$ from HLL to achieve $\text{RemoveNoise}(\mathbf{c}_{\text{att}}^\top) = \text{RndPad}_A(\mathbf{s}_r) - C(\mathbf{x}) \mathbf{s}_r^\top \mathbf{G}$ where $\text{RndPad}_A(\mathbf{s}_r) = \text{RemoveNoise}(\mathbf{s}_r^\top \mathbf{A})$.
2. **Structure restoration.** Here we transform the noiseless encoding achieved above

to an attribute encoding with a smaller noise, following the blueprint of [122].

Before describing the above steps in detail, we first state a few tools, theorems and lemmas that will be useful later.

4.3.1 Preparation

Rounding Function. Let $\text{Rnd} : \mathbb{Z}_q \rightarrow \mathbb{Z}_q$ be a rounding function defined as $\text{Rnd}(x) = (q/p) \lfloor p/q \cdot x \rfloor$, where p divides q . We make the following observations for the function Rnd .

Claim 4.6. For $x \in \mathbb{Z}_q$, $\text{Rnd}(x) = x + e_{\text{Rnd}}$ where $|e_{\text{Rnd}}| \leq \frac{q}{2p}$.

Proof. The rounding function $\text{Rnd}(x)$ first rounds $\frac{p}{q} \cdot x$ to the nearest integer, then scales it up to \mathbb{Z}_q by multiplying with $\frac{q}{p}$. Let m be the nearest integer to $\frac{p}{q} \cdot x$. Then, $\text{Rnd}(x) = m \cdot \frac{q}{p}$. The error e_{Rnd} introduced by the rounding process is defined as $e_{\text{Rnd}} = \text{Rnd}(x) - x$.

When $\frac{p}{q} \cdot x$ is rounded to m , the error introduced is at most $\frac{1}{2}$. Scaling this error up to \mathbb{Z}_q by multiplying with $\frac{q}{p}$, the maximum error becomes $\frac{q}{p} \cdot \frac{1}{2} = \frac{q}{2p}$. Therefore, we have $|e_{\text{Rnd}}| \leq \frac{q}{2p}$. ■

Claim 4.7. For random $x, x' \in \mathbb{Z}_q$ such that $|x - x'|$ is negligibly smaller than $\frac{q}{p}$, we have $\text{Rnd}(x) - \text{Rnd}(x') = 0$ with all but negligible probability.

Proof. Let ϵ represent the negligible quantity, such that $|x - x'| \leq \epsilon \cdot \frac{q}{p}$. The function Rnd operates by scaling x to the \mathbb{Z}_p space, rounding it, and then scaling back to \mathbb{Z}_q . This creates intervals in \mathbb{Z}_q of size $\frac{q}{p}$ that map to each integer in \mathbb{Z}_p . The probability that x and x' fall into different rounding intervals is related to how close x is to the boundary of a rounding interval. If x and x' are within $\epsilon \times \frac{q}{p}$ of each other, for a small ϵ , and they straddle a rounding boundary, then they round to different values.

Given a uniform distribution of x and x' over \mathbb{Z}_q , the probability of a number falling within $\epsilon \times \frac{q}{p}$ of a rounding boundary is approximately 2ϵ because there are two boundaries (upper and lower) for each interval, and each boundary has a "risk zone" of $\epsilon \times \frac{q}{p}$ around

it.

Therefore, $\Pr[\text{Rnd}(x) \neq \text{Rnd}(x') \mid |x - x'| \leq \epsilon \cdot (q/p)] \leq 2\epsilon$, which is negligible. ■

Theorem 4.8 (Noise Removal: [122] Construction 1). *There exists a deterministic procedure $\text{RemoveNoise}(\cdot)$ such that for an input $\mathbf{u}^\top = \mathbf{s}^\top (\mathbf{A} - x\mathbf{G}) + \mathbf{e} \in \mathbb{Z}_q^{1 \times m}$, where $\mathbf{A} \in \mathbb{Z}_q^{(n+1) \times m}$, $x \in \{0, 1\}$, $\|\mathbf{s}\| \leq B$, $\|\mathbf{e}\| \leq B$,*

$$\text{RemoveNoise}(\mathbf{u}^\top) = \text{RndPad}_{\mathbf{A}}(\mathbf{s}) - x\mathbf{s}^\top \mathbf{G} \in \mathbb{Z}_q^{1 \times m}$$

where $\text{RndPad}_{\mathbf{A}}(\mathbf{s}) = \text{RemoveNoise}(\mathbf{s}^\top \mathbf{A})$.

Lemma 4.9. *A canonical Boolean circuit of $\text{RndPad}_{\mathbf{A}}(\cdot)$ is of depth $O(\log n \log \log q)$ and can be efficiently generated from $\mathbf{A} \in \mathbb{Z}_q^{(n+1) \times m}$.*

Theorem 4.10 (Bootstrapping: [122] Theorem 12). *It works as follows:*

- *The secret is $\mathbf{s} = (\bar{\mathbf{s}}^\top, -1)^\top \in \mathbb{Z}^{n+1}$ with $\text{bits}(\mathbf{s}) \in \{0, 1\}^{(n+1)\lceil \log q \rceil}$. The circular ciphertext is*

$$\mathbf{S} = \left(\bar{\mathbf{s}}^\top \bar{\mathbf{A}}_{\text{fhe}} + \mathbf{e}_{\text{fhe}}^\top \right) (\mathbf{R}_1, \dots, \mathbf{R}_{(n+1)\lceil \log q \rceil}) - \text{bits}(\mathbf{s}) \otimes \mathbf{G} \in \mathbb{Z}_q^{(n+1) \times m(n+1)\lceil \log q \rceil}$$

where $\bar{\mathbf{A}}_{\text{fhe}} \in \mathbb{Z}_q^{(n+1) \times m}$, $\mathbf{e}_{\text{fhe}}^\top \in \mathbb{Z}^m$, and $\mathbf{R}_i \in \mathbb{Z}_q^{m \times m}$ for $1 \leq i \leq (n+1)\lceil \log q \rceil$. Let $L_S = m(n+1)^2 \lceil \log q \rceil^2$, so that $\text{bits}(\mathbf{S}) \in \{0, 1\}^{1 \times L_S}$.

- *The circular encoding is $\mathbf{s}^\top (\mathbf{A}_{\text{circ}} - (1, \text{bits}(\mathbf{S})) \otimes \mathbf{G}) + \mathbf{e}_{\text{circ}}^\top$ where $\mathbf{A}_{\text{circ}} \in \mathbb{Z}_q^{(n+1) \times (L_S+1)m}$ and $\mathbf{e}_{\text{circ}} \in \mathbb{Z}^{(L_S+1)m}$.*
- *There are efficient deterministic algorithms*

$$\text{EvalRndPad}(\mathbf{A}_{\text{circ}}, \mathbf{A}) = H_{\mathbf{A}}^{\text{RndPad}} \quad \text{and} \quad \text{EvalRndPadS}(\mathbf{A}_{\text{circ}}, \mathbf{A}, \mathbf{S}) = H_{\mathbf{A}, \mathbf{S}}^{\text{RndPad}}$$

such that

$$\left\| (H_{\mathbf{A}}^{\text{RndPad}})^\top \right\|, \left\| (H_{\mathbf{A}, \mathbf{S}}^{\text{RndPad}})^\top \right\| \leq 2^{O(\log^5 \lambda)}.$$

Moreover, when \mathbf{S} is indeed of the correct form for $\text{RndPad}_{\mathbf{A}}(\cdot)$ defined in Theorem 4.8,

$$\begin{aligned} & \mathbf{s}^\top (\mathbf{A}_{\text{circ}} - (1, \text{bits}(\mathbf{S})) \otimes \mathbf{G}) H_{\mathbf{A}, \mathbf{S}}^{\text{RndPad}} \\ &= \mathbf{s}^\top \mathbf{A}_{\text{circ}} H_{\mathbf{A}}^{\text{RndPad}} - \text{RndPad}_{\mathbf{A}}(\mathbf{s}) + \mathbf{e}_{\text{fhe}}^\top \mathbf{R}_{\text{RndPad}_{\mathbf{A}}} \end{aligned}$$

where $\left\| \mathbf{R}_{\text{RndPad}_{\mathbf{A}}}^\top \right\| \leq 2^{O(\log^4 \lambda)}$.

Theorem 4.11 (Unbounded Homomorphic Evaluation: [122] Construction 2). *Let*

$C : \{0, 1\}^L \rightarrow \{0, 1\}$ and $\mathbf{x} \in \{0, 1\}^L$. Suppose

$$\mathbf{s} = (\bar{\mathbf{s}}^\top, -1)^\top, \mathbf{R} \in \{0, 1\}^{m \times m(n+1) \log q}, \mathbf{S} = \begin{pmatrix} \bar{\mathbf{A}}_{\text{fhe}} \\ \bar{\mathbf{s}}^\top \bar{\mathbf{A}}_{\text{fhe}} + \mathbf{e}_{\text{fhe}}^\top \end{pmatrix} \mathbf{R} - \text{bits}(\mathbf{s}) \otimes \mathbf{G}$$

$$\mathbf{c}_{\text{att}}^\top = \mathbf{s}^\top (\mathbf{A}_{\text{att}} + (1, \mathbf{x}) \otimes \mathbf{G}) + \mathbf{e}_{\text{att}}^\top, \mathbf{c}_{\text{circ}}^\top = \mathbf{s}^\top (\mathbf{A}_{\text{circ}} - (1, \text{bits}(\mathbf{S})) \otimes \mathbf{G}) + \mathbf{e}_{\text{circ}}^\top,$$

where $\|\mathbf{s}\| \leq B, \|\mathbf{e}_{\text{att}}\|, \|\mathbf{e}_{\text{fhe}}^\top\|, \|\mathbf{e}_{\text{circ}}\| \leq 2^{-\lambda} B$.

There are two efficient algorithms

$$\text{UEvalC}(\mathbf{A}_{\text{att}}, \mathbf{A}_{\text{circ}}, C) = \mathbf{A}_C, \quad \text{UEvalCX}(\mathbf{A}_{\text{att}}, \mathbf{c}_{\text{att}}^\top, \mathbf{A}_{\text{circ}}, \mathbf{c}_{\text{circ}}^\top, C, \mathbf{x}, \mathbf{S}) = \mathbf{c}_{C, \mathbf{x}}^\top$$

satisfying, with all but negligible probability, $\mathbf{c}_{C, \mathbf{x}}^\top = \mathbf{s}^\top (\mathbf{A}_C - C(\mathbf{x})\mathbf{G}) + \mathbf{e}_{C, \mathbf{x}}^\top$ and

$$\|\mathbf{e}_{C, \mathbf{x}}\| \leq (\|\mathbf{e}_{\text{fhe}}\| + \|\mathbf{e}_{\text{circ}}\|) \cdot 2^{\text{poly}(\log q)} \leq B.$$

4.3.2 Noise Removal for Randomized Encoding

Here we use the `RemoveNoise`(\cdot) procedure from Theorem 4.8 to remove noise from a randomized attribute encoding $\mathbf{c}_{\text{att}}^\top = \mathbf{s}_r^\top (\mathbf{A} - C(\mathbf{x})\mathbf{G}) + \mathbf{e}^\top \in \mathbb{Z}_q^{1 \times m}$ with $\|\mathbf{s}_r\|, \|\mathbf{e}\| \leq B$. Here we have $\mathbf{s}_r^\top = \mathbf{s}^\top (\mathbf{I}_{n+1} \otimes \mathbf{r})$, $\mathbf{s} \in \mathbb{Z}^{m(n+1)}$, $\mathbf{r} \in \mathbb{Z}^m$, $\mathbf{A} \in \mathbb{Z}_q^{(n+1) \times m}$, and $\mathbf{e} \in \mathbb{Z}^m$.

We have the following using Theorem 4.8.

$$\text{RemoveNoise}(\mathbf{s}_r^\top (\mathbf{A} - C(\mathbf{x})\mathbf{G}) + \mathbf{e}^\top) = \text{RndPad}_A(\mathbf{s}_r) - C(\mathbf{x})\mathbf{s}_r^\top \mathbf{G}.$$

4.3.3 Randomized Bootstrapping a.k.a Structure Restoration

To support unbounded evaluation, we transform the noiseless encoding achieved in Section 4.3.2 back to an attribute encoding with smaller noise, following the blueprint of [122]. In particular, we want to compute $\mathbf{s}_r^\top \mathbf{A}'_C - \text{RndPad}_A(\mathbf{s}_r) + (\mathbf{e}')^\top$ for some publicly computable matrix \mathbf{A}'_C and a small noise \mathbf{e}' whose norm is within the bound of removable noise. Note that, as discussed in the technical overview, this cannot be directly computed during the key generation or encryption process as it requires the knowledge of both the keygen randomness \mathbf{r} and the encryption randomness \mathbf{s} . We divide the structure

restoration into two parts.

Step 1: Computing Advice for HLL. We define a new algorithm ComputeAdvice_r that enables us to compute an extra *noisy* FHE circular encryption $\mathbf{S}'_r = \text{hct}_{s_r}(\mathbf{s}_r) + \mathbf{err}_S$ using key \mathbf{s}_r and a circular encoding $\mathbf{E}'_r = \mathbf{s}_r^\top (\mathbf{A}_{\text{circ}} - (1, \text{bits}(\mathbf{S}'_r)) \otimes \mathbf{G}) + \mathbf{e}_E^\top$ where $\mathbf{s}_r^\top = \mathbf{s}^\top (\mathbf{I}_{n+1} \otimes \mathbf{r}) = (\bar{\mathbf{s}}_r^\top, -1)^\top$ ³ is the FHE secret key.

Step 2: HLL Bootstrapping. We use \mathbf{S}'_r and \mathbf{E}'_r obtained above to compute $\mathbf{s}_r^\top \mathbf{A} - \text{RndPad}_A(\mathbf{s}_r)$ similarly to HLL.

Notations. We set $m_T = m(m(n+1)\lceil \log q \rceil + \lambda)$, $L_T = m(n+1)(m(n+1)\lceil \log q \rceil + \lambda)\lceil \log q \rceil$, $m_S = m(n+1)\lceil \log q \rceil$, $L_S = m(n+1)^2\lceil \log q \rceil^2$, and $\ell = L_S(1 + m \log q)$.

Ingredients. We require the following tools.

1. A rounding function $\text{Rnd} : \mathbb{Z}_q \rightarrow \mathbb{Z}_q$ where $\text{Rnd}(x) = \frac{q}{p} \left\lfloor \frac{px}{q} \right\rfloor$, where p is an integer that divides q .
2. Three pseudorandom functions: $\text{PRF}_1 : \{0, 1\}^\lambda \times \mathbb{Z}^m \rightarrow [-\sigma', \sigma']^{1 \times m}$, $\text{PRF}_2 : \{0, 1\}^\lambda \times \mathbb{Z}^m \rightarrow \{0, 1\}^{m \times m_S}$ and $\text{PRF}_3 : \{0, 1\}^\lambda \times \mathbb{Z}^m \rightarrow [-\sigma', \sigma']^{1 \times (L_S+1)m}$.

Next, we describe our ComputeAdvice_r procedure.

Step 1: Procedure $\text{ComputeAdvice}_r(\bar{\mathbf{A}}_{\text{fhe}}, \mathbf{A}_{\text{circ}}, \mathbf{A}_{\text{path}}, \mathbf{T}, \mathbf{D}, \mathbf{c}_1^\top)$. The algorithm ComputeAdvice_r , with $\mathbf{r} \in \mathbb{Z}_q^m$ hardwired, has the following functionality.

Input: It takes as input $(\bar{\mathbf{A}}_{\text{fhe}}, \mathbf{A}_{\text{circ}}, \mathbf{A}_{\text{path}}, \mathbf{T}, \mathbf{D}, \mathbf{c}_1^\top)$, where $\bar{\mathbf{A}}_{\text{fhe}} \in \mathbb{Z}_q^{n \times m}$, $\mathbf{A}_{\text{circ}} \in \mathbb{Z}_q^{(n+1) \times (L_S+1)m}$,

$$\mathbf{A}_{\text{path}} \in \mathbb{Z}_q^{(n+1) \times (L_T+1)m}, \mathbf{T} = \begin{pmatrix} \bar{\mathbf{A}}'_{\text{fhe}} \\ \bar{\mathbf{t}}^\top \bar{\mathbf{A}}'_{\text{fhe}} + (\mathbf{e}'_{\text{fhe}})^\top \end{pmatrix} \mathbf{R} - \text{bits}(\mathbf{s}, \text{sd}) \otimes \mathbf{G}, \in \mathbb{Z}_q^{(n+1) \times m_T}$$

³We sample \mathbf{s} and \mathbf{r} in such a way that $\mathbf{s}_r^\top = (\bar{\mathbf{s}}_r^\top, -1)^\top$ holds.

$$\mathbf{D} = \mathbf{t}^\top (\mathbf{A}_{\text{path}} - (1, \text{bits}(\mathbf{T})) \otimes \mathbf{G}) + \mathbf{e}_{\mathbf{D}}^\top \in \mathbb{Z}_q^{(L_T+1)m}, \quad \mathbf{c}_1^\top = \mathbf{t}^\top (\mathbf{A}_{\mathbf{r}}) + \mathbf{e}_1^\top \in \mathbb{Z}_q^\ell,$$

for $\mathbf{t} = (\bar{\mathbf{t}}^\top, -1)^\top \in \mathbb{Z}^{n+1}$, $\bar{\mathbf{A}}'_{\text{fhe}} \in \mathbb{Z}_q^{(n+1) \times m}$, $\mathbf{e}'_{\text{fhe}} \in \mathbb{Z}^m$, $\mathbf{R} \in \{0, 1\}^{m \times m_T}$, and $\mathbf{A}_{\mathbf{r}} = \mathbf{A}_{\text{path}} \mathbf{H}_{\mathbf{F}} \in \mathbb{Z}_q^{(n+1) \times \ell}$. Here $\mathbf{H}_{\mathbf{F}}$ is a short matrix computed w.r.t. the function F defined in Figure 4.1.

Output: It outputs $(\mathbf{S}'_{\mathbf{r}} \in \mathbb{Z}_q^{(n+1) \times m_S}, \mathbf{E}'_{\mathbf{r}} \in \mathbb{Z}_q^{1 \times (L_S+1)m})$, where

$$\mathbf{S}'_{\mathbf{r}} = \begin{pmatrix} \bar{\mathbf{A}}_{\text{fhe}} \\ \bar{\mathbf{s}}_{\mathbf{r}}^\top \bar{\mathbf{A}}_{\text{fhe}} + \mathbf{e}_{\text{fhe}}^\top \end{pmatrix} \mathbf{R} - \text{bits}(\mathbf{s}_{\mathbf{r}}) \otimes \mathbf{G} + \mathbf{err}_{\mathbf{S}}, \quad \mathbf{E}'_{\mathbf{r}} = \mathbf{s}_{\mathbf{r}}^\top (\mathbf{A}_{\text{circ}} - (1, \text{bits}(\mathbf{S}'_{\mathbf{r}})) \otimes \mathbf{G}) + \mathbf{e}_{\text{circ}}^\top$$

for $\mathbf{s}_{\mathbf{r}}^\top = (\bar{\mathbf{s}}_{\mathbf{r}}^\top, -1) \in \mathbb{Z}^{n+1}$, $\mathbf{e}_{\text{fhe}} \in \mathbb{Z}^m$, $\mathbf{err}_{\mathbf{S}}^4 \in \mathbb{Z}^{(n+1) \times m_S}$, and $\mathbf{e}_{\text{circ}} \in \mathbb{Z}_q^{1 \times (L_S+1)m}$.

If the terms \mathbf{e}'_{fhe} , $\mathbf{e}_{\mathbf{D}}$ and \mathbf{e}_1 are short, then \mathbf{e}_{fhe} , $\mathbf{err}_{\mathbf{S}}$, and \mathbf{e}_{circ} are short. This will be formally proven in Theorem 4.12.

The algorithm `ComputeAdvicer` works as follows.

1. Defining function to evaluate. Define function $F = F_{\mathbf{r}}(\cdot, \cdot)$ as in Figure 4.1. Using the fact that the computation of modular inner product and the PRF is in NC_1 , we analyse the depth of F as follows.
Step 1 and 3 can be implemented by a circuit of depth $O(\log m \log \log q + \log n \log \log q)$. Step 2 can be implemented by a circuit of depth $O(\log n \log \log q) + O(\log(m \log q)) \leq O(\log m \log \log q + \log n \log \log q)$, where the first component is w.r.t. inner product and second for the addition. Similarly step 4 can be implemented by a circuit of depth $O(\log n \log \log q)$. Thus the entire computation can be implemented with depth $d_F = O(\log m \log \log q + \log n \log \log q)$.
2. Defining circuits for homomorphic evaluation. Here we define homomorphic evaluation circuits for function F .
 - Define $\mathbf{VEval}_F = \text{MakeVEvalCkt}(n, m, q, F) \in \mathbb{Z}_q^{n+1 \times \ell}$. From Theorem 4.5, the depth of \mathbf{VEval}_F is

$$(d_F O(\log m \log \log q) + O(\log^2 \log q)) = O(\log^2 m \log^2 \log q + \log m \log n \log^2 \log q).$$

⁴This is an extra noise term in the fhe ciphertext. However, we bound it such that it is still correctly decryptable by secret key $\mathbf{s}_{\mathbf{r}}$.

Function F

Hardwired constants: $\mathbf{r} \in \mathcal{D}_{\mathbb{Z}, \sigma_{\mathbf{r}}}^m$.

On input $(\mathbf{s} \in \mathbb{Z}^{m(n+1)}, \mathbf{sd} \in \{0, 1\}^\lambda)$, proceed as follows:

- a) Set $\mathbf{s}_{\mathbf{r}}^\top := \mathbf{s}^\top (\mathbf{I}_{n+1} \otimes \mathbf{r})$ and parse $\mathbf{s}_{\mathbf{r}} = (\bar{\mathbf{s}}_{\mathbf{r}}^\top, -1)^\top$.
- b) Compute $\mathbf{a}_{\mathbf{r}}^\top = \bar{\mathbf{s}}_{\mathbf{r}}^\top \bar{\mathbf{A}}_{\text{fhe}} + \text{PRF}_1(\mathbf{sd}, \mathbf{r})$.
- c) Compute $\mathbf{S}_{\mathbf{r}} = \begin{pmatrix} \bar{\mathbf{A}}_{\text{fhe}} \\ \mathbf{a}_{\mathbf{r}}^\top \end{pmatrix} \text{PRF}_2(\mathbf{sd}, \mathbf{r}) - \text{bits}(\mathbf{s}_{\mathbf{r}}) \otimes \mathbf{G} \in \mathbb{Z}_q^{(n+1) \times m(n+1) \lceil \log q \rceil}$.
Parse $\mathbf{S}_{\mathbf{r}} \in \mathbb{Z}_q^{1 \times m(n+1)^2 \lceil \log q \rceil}$ and set $\bar{\mathbf{S}}_{\mathbf{r}} = \text{Rnd}(\mathbf{S}_{\mathbf{r}})$.
- d) Compute $\mathbf{E}_{\mathbf{r}} := \mathbf{s}_{\mathbf{r}}^\top (\mathbf{A}_{\text{circ}} - (1, \text{bits}(\bar{\mathbf{S}}_{\mathbf{r}})) \otimes \mathbf{G}) + \text{PRF}_3(\mathbf{sd}, \mathbf{r}) \in \mathbb{Z}_q^{1 \times (L_S+1)m}$.
- e) Output $(\mathbf{S}_{\mathbf{r}}, \mathbf{E}_{\mathbf{r}}) \in \mathbb{Z}_q^{1 \times \ell}$.

Figure 4.1: Function F

$$\begin{aligned} - \text{ Compute } \mathbf{H}_{\mathbf{F}} &= \text{MEvalC}(\mathbf{A}_{\text{path}}, \mathbf{VEval}_{\mathbf{F}}), & \mathbf{H}_{\mathbf{F}, \mathbf{T}} &= \\ & \text{MEvalCX}(\mathbf{A}_{\text{path}}, \mathbf{VEval}_{\mathbf{F}}, \mathbf{T}) \in \mathbb{Z}^{(L_{\mathbf{T}}+1)m \times \ell}. \end{aligned}$$

Using the depth bound from Section 4.2.8, we have

$$\begin{aligned} \|(\mathbf{H}_{\mathbf{F}})^\top\|, \|(\mathbf{H}_{\mathbf{F}, \mathbf{T}})^\top\| &\leq (m+2)^{d_{\mathbf{VEval}_{\mathbf{F}}} \lceil \log q \rceil} \\ &= (m+2)^{O(\log^2 m \log^2 \log q + \log m \log n \log^2 \log q) \lceil \log q \rceil} \\ &\leq 2^{\log^5 \lambda}. \end{aligned}$$

3. Compute the circuit homomorphically on \mathbf{D} . Here we use the homomorphic evaluation circuits to compute on \mathbf{D} .

$$\begin{aligned} \mathbf{D} \cdot \mathbf{H}_{\mathbf{F}, \mathbf{T}} &= (\mathbf{t}^\top (\mathbf{A}_{\text{path}} - (1, \text{bits}(\mathbf{T})) \otimes \mathbf{G}) + \mathbf{e}_{\mathbf{D}}^\top) \mathbf{H}_{\mathbf{F}, \mathbf{T}} \\ &= \mathbf{t}^\top \mathbf{A}_{\text{path}} \mathbf{H}_{\mathbf{F}} - \mathbf{t}^\top \mathbf{VEval}_{\mathbf{F}}(\text{bits}(\mathbf{T})) + \mathbf{e}_{\mathbf{D}}^\top \mathbf{H}_{\mathbf{F}, \mathbf{T}} \\ &= \mathbf{t}^\top \mathbf{A}_{\mathbf{r}} - \mathbf{F}(\mathbf{s}, \mathbf{sd}) + (\mathbf{e}'_{\text{fhe}})^\top \mathbf{R}_{\mathbf{F}} + \mathbf{e}_{\mathbf{D}}^\top \mathbf{H}_{\mathbf{F}, \mathbf{T}} \\ &= \mathbf{t}^\top \mathbf{A}_{\mathbf{r}} - (\mathbf{S}_{\mathbf{r}}, \mathbf{E}_{\mathbf{r}}) + \mathbf{e}_{\mathbf{F}}^\top. \end{aligned}$$

where $\mathbf{e}_{\mathbf{F}}^\top = (\mathbf{e}'_{\text{fhe}})^\top \mathbf{R}_{\mathbf{F}} + \mathbf{e}_{\mathbf{D}}^\top \mathbf{H}_{\mathbf{F}, \mathbf{T}}$ and by Theorem 4.5, we have

$$\begin{aligned} \|\mathbf{R}_{\mathbf{F}}^\top\| &\leq (m+2)^{d_{\mathbf{F}} \lceil \log q \rceil} \cdot \max_{i \in [m(n+1) \lceil \log q \rceil + \lambda]} \|\mathbf{R}_i^\top\| \\ &\leq (m+2)^{d_{\mathbf{F}} \lceil \log q \rceil} \cdot m = (m+2)^{d_{\mathbf{F}} \lceil \log q \rceil} \cdot 3(n+1) \lceil \log q \rceil \\ &\leq (m+2)^{d_{\mathbf{F}}} O(\log q) \leq 2^{O(\log^4 \lambda)}. \end{aligned}$$

4. Cancel out masking term. Compute $\mathbf{c}_1^\top - \mathbf{D} \cdot \mathbf{H}_{\mathbf{F}, \mathbf{T}}$.

$$\begin{aligned} & \mathbf{c}_1^\top - \mathbf{D} \cdot \mathbf{H}_{\mathbf{F}, \mathbf{T}} \\ &= \mathbf{t}^\top (\mathbf{A}_{\mathbf{r}}) + \mathbf{e}_1^\top - (\mathbf{t}^\top \mathbf{A}_{\mathbf{r}} - (\mathbf{S}_{\mathbf{r}}, \mathbf{E}_{\mathbf{r}}) + \mathbf{e}_{\mathbf{F}}^\top) \\ &= (\mathbf{S}_{\mathbf{r}}, \mathbf{E}_{\mathbf{r}}) + \mathbf{e}^\top = (\mathbf{S}_{\mathbf{r}} + \mathbf{e}_{\mathbf{S}}^\top, \mathbf{E}_{\mathbf{r}} + \mathbf{e}_{\mathbf{E}}^\top) \end{aligned}$$

where $\mathbf{e}^\top = \mathbf{e}_1^\top - \mathbf{e}_{\mathbf{F}}^\top \in \mathbb{Z}^{1 \times \ell}$ and $\|\mathbf{e}_{\mathbf{S}}\|, \|\mathbf{e}_{\mathbf{E}}\| \leq \|\mathbf{e}\|$.

5. Output Output $\mathbf{S}'_{\mathbf{r}} = \text{Rnd}(\mathbf{S}_{\mathbf{r}} + \mathbf{e}_{\mathbf{S}}^\top)$ and $\mathbf{E}'_{\mathbf{r}} = \mathbf{E}_{\mathbf{r}} + \mathbf{e}_{\mathbf{E}}^\top$.

We encapsulate the property of our algorithm $\text{ComputeAdvice}_{\mathbf{r}}$ in the following lemma.

Lemma 4.12 ($\text{ComputeAdvice}_{\mathbf{r}}$). *If σ' , $\|\mathbf{e}'_{\text{the}}\|, \|\mathbf{e}_{\mathbf{D}}\| \leq 2^{-2\lambda} B'$, $\|\mathbf{e}_1\| \leq 2^{-\lambda/2} B'$, $q = B' p \lambda^{\omega(1)}$, $\|\mathbf{s}_{\mathbf{r}}\| \leq 2^{-\lambda} B'$, for a bound $B' = 2^{-4\lambda} B$ where $B < q/4$ is the bound on the removable error, then for procedure $\text{ComputeAdvice}_{\mathbf{r}}$ defined above, the output $(\mathbf{S}'_{\mathbf{r}}, \mathbf{E}'_{\mathbf{r}})$ is suitable for HLL Bootstrapping. Specifically, we have*

- $\bar{\mathbf{S}}_{\mathbf{r}} = \mathbf{S}'_{\mathbf{r}}$ with overwhelming probability.
- $\bar{\mathbf{S}}_{\mathbf{r}}$ and \mathbf{S}' are valid the ciphertexts encrypting $\text{bits}(\mathbf{s}_{\mathbf{r}})$ using key $\mathbf{s}_{\mathbf{r}}$ with decryption error $\leq 2^{-\lambda} B$
- The error in $\mathbf{E}'_{\mathbf{r}}$ is bounded as $\|\mathbf{e}_{\text{circ}}\| \leq 2^{-\lambda} B$.

Proof. We show that the output of $\text{ComputeAdvice}_{\mathbf{r}}$ is valid for HLL bootstrapping.

First we note that since $\|\mathbf{e}'_{\text{the}}\|, \|\mathbf{e}_{\mathbf{D}}\| \leq 2^{-2\lambda} B'$ and $\|\mathbf{R}_{\mathbf{F}}^\top\|, \|(\mathbf{H}_{\mathbf{F}, \mathbf{T}})^\top\| \leq 2^{\text{poly}(\log \lambda)}$, from Step 3 of $\text{ComputeAdvice}_{\mathbf{r}}$ we have $\|\mathbf{e}_{\mathbf{F}}\| \leq 2^{-3\lambda/2} B'$. Also, since $\|\mathbf{e}_1\| \leq 2^{-\lambda/2} B'$, from Step 4 we have $\|\mathbf{e}_{\mathbf{S}}\|, \|\mathbf{e}_{\mathbf{E}}\| \leq \|\mathbf{e}\| \leq \|\mathbf{e}_1\| + \|\mathbf{e}_{\mathbf{F}}\| \leq B'$.

- Equality of $\bar{\mathbf{S}}_{\mathbf{r}}$ and $\mathbf{S}'_{\mathbf{r}}$. We have $\bar{\mathbf{S}}_{\mathbf{r}} = \text{Rnd}(\mathbf{S}_{\mathbf{r}})$ and $\mathbf{S}'_{\mathbf{r}} = \text{Rnd}(\mathbf{S}_{\mathbf{r}} + \mathbf{e}_{\mathbf{S}}^\top)$. From Claim 4.7, we know that $\text{Rnd}(\mathbf{S}_{\mathbf{r}}) = \text{Rnd}(\mathbf{S}_{\mathbf{r}} + \mathbf{e}_{\mathbf{S}}^\top)$ with overwhelming probability if $\|(\mathbf{S}_{\mathbf{r}} - (\mathbf{S}_{\mathbf{r}} + \mathbf{e}_{\mathbf{S}}^\top))^\top\| = \|\mathbf{e}_{\mathbf{S}}\| \leq \epsilon \frac{q}{p}$ for some negligible parameter ϵ . Note that $\|\mathbf{e}_{\mathbf{S}}\| \leq B' \leq \frac{1}{\lambda^{\omega(1)}} q/p$. Thus we have $\text{Rnd}(\mathbf{S}_{\mathbf{r}}) = \text{Rnd}(\mathbf{S}_{\mathbf{r}} + \mathbf{e}_{\mathbf{S}}^\top)$ with overwhelming probability.

- Validity of $\bar{\mathbf{S}}_{\mathbf{r}}$ and $\mathbf{S}'_{\mathbf{r}}$.

First we show that $\bar{\mathbf{S}}_{\mathbf{r}} = \mathbf{S}_{\mathbf{r}} + \mathbf{err}_1 \in \mathbb{Z}_q^{(n+1) \times m(n+1) \lceil \log q \rceil}$ is a valid the encryption of

\mathbf{s}_r . Consider the decryption equation

$$\begin{aligned}\mathbf{s}_r^\top \bar{\mathbf{S}}_r &= \mathbf{s}_r^\top \mathbf{S}_r + \mathbf{s}_r^\top \mathbf{err}_1 \\ &= -\text{PRF}_1(\text{sd}, \mathbf{r})\text{PRF}_2(\text{sd}, \mathbf{r}) - \mathbf{s}_r(\text{bits}(\mathbf{s}_r) \otimes \mathbf{G}) + \mathbf{s}_r^\top \mathbf{err}_1 \\ &= -\mathbf{s}_r(\text{bits}(\mathbf{s}_r) \otimes \mathbf{G}) + \mathbf{err}_{\text{fhe}}^\top.\end{aligned}$$

which can be used to extract $\text{bits}(\mathbf{s}_r)$ via tensoring by $\mathbf{G}^{-1}(\lfloor q/2 \rfloor_{\ell_{n+1}})$ if $\|\mathbf{err}_{\text{fhe}}\| \leq 2^{-\lambda}B < q/4$, where $\mathbf{err}_{\text{fhe}} = \mathbf{s}_r^\top \mathbf{err}_1 - \text{PRF}_1(\text{sd}, \mathbf{r})\text{PRF}_2(\text{sd}, \mathbf{r})$.

Note that Claim 4.6 implies $\|\mathbf{err}_1\| \leq \frac{q}{2p} \leq B'\lambda^{\omega(1)}$. Then, we have $\|\mathbf{err}_{\text{fhe}}\| \leq B' \cdot B'\lambda^{\omega(1)} + O(m)\sigma' \leq \lambda^{\omega(1)}(B')^2 + O(m)2^{-\lambda}B' \leq 2^{-\lambda}B$.

Next, we show that $\mathbf{S}'_r = \mathbf{S}_r + \mathbf{es} + \mathbf{err}_2 \in \mathbb{Z}_q^{(n+1) \times m(n+1)\lceil \log q \rceil}$ is a valid fhe encryption of \mathbf{s}_r . Consider the decryption equation

$$\begin{aligned}\mathbf{s}_r^\top \mathbf{S}'_r &= \mathbf{s}_r^\top \mathbf{S}_r + \mathbf{s}_r^\top (\mathbf{es} + \mathbf{err}_2) \\ &= -\mathbf{s}_r^\top (\text{bits}(\mathbf{s}_r) \otimes \mathbf{G}) + (\mathbf{err}'_{\text{fhe}})^\top.\end{aligned}$$

which can be used to extract $\text{bits}(\mathbf{s}_r)$ via tensoring by $\mathbf{G}^{-1}(\lfloor q/2 \rfloor_{\ell_{n+1}})$ if $\|\mathbf{err}'_{\text{fhe}}\| \leq 2^{-\lambda}B < q/4$, where $(\mathbf{err}'_{\text{fhe}})^\top = \mathbf{s}_r^\top (\mathbf{es} + \mathbf{err}_2) - \text{PRF}_1(\text{sd}, \mathbf{r})\text{PRF}_2(\text{sd}, \mathbf{r})$ and $\|\mathbf{err}'_{\text{fhe}}\| \leq 2^{-\lambda}B$. The analysis is similar to $\|\mathbf{err}_{\text{fhe}}\|$, hence omitted.

- Error bound in \mathbf{E}'_r . We have $\mathbf{E}'_r = \mathbf{s}_r^\top (\mathbf{A}_{\text{circ}} - (1, \text{bits}(\bar{\mathbf{S}}_r)) \otimes \mathbf{G}) + \mathbf{e}_{\text{circ}}^\top$, where $\mathbf{e}_{\text{circ}}^\top = \text{PRF}_3(\text{sd}, \mathbf{r}) + \mathbf{e}_E^\top$. We have $\|\mathbf{e}_{\text{circ}}\| \leq 2^{-\lambda}B' + B' \leq 2^{-\lambda}B$.

■

Step 2: HLL Bootstrapping On input $\mathbf{A}_{\text{circ}}, \mathbf{S}'_r = \text{hct}_{\mathbf{s}_r}(\mathbf{s}_r) + \mathbf{err}_S$, and $\mathbf{E}'_r = \mathbf{s}_r^\top (\mathbf{A}_{\text{circ}} - (1, \text{bits}(\mathbf{S}'_r)) \otimes \mathbf{G}) + \mathbf{e}_{\text{circ}}^\top$ ⁵, it does the following⁶.

1. Defining circuits for homomorphic evaluation. Define the following evaluation circuits/matrices corresponding to the function $\text{RndPad}_A(\cdot)$.

– Define $\text{VEval}_{\text{RndPad}_A} = \text{MakeVEvalCkt}(\text{RndPad}_A)$.

– Compute $\mathbf{H}_A^{\text{RndPad}} \in \mathbb{Z}^{(L_S+1)m \times m}$ and $\mathbf{H}_{A, \mathbf{S}'_r}^{\text{RndPad}} \in \mathbb{Z}^{(L_S+1)m \times m}$.

$$\begin{aligned}\mathbf{H}_A^{\text{RndPad}} &= \text{MEvalC}(\mathbf{A}_{\text{circ}}, \text{VEval}_{\text{RndPad}_A}) \in \mathbb{Z}^{(L_S+1)m \times m} \\ \mathbf{H}_{A, \mathbf{S}'_r}^{\text{RndPad}} &= \text{MEvalCX}(\mathbf{A}_{\text{circ}}, \text{VEval}_{\text{RndPad}_A}, \mathbf{S}'_r) \in \mathbb{Z}^{(L_S+1)m \times m}.\end{aligned}$$

Here $\left\| (\mathbf{H}_A^{\text{RndPad}})^\top \right\|, \left\| (\mathbf{H}_{A, \mathbf{S}'_r}^{\text{RndPad}})^\top \right\| \leq 2^{O(\log^5 \lambda)}$.

⁵We replace $\bar{\mathbf{S}}_r$ with \mathbf{S}'_r in the \mathbf{E}'_r , since $\bar{\mathbf{S}}_r = \mathbf{S}'_r$ with overwhelming probability.

⁶Note that this doesn't straightforward follows from HLL. The error terms are different in the circular ciphertext \mathbf{S}'_r . However, we can still get the final error with the desired bounds as analysed below.

2. Compute the matrix homomorphically on $\mathbf{E}'_{\mathbf{r}}$. We have

$$\begin{aligned}\mathbf{E}'_{\mathbf{r}} \cdot \mathbf{H}_{\mathbf{A}, \mathbf{S}'_{\mathbf{r}}}^{\text{RndPad}} &= \left(\mathbf{s}_{\mathbf{r}}^{\top} (\mathbf{A}_{\text{circ}} - (1, \text{bits}(\mathbf{S}'_{\mathbf{r}})) \otimes \mathbf{G}) + \mathbf{e}_{\text{circ}}^{\top} \right) \mathbf{H}_{\mathbf{A}, \mathbf{S}'_{\mathbf{r}}}^{\text{RndPad}} \\ &= \mathbf{s}_{\mathbf{r}}^{\top} \mathbf{A}_{\text{circ}} \mathbf{H}_{\mathbf{A}}^{\text{RndPad}} - \text{RndPad}_{\mathbf{A}}(\mathbf{s}_{\mathbf{r}}) + \mathbf{err}^{\top}\end{aligned}$$

where, using the fact that depth of RndPad is $O(\log n \log \log q)$, we have

$$\begin{aligned}\|\mathbf{err}\| &\leq (m+2)^{O(\log n \log \log q)} \cdot m \cdot (\|\text{PRF}_1(\mathbf{sd}, \mathbf{r})^{\top}\| + \|(\mathbf{s}_{\mathbf{r}}^{\top}(\mathbf{e}_{\mathbf{S}} + \mathbf{err}_2))^{\top}\|) \\ &\quad + \|\mathbf{e}_{\text{circ}}\| \cdot 2^{O(\log^5 \lambda)}.\end{aligned}$$

Assuming $\|\text{PRF}_1(\mathbf{sd}, \mathbf{r})^{\top}\|, \|(\mathbf{s}_{\mathbf{r}}^{\top}(\mathbf{e}_{\mathbf{S}} + \mathbf{err}_2))^{\top}\|, \|\mathbf{e}_{\text{circ}}\| \leq 2^{-\lambda}B$, we get $\|\mathbf{err}\| \leq 2^{-\lambda/2}B$.

3. Output $\mathbf{s}_{\mathbf{r}}^{\top} \mathbf{A}_C - \text{RndPad}_{\mathbf{A}}(\mathbf{s}_{\mathbf{r}}) + \mathbf{err}^{\top}$, where $\mathbf{A}_C = \mathbf{A}_{\text{circ}} \mathbf{H}_{\mathbf{A}}^{\text{RndPad}}$.

Restoring the structure. We note that combining the outputs from Section 4.3.2 and 4.3.3, we get

$$\mathbf{s}_{\mathbf{r}}^{\top} \mathbf{A}_C - C(\mathbf{x})\mathbf{s}_{\mathbf{r}}^{\top} \mathbf{G} + \mathbf{err}^{\top}, \quad \text{where } \|\mathbf{err}\| \leq 2^{-\lambda/2}B.$$

4.4 CIPHERTEXT POLICY ABE FOR UNBOUNDED DEPTH CIRCUITS

In this section we present CP-ABE for unbounded depth and size circuits.

4.4.1 Construction

We construct a CPABE scheme for circuit class $\mathcal{C} = \{C : \{0, 1\}^L \rightarrow \{0, 1\}\}$.

Setup($1^\lambda, 1^L$). The setup algorithm does the following.

1. Set $L_S = m(n+1)^2 \lceil \log q \rceil^2$, $L_T = m(n+1)(m(n+1) \lceil \log q \rceil + \lambda) \lceil \log q \rceil$, and $\ell = L_S(1 + m \log q)$.
2. Sample: $(\mathbf{B}, \tau) \leftarrow \text{TrapGen}(1^{(n+1)(m+2)}, 1^{(n+1)(m+2)w}, q)$, $\bar{\mathbf{A}}_{\text{fhe}} \leftarrow \mathbb{Z}_q^{n \times m}$, $\mathbf{A}_0 \leftarrow \mathbb{Z}_q^{(n+1) \times m}$, $\mathbf{A}_{\text{path}} \leftarrow \mathbb{Z}_q^{(n+1) \times (L_T+1)m}$, $\mathbf{A}_{\text{att}} \leftarrow \mathbb{Z}_q^{(n+1) \times (L+1)m}$, $\mathbf{A}_{\text{circ}} \leftarrow \mathbb{Z}_q^{(n+1) \times (L_S+1)m}$, $\mathbf{u} \leftarrow \mathcal{D}_{\mathbb{Z}, \sigma_{\mathbf{r}}}^m$, where $w \in O(\log q)$.

3. Output $\text{mpk} := (\bar{\mathbf{A}}_{\text{fhe}}, \mathbf{A}_0, \mathbf{A}_{\text{path}}, \mathbf{A}_{\text{att}}, \mathbf{A}_{\text{circ}}, \mathbf{B}, \mathbf{u})^7$ and $\text{msk} := \tau$.

$\text{KeyGen}(\text{msk}, \mathbf{x})$. The key generation algorithm does the following.

1. For $1 \leq j \leq m - 1$, sample $\mathbf{r}_j \leftarrow \mathcal{D}_{\mathbb{Z}, \sigma_{\mathbf{r}}}$. Set $\mathbf{r} = (\mathbf{r}_1, \dots, \mathbf{r}_m)$ such that $\sum_{j \in [m]} \mathbf{r}_j = \mathbf{1}$.
2. Define function F as in Figure 4.1 using \mathbf{r} , $\text{VEval}_F = \text{MakeVEvalCkt}(n, m, q, F) \in \mathbb{Z}_q^{(n+1) \times \ell}$ and compute $\mathbf{H}_F = \text{MEvalC}(\mathbf{A}_{\text{path}}, \text{VEval}_F)$. Set $\mathbf{A}_{\mathbf{r}} = \mathbf{A}_{\text{path}} \mathbf{H}_F$.

3. Sample

$$\mathbf{K} \leftarrow \mathbf{B}_{\tau}^{-1} \begin{pmatrix} \mathbf{A}_0 \mathbf{r} & & \\ & (\mathbf{A}_{\text{att}} - (\mathbf{1}, \mathbf{x}^{\top}) \otimes \mathbf{G}) \otimes \mathbf{r} & \\ & & \mathbf{A}_{\mathbf{r}} \end{pmatrix}$$

4. Output $\text{sk} := (\mathbf{K}, \mathbf{r})$.

$\text{Enc}(\text{mpk}, C, \mu)$. The encryption algorithm does the following.

1. Sample $\mathbf{s}_0 \leftarrow \mathcal{D}_{\mathbb{Z}, \sigma_s}^{(n+1)}$, $\mathbf{s}_j \leftarrow \mathcal{D}_{\mathbb{Z}, \sigma_s}^m$ for $1 \leq j \leq n$, $\bar{\mathbf{t}} \leftarrow \mathcal{D}_{\mathbb{Z}, \sigma_s}^n$, $\mathbf{e}'_0 \leftarrow \mathcal{D}_{\mathbb{Z}, \sigma_0}^{(n+1)w}$, $\mathbf{e}'_{\text{att}} \leftarrow \mathcal{D}_{\mathbb{Z}, \sigma_{\text{att}}}^{m(n+1)w}$, $\mathbf{e}'_1 \leftarrow \mathcal{D}_{\mathbb{Z}, \sigma_1}^{(n+1)w}$, $\mathbf{e}_{\text{msg}} \leftarrow \mathcal{D}_{\mathbb{Z}, \sigma_{\text{msg}}}^m$. Set $\mathbf{s} = (\mathbf{s}_1^{\top}, \dots, \mathbf{s}_n^{\top}, -\mathbf{1}_m)^{\top}$, $\mathbf{t} = (\bar{\mathbf{t}}^{\top}, -1)^{\top}$, $\mathbf{e}_{\mathbf{B}}^{\top} = ((\mathbf{e}'_0)^{\top}, (\mathbf{e}'_{\text{att}})^{\top}, (\mathbf{e}'_1)^{\top})$.

2. Compute

$$\mathbf{c}_{\mathbf{B}}^{\top} := (\mathbf{s}_0^{\top} \mid \mathbf{s}^{\top} \mid \mathbf{t}^{\top}) \mathbf{B} + \mathbf{e}_{\mathbf{B}}^{\top}, \quad \text{ct}_{\text{msg}}^{\top} := \mathbf{s}_0^{\top} \mathbf{A}_0 + \mathbf{s}^{\top} (\mathbf{A}_C \mathbf{u} \otimes \mathbf{I}_m) + \mu \cdot \mathbf{g}^{\top} + \mathbf{e}_{\text{msg}}^{\top}$$

where $\mathbf{A}_C \leftarrow \text{UEvalC}(\mathbf{A}_{\text{circ}}, \mathbf{A}_{\text{att}}, C)$.

3. For $\text{sd} \leftarrow \{0, 1\}^{\lambda}$ and $\mathbf{R} = (\mathbf{R}_1, \dots, \mathbf{R}_{m(n+1)\lceil \log q \rceil + \lambda})$, compute the ciphertext as follows.

$$\bar{\mathbf{A}}'_{\text{fhe}} \leftarrow \mathbb{Z}_q^{n \times m}, \quad \mathbf{e}'_{\text{fhe}} \leftarrow \mathcal{D}_{\mathbb{Z}, \sigma'}^m, \quad \mathbf{A}_{\mathbf{t}} := (\bar{\mathbf{A}}'_{\text{fhe}} \bar{\mathbf{t}}^{\top} \bar{\mathbf{A}}'_{\text{fhe}} + (\mathbf{e}'_{\text{fhe}})^{\top})^{\top},$$

$$\mathbf{T} = \mathbf{A}_{\mathbf{t}} \mathbf{R} - \text{bits}(\mathbf{s}, \text{sd}) \otimes \mathbf{G} \in \mathbb{Z}_q^{(n+1) \times m(m(n+1)\lceil \log q \rceil + \lambda)}$$

where $\mathbf{R}_i \leftarrow \{0, 1\}^{m \times m}$ for all $1 \leq i \leq (m(n+1)\lceil \log q \rceil + \lambda)$.

Set $L_T = m(n+1)(m(n+1)\lceil \log q \rceil + \lambda)\lceil \log q \rceil$ so that $\text{bits}(\mathbf{T}) \in \{0, 1\}^{L_T}$.

⁷All the algorithms take mpk implicitly.

4. Computes circular encoding as follows.

$$\mathbf{e}_D \leftarrow \mathcal{D}_{\mathbb{Z}, \sigma'}^{(L_T+1)m}, \quad \mathbf{D} = \mathbf{t}^\top (\mathbf{A}_{\text{path}} - (1, \text{bits}(\mathbf{T})) \otimes \mathbf{G}) + \mathbf{e}_D^\top.$$

5. Output $\text{ct} = (\mathbf{c}_B, \mathbf{c}_{\text{msg}}, \mathbf{T}, \mathbf{D})$.

Dec(sk, x, ct, C). The decryption algorithm does the following.

1. Parse $\text{sk} = (\mathbf{K}, \mathbf{r}^\top)$ and $\text{ct} = (\mathbf{A}_t, \mathbf{c}_B, \mathbf{c}_{\text{msg}}, \mathbf{T}, \mathbf{D})$.
2. Compute $\text{ct}_B^\top \cdot \mathbf{K}$ and parse it as $(\mathbf{c}_0^\top \quad \mathbf{c}_{\text{att}}^\top \quad \mathbf{c}_1^\top)$ where $\mathbf{c}_0 \in \mathbb{Z}_q$, $\mathbf{c}_{\text{att}} \in \mathbb{Z}_q^{(L+1)m}$, $\mathbf{c}_1 \in \mathbb{Z}_q^\ell$.
3. Compute $(\mathbf{S}'_r, \mathbf{E}'_r) \leftarrow \text{ComputeAdvice}_r(\bar{\mathbf{A}}_{\text{fhe}}, \mathbf{A}_{\text{circ}}, \mathbf{A}_{\text{path}}, \mathbf{T}, \mathbf{D}, \mathbf{c}_1^\top)$.
4. Set $\mathbf{c}_{\text{circ}}^\top := \mathbf{E}'_r$ and compute $\mathbf{c}_{C,x}^\top \leftarrow \text{UEvalCX}(\mathbf{A}_{\text{att}}, \mathbf{c}_{\text{att}}^\top, \mathbf{A}_{\text{circ}}, \mathbf{c}_{\text{circ}}^\top, C, \mathbf{x}, \mathbf{S}'_r)$.
5. Output $\text{ct}_{\text{msg}}^\top \cdot \mathbf{r} - \mathbf{c}_0^\top - \mathbf{c}_{C,x}^\top \cdot \mathbf{u}$

Parameters. We set our parameters as follows.

$$n = \text{poly}(\lambda), \quad q = 2^{14\lambda} \lambda^{\omega(1)}, \quad p = 2^{10\lambda}, \quad m = O(n \log q), \quad B = 2^{7\lambda},$$

$$\sigma_r = \text{poly}(\lambda), \quad \sigma_s = \sigma' = 2^\lambda, \quad \sigma_0 = 2^{7\lambda}, \quad \sigma_1 = 2^{2\lambda} \lambda^{\omega(1)}, \quad \sigma_{\text{msg}} = 2^{5\lambda}, \quad \sigma_{\text{att}} = 2^{5\lambda},$$

$$\tau = O(\sqrt{(n+1)(m+2) \log q}), \quad \chi_0 = 2^{6\lambda} \lambda^{\omega(1)}, \quad \chi_1 = 2^{2\lambda} \lambda^{\omega(1)}, \quad \chi_{\text{att}} = 2^{3\lambda}$$

where $\chi_0, \chi_1, \chi_{\text{att}}$ appear only in the security proof.

Efficiency. Using the above set parameters, we have

$$|\text{mpk}| = \text{poly}(\lambda, L), \quad |\text{sk}| = \text{poly}(\lambda, L), \quad |\text{ct}| = \text{poly}(\lambda).$$

Correctness We show the correctness step by step below.

- First, we note that

$$\begin{aligned} \mathbf{ct}_B \cdot \mathbf{K} &= \left((\mathbf{s}_0^\top \mid \mathbf{s}^\top \mid \mathbf{t}^\top) \mathbf{B} + \mathbf{e}_B^\top \right) \mathbf{B}_\tau^{-1} \begin{pmatrix} \mathbf{A}_0 \mathbf{r} \\ (\mathbf{A}_{\text{att}} - (1, \mathbf{x}^\top) \otimes \mathbf{G}) \otimes \mathbf{r} \\ \mathbf{A}_r \end{pmatrix} \\ (\mathbf{c}_0^\top \mid \mathbf{c}_{\text{att}}^\top \mid \mathbf{c}_1^\top) &= (\mathbf{s}_0^\top (\mathbf{A}_0 \mathbf{r}) + \mathbf{e}_0^\top \mid \mathbf{s}^\top ((\mathbf{A}_{\text{att}} - (1, \mathbf{x}^\top) \otimes \mathbf{G}) \otimes \mathbf{r}) + \mathbf{e}_{\text{att}}^\top \mid \mathbf{t}^\top (\mathbf{A}_r) + \mathbf{e}_1^\top) \\ \text{where } \|\mathbf{e}_0\| &\leq \tau \|\mathbf{e}'_0\|, \|\mathbf{e}_{\text{att}}\| \leq \tau \|\mathbf{e}'_{\text{att}}\|, \|\mathbf{e}_1\| \leq \tau \|\mathbf{e}'_1\|. \end{aligned}$$

- Next, we have $\|\mathbf{s}_r\| \leq \sigma_s \sigma_r \cdot m \leq 2^\lambda \cdot \text{poly}(\lambda) \cdot m \leq B$ and $q/p = 2^{3\lambda} \lambda^{\omega(1)}$.
- Correctness of ComputeAdvice_r . From our parameter setting, we have $\|\mathbf{e}_1\| \leq \tau \cdot 2^{2\lambda} \lambda^{\omega(1)} \sqrt{\lambda}$, $\|\mathbf{e}_D\|, \|\mathbf{e}'_{\text{fhe}}\| \leq 2^\lambda \sqrt{\lambda}$, where $\mathbf{e}_1, \mathbf{e}_D, \mathbf{e}'_{\text{fhe}}$ are the error terms present in the input to procedure ComputeAdvice_r . Next, we observe the following.

- The error term in Step 3 is $\mathbf{e}_F^\top = (\mathbf{e}')^\top \mathbf{R}_F + \mathbf{e}_D^\top \mathbf{H}_{F,T}$. We have $\|\mathbf{e}_F\| \leq 2^{\lambda + \text{poly}(\log \lambda)} \sqrt{\lambda}$.
- The error term in Step 4 is $\mathbf{e}^\top = \mathbf{e}_1^\top - \mathbf{e}_F^\top \in \mathbb{Z}^{1 \times \ell}$. We have $\|\mathbf{e}\| \leq 2^{3\lambda}$. This implies $\|\mathbf{e}_S\|, \|\mathbf{e}_E\| \leq 2^{3\lambda}$.
- The decryption error in $\bar{\mathbf{S}}_r$ is $\mathbf{err}_{\text{fhe}} = \mathbf{s}_r^\top \mathbf{err}_1 - \text{PRF}_1(\text{sd}, \mathbf{r}) \text{PRF}_2(\text{sd}, \mathbf{r})$. We have $\|\mathbf{err}_{\text{fhe}}\| \leq O(m^2) \cdot \text{poly}(\lambda) 2^{5\lambda} \lambda^{\omega(1)} + m \cdot 2^\lambda \leq 2^{6\lambda}$, where we use the fact that the error introduced by $\text{Rnd}(\cdot)$ is bounded by $q/2p$.
- The decryption error in \mathbf{S}'_r is $(\mathbf{err}'_{\text{fhe}})^\top = \mathbf{s}_r^\top (\mathbf{e}_S + \mathbf{err}_2) - \text{PRF}_1(\text{sd}, \mathbf{r}) \text{PRF}_2(\text{sd}, \mathbf{r})$. We have $\|\mathbf{err}'_{\text{fhe}}\| \leq 2^{5\lambda} + O(m^2) \cdot \text{poly}(\lambda) 2^{5\lambda} \lambda^{\omega(1)} + m \cdot 2^\lambda \leq 2^{6\lambda}$, where we use the fact that the error introduced by $\text{Rnd}(\cdot)$ is bounded by $q/2p$ and $\|\mathbf{s}_r\| \cdot \|\mathbf{e}_S\| \leq 2^{5\lambda}$.
- Next, we note that $\|\mathbf{e}_S\| \leq 2^{3\lambda} \leq \frac{1}{2^\lambda} \cdot \frac{q}{p}$. This implies that $\text{Rnd}(\mathbf{S}_r) = \text{Rnd}(\mathbf{S}_r + \mathbf{e}_S^\top)$ with overwhelming probability.
- The error term in \mathbf{E}'_r is $\mathbf{e}_{\text{circ}}^\top = \text{PRF}_3(\text{sd}, \mathbf{r}) + \mathbf{e}_E^\top$. We have $\|\mathbf{e}_{\text{circ}}\| \leq 2^\lambda + 2^{3\lambda} \leq 2^{-\lambda} B$.

Using the above observations, we note that the procedure ComputeAdvice_r outputs \mathbf{S}'_r and \mathbf{E}'_r where \mathbf{E}'_r encode \mathbf{S}'_r with all but negligible probability.

- Next, note that for $\mathbf{c}_{\text{att}}^\top = \mathbf{s}^\top ((\mathbf{A}_{\text{att}} - (1, \mathbf{x}^\top) \otimes \mathbf{G}) \otimes \mathbf{r}) + \mathbf{e}_{\text{att}}^\top = \mathbf{s}_r^\top (\mathbf{A}_{\text{att}} - (1, \mathbf{x}^\top) \otimes \mathbf{G}) + \mathbf{e}_{\text{att}}^\top$, where $\|\mathbf{s}_r\|, \|\mathbf{e}_{\text{att}}\| \leq B$. So, by the correctness of UEvalCX , we have

$$\mathbf{ct}_{C,x}^\top = \mathbf{s}_r^\top (\mathbf{A}_C - C(\mathbf{x})\mathbf{G}) + \mathbf{e}_{C,x}^\top = \mathbf{s}(\mathbf{I}_{n+1} \otimes \mathbf{r})(\mathbf{A}_C - C(\mathbf{x})\mathbf{G}) + \mathbf{e}_{C,x}^\top,$$

where, from Step 2 of procedure HLL Bootstrapping, it follows that $\|\mathbf{e}_{C,x}\| \leq (m +$

$2^{O(\log n \log \log q)} \cdot m \cdot (\|\text{PRF}_1(\text{sd}, \mathbf{r})^\top\| + \|(\mathbf{s}_r^\top (\mathbf{e}_s + \mathbf{err}_2))^\top\|) + \|\mathbf{e}_{\text{circ}}\| \cdot 2^{O(\log^5 \lambda)}$.
 From our parameter setting and the error analysis above, $\|\mathbf{e}_{C,\mathbf{x}}\| \leq 2^{6\lambda} \leq 2^{-\lambda/2} B$.

– Next, for C, \mathbf{x} such that $C(\mathbf{x}) = 0$, we have the following.

$$\begin{aligned} & \mathbf{ct}_{\text{msg}}^\top \cdot \mathbf{r} - \mathbf{c}_0^\top - \mathbf{ct}_{C,\mathbf{x}}^\top \cdot \mathbf{u} \\ &= (\mathbf{s}_0^\top \mathbf{A}_0 + \mathbf{s}^\top (\mathbf{A}_C \mathbf{u} \otimes \mathbf{I}_m) + \mu \cdot \mathbf{g}^\top + \mathbf{e}_{\text{msg}}^\top) \cdot \mathbf{r} - \mathbf{s}_0^\top (\mathbf{A}_0 \mathbf{r}) - \mathbf{e}_0^\top \\ & \quad - \mathbf{s}^\top (\mathbf{I}_{n+1} \otimes \mathbf{r}) \mathbf{A}_C \mathbf{u} - \mathbf{e}_{C,\mathbf{x}}^\top \mathbf{u} \\ &= \mathbf{s}^\top (\mathbf{A}_C \mathbf{u} \otimes \mathbf{I}_m) \mathbf{r} + \mu \cdot \mathbf{g}^\top \mathbf{r} + \mathbf{e}_{\text{msg}}^\top \mathbf{r} - \mathbf{e}_0^\top - \mathbf{s}^\top (\mathbf{A}_C \mathbf{u} \otimes \mathbf{I}_m) (\mathbf{1} \otimes \mathbf{r}) - \mathbf{e}_{C,\mathbf{x}}^\top \mathbf{u} \\ &\approx \mu \lfloor q/2 \rfloor + \mathbf{e}^\top. \end{aligned}$$

where $\mathbf{e}^\top = \mathbf{e}_{\text{msg}}^\top \mathbf{r} - \mathbf{e}_0^\top - \mathbf{e}_{C,\mathbf{x}}^\top \mathbf{u}$ and $\|\mathbf{e}\| \leq 2^{5\lambda} \sqrt{\lambda} \cdot m \cdot \text{poly}(\lambda) - O(m) \cdot 2^\lambda \sqrt{\lambda} - 2^{6\lambda} \cdot \text{poly}(\lambda) \leq 2^{7\lambda} < q/4$, which is within the bounds of rounding. Hence the decryption outputs μ with all but negligible probability.

4.4.2 Our Assumption

We detail our assumptions next. We introduce a circular version of the Tensor LWE assumption [169] in similar spirit to the Evasive Circular assumption [122].

Assumption 4.13 (Circular Tensor LWE). Let $\mathbf{s} = (\mathbf{s}_1^\top, \dots, \mathbf{s}_n^\top, -\mathbf{1}_m)^\top$ where $\mathbf{s}_i \leftarrow \mathcal{D}_{\mathbb{Z}, \sigma_s}^m$, $\mathbf{r} = (\mathbf{r}_1, \dots, \mathbf{r}_m)$, where $\mathbf{r}_i \leftarrow \mathcal{D}_{\mathbb{Z}, \sigma_r}$ such that $\sum_{i \in [m]} \mathbf{r}_i = \mathbf{1}$. For $i \in [Q]$, set $\mathbf{s}_{\mathbf{r}_i}^\top = \mathbf{s}^\top (\mathbf{I}_{n+1} \otimes \mathbf{r}_i) = (\bar{\mathbf{s}}_{\mathbf{r}_i}^\top, -1) \in \mathbb{Z}^{n+1}$. Suppose

$$\mathbf{A}_{\text{att}} = \begin{pmatrix} \bar{\mathbf{A}}_{\text{att}} \\ \mathbf{a}_{\text{att}}^\top \end{pmatrix} \leftarrow \mathbb{Z}_q^{(n+1) \times (L+1)m}, \bar{\mathbf{A}}_{\text{fhe}} \leftarrow \mathbb{Z}_q^{n \times m}, \mathbf{A}_{\text{circ}} = \begin{pmatrix} \bar{\mathbf{A}}_{\text{circ}} \\ \mathbf{a}_{\text{circ}}^\top \end{pmatrix} \leftarrow \mathbb{Z}_q^{(n+1) \times m_{\text{circ}}}$$

and for $i \in [Q]$, $\sigma_{\text{att}}, \sigma' \ll q$, suppose

$$\begin{aligned} & \mathbf{e}_{\text{att},i} \leftarrow \mathcal{D}_{\mathbb{Z}, \sigma_{\text{att}}}^{(L+1)m}, \mathbf{e}_{\text{fhe},i} \leftarrow [-\sigma', \sigma']^m, \mathbf{e}_{\text{circ},i} \leftarrow [-\sigma', \sigma']^{m_{\text{circ}}} \\ & \delta_{\text{att},i} \leftarrow \mathbb{Z}_q^{(L+1)m}, \delta_{\text{fhe},i} \leftarrow \mathbb{Z}_q^m, \delta_{\text{circ},i} \leftarrow \mathbb{Z}_q^{m_{\text{circ}}}, \Delta_i \leftarrow \mathbb{Z}_q^{(n+1) \times m_s} \\ & \mathbf{R}_i \leftarrow \{0, 1\}^{m \times m_s}, \mathbf{S}_{\mathbf{r}_i} = \begin{pmatrix} \bar{\mathbf{A}}_{\text{fhe}} \\ \bar{\mathbf{s}}_{\mathbf{r}_i}^\top \bar{\mathbf{A}}_{\text{fhe}} + \mathbf{e}_{\text{fhe},i}^\top \end{pmatrix} \mathbf{R}_i - \text{bits}(\mathbf{s}_{\mathbf{r}_i}) \otimes \mathbf{G} \in \mathbb{Z}_q^{(n+1) \times m_s} \end{aligned}$$

where the parameters are functions of λ .

For all $\mathbf{x}_1, \dots, \mathbf{x}_Q \in \{0, 1\}^L$, the *circular tensor* LWE assumption states that

$$\left(1^\lambda, \mathbf{A}_{\text{att}}, \mathbf{A}_{\text{circ}}, \bar{\mathbf{A}}_{\text{fhe}}, \left\{ \begin{array}{l} \bar{\mathbf{s}}_{\mathbf{r}_i}^\top (\bar{\mathbf{A}}_{\text{att}} - (1, \mathbf{x}_i^\top) \otimes \bar{\mathbf{G}}) + \mathbf{e}_{\text{att},i}^\top, \quad \bar{\mathbf{s}}_{\mathbf{r}_i}^\top \bar{\mathbf{A}}_{\text{fhe}} + \mathbf{e}_{\text{fhe},i}^\top, \\ \mathbf{S}_{\mathbf{r}_i}, \quad \bar{\mathbf{s}}_{\mathbf{r}_i}^\top (\bar{\mathbf{A}}_{\text{circ}} - (1, \text{bits}(\mathbf{S}_i)) \otimes \bar{\mathbf{G}}) + \mathbf{e}_{\text{circ},i}^\top, \quad \mathbf{r}_i \end{array} \right\}_{i \in [Q]} \right)_{\lambda \in \mathbb{N}} \quad (4.5)$$

$$\approx \left(1^\lambda, \mathbf{A}_{\text{att}}, \mathbf{A}_{\text{circ}}, \bar{\mathbf{A}}_{\text{fhe}}, \left\{ \delta_{\text{att},i}^\top, \quad \delta_{\text{fhe},i}^\top, \quad \Delta_i, \quad \delta_{\text{circ},i}^\top \right\}_{i \in [Q]} \right)_{\lambda \in \mathbb{N}} \quad (4.6)$$

Remark 16. Here, we choose \mathbf{e}_{fhe} and \mathbf{e}_{circ} from uniform distribution over an interval rather than from usual Gaussian distribution. This change makes little difference to the assumption, since when the modulus q is super-polynomial, we can reduce the problem with Gaussian distribution to the one with uniform distribution with (superpolynomially) larger width and vice versa, due to the smudging (Lemma 2.3).

4.4.3 Security Proof

Theorem 4.14. Assuming evasive LWE (assumption 2.6) for the sampler class \mathcal{SC} induced by the sampler as defined in Equation (4.7), circular tensor LWE (assumption 4.13) and LWE (assumption 2.5), there exists a CP-ABE scheme for unbounded depth and unbounded size circuits which satisfies very selective security (Definition 2.7).

Proof. Suppose the ABE adversary \mathcal{A} with randomness $\text{coins}_{\mathcal{A}}$ queries C and $\mathbf{x}_1, \dots, \mathbf{x}_Q$ such that $C(\mathbf{x}_1) = \dots = C(\mathbf{x}_Q) = 1$. The view of the adversary when bit μ is encoded is:

$$\left(\begin{array}{l} \text{mpk} = (\bar{\mathbf{A}}_{\text{fhe}}, \mathbf{A}_0, \mathbf{A}_{\text{path}}, \mathbf{A}_{\text{att}}, \mathbf{A}_{\text{circ}}, \mathbf{B}, \mathbf{u}), \\ \mathbf{c}_{\mathbf{B}}^\top = (\mathbf{s}_0^\top \mid \mathbf{s}^\top \mid \mathbf{t}^\top) \mathbf{B} + \mathbf{e}_{\mathbf{B}}^\top, \quad \text{ct}_{\text{msg}}^\top = \mathbf{s}_0^\top \mathbf{A}_0 + \mathbf{s}^\top (\mathbf{A}_C \mathbf{u} \otimes \mathbf{I}_m) + \mu \cdot \mathbf{g}^\top + \mathbf{e}_{\text{msg}}^\top \\ \mathbf{T} = \mathbf{A}_t \mathbf{R} - \text{bits}(\mathbf{s}, \text{sd}) \otimes \mathbf{G}, \quad \mathbf{D} = \mathbf{t}^\top (\mathbf{A}_{\text{path}} - (1, \text{bits}(\mathbf{T})) \otimes \mathbf{G}) + \mathbf{e}_{\mathbf{D}}^\top, \quad \{\mathbf{K}_i, \quad \mathbf{r}_i\}_{i \in [Q]} \end{array} \right)$$

where for $i \in [Q]$, we have

$$\mathbf{K}_i = \mathbf{B}_\tau^{-1} \begin{pmatrix} \mathbf{A}_0 \mathbf{r}_i \\ (\mathbf{A}_{\text{att}} - (1, \mathbf{x}^\top) \otimes \mathbf{G}) \otimes \mathbf{r}_i \\ \mathbf{A}_{\mathbf{r}_i} \end{pmatrix}$$

Note that to prove $\text{ct}_{\text{msg}}^\top = \mathbf{s}_0^\top \mathbf{A}_0 + \mathbf{s}^\top (\mathbf{A}_C \mathbf{u} \otimes \mathbf{I}_m) + \mu \cdot \mathbf{g}^\top + \mathbf{e}_{\text{msg}}^\top$ is indistinguishable

from random, it suffices to show that $\mathbf{c}_{\text{msg}}^\top = \mathbf{s}_0^\top \mathbf{A}_0 + \mathbf{s}^\top (\mathbf{A}_C \mathbf{u} \otimes \mathbf{I}_m) + \mathbf{e}_{\text{msg}}^\top$. We prove the pseudorandomness of the above distribution by replacing $\mathbf{c}_{\text{msg}}^\top$ with $\mathbf{c}_{\text{msg}}^\top$.

We invoke the Theorem 2.7 variant of evasive LWE hardness assumption for a matrix \mathbf{B} with Gaussian parameter τ and a sampler Samp that outputs $(\mathbf{S}, \mathbf{P}, \text{aux} = (\text{aux}_1, \text{aux}_2))$ defined as follows.

$$\begin{aligned}
\text{aux}_1 &= \left(\begin{array}{l} \mathbf{c}_{\text{msg}} = (\mathbf{s}_0, \mathbf{s}) \begin{pmatrix} \mathbf{A}_0 \\ \mathbf{A}_C \mathbf{u} \otimes \mathbf{I}_m \end{pmatrix}, \mathbf{T} = \begin{pmatrix} \bar{\mathbf{A}}_{\text{fhe}} \\ \bar{\mathbf{t}}^\top \bar{\mathbf{A}}_{\text{fhe}} + \bar{\mathbf{e}}_{\text{fhe}}^\top \end{pmatrix} \mathbf{R} - \text{bits}(\mathbf{s}, \text{sd}) \otimes \mathbf{G}, \\ \mathbf{D} = \mathbf{t}^\top (\mathbf{A}_{\text{path}} - (1, \text{bits}(\mathbf{T})) \otimes \mathbf{G}) + \mathbf{e}_{\mathbf{D}}^\top \end{array} \right) \\
\text{aux}_2 &= (\mathbf{x}_1, \dots, \mathbf{x}_Q, C, \text{coins}_{\mathcal{A}}, \mathbf{r}_1, \dots, \mathbf{r}_Q, \bar{\mathbf{A}}_{\text{fhe}}, \mathbf{A}_0, \mathbf{A}_{\text{path}}, \mathbf{A}_{\text{att}}, \mathbf{A}_{\text{circ}}, \mathbf{u}) \\
\mathbf{S} &= (\mathbf{s}_0^\top \mid \mathbf{s}^\top \mid \mathbf{t}^\top) \\
\mathbf{P}_0 &= (\mathbf{A}_0 \mathbf{r}_1, \dots, \mathbf{A}_0 \mathbf{r}_Q) \\
\mathbf{P}_1 &= ((\mathbf{A}_{\text{att}} - (1, \mathbf{x}^\top) \otimes \mathbf{G}) \otimes \mathbf{r}_1, \dots, (\mathbf{A}_{\text{att}} - (1, \mathbf{x}^\top) \otimes \mathbf{G}) \otimes \mathbf{r}_Q) \\
\mathbf{P}_2 &= (\mathbf{A}_{\mathbf{r}_1}, \dots, \mathbf{A}_{\mathbf{r}_Q}) \\
\mathbf{P} &= \begin{pmatrix} \mathbf{P}_0 & & \\ & \mathbf{P}_1 & \\ & & \mathbf{P}_2 \end{pmatrix} \tag{4.7}
\end{aligned}$$

We set $\text{aux}_0 = (\mathbf{x}_1, \dots, \mathbf{x}_Q, C, \text{coins}_{\mathcal{A}})$.

By applying evasive LWE, it suffices to show pseudorandomness of the following distribution, given aux_0 ,

$$\left(\begin{array}{l} \text{mpk} = (\bar{\mathbf{A}}_{\text{fhe}}, \mathbf{A}_0, \mathbf{A}_{\text{path}}, \mathbf{A}_{\text{att}}, \mathbf{A}_{\text{circ}}, \mathbf{B}, \mathbf{u}) \\ \mathbf{c}_{\mathbf{B}}^\top = (\mathbf{s}_0^\top \mid \mathbf{s}^\top \mid \mathbf{t}^\top) \mathbf{B} + \mathbf{e}_{\mathbf{B}}^\top, \mathbf{c}_{\text{msg}}^\top = \mathbf{s}_0^\top \mathbf{A}_0 + \mathbf{s}^\top (\mathbf{A}_C \mathbf{u} \otimes \mathbf{I}_m) + \mathbf{e}_{\text{msg}}^\top, \\ \mathbf{T} = \mathbf{A}_{\mathbf{t}} \mathbf{R} - \text{bits}(\mathbf{s}, \text{sd}) \otimes \mathbf{G}, \mathbf{D} = \mathbf{t}^\top (\mathbf{A}_{\text{path}} - (1, \text{bits}(\mathbf{T})) \otimes \mathbf{G}) + \mathbf{e}_{\mathbf{D}}^\top, \\ \{\mathbf{c}_{0,i}^\top = \mathbf{s}_0^\top (\mathbf{A}_0 \mathbf{r}_i) + \mathbf{e}_{0,i}^\top, \mathbf{c}_{\text{att},i}^\top = \mathbf{s}_{\mathbf{r}_i}^\top (\mathbf{A}_{\text{att}} - (1, \mathbf{x}_i^\top) \otimes \mathbf{G}) + \mathbf{e}_{\text{att},i}^\top, \mathbf{c}_{1,i}^\top = \mathbf{t}^\top \mathbf{A}_{\mathbf{r}_i} + \mathbf{e}_{1,i}^\top, \mathbf{r}_i\}_{i \in [Q]} \end{array} \right)$$

where $\mathbf{e}_{0,i}^\top \leftarrow \mathcal{D}_{\mathbb{Z}, \sigma_0}$, $\mathbf{e}_{\text{att},i}^\top \leftarrow \mathcal{D}_{\mathbb{Z}, \sigma_{\text{att}}}^{(L+1)m}$, $\mathbf{e}_{1,i}^\top \leftarrow \mathcal{D}_{\mathbb{Z}, \sigma_1}^\ell$ for $i \in [Q]$.

Hyb_0 . This is the distribution as specified in Section 4.4.3.

Hyb₁. This hybrid is same as Hyb₀, except we compute $\mathbf{c}_{1,i}^\top$ as a function of $\mathbf{T}, \mathbf{D}, \mathbf{S}_{\mathbf{r}_i}, \mathbf{E}_{\mathbf{r}_i}$.

Specifically, we compute

$$\mathbf{c}_{1,i}^\top := \mathbf{D} \mathbf{H}_{\mathbf{F},\mathbf{T}} + (\mathbf{S}_{\mathbf{r}_i}, \mathbf{E}_{\mathbf{r}_i}) + \mathbf{e}_{1,i}^\top, \text{ for } i \in [Q], \quad (4.8)$$

We claim that Hyb₀ and Hyb₁ are statistically indistinguishable. To see this, note that:

- $\mathbf{t}^\top(\mathbf{A}_{\mathbf{r}_i}) = \mathbf{D} \mathbf{H}_{\mathbf{F},\mathbf{T}} + (\mathbf{S}_{\mathbf{r}_i}, \mathbf{E}_{\mathbf{r}_i}) + \mathbf{e}_{\mathbf{F},i}^\top$, where $\|\mathbf{e}_{\mathbf{F},i}\| \leq 2^{\lambda+\text{poly}(\log \lambda)} \sqrt{\lambda} \leq 2^{2\lambda}$, by our parameter setting.
- We have $\mathbf{e}_{1,i}^\top \approx_s \mathbf{e}_{\mathbf{F},i}^\top + \mathbf{e}_{1,i}^\top$ by noise flooding (Theorem 2.3) since $\lambda^{\omega(1)} \|\mathbf{e}_{\mathbf{F},i}\| \leq \lambda^{\omega(1)} 2^{2\lambda} = \chi_1$.

Thus it suffices to show pseudorandomness of the following, given aux_0 ,

$$\left(\begin{array}{l} \text{mpk} = (\bar{\mathbf{A}}_{\text{fhe}}, \mathbf{A}_0, \mathbf{A}_{\text{path}}, \mathbf{A}_{\text{att}}, \mathbf{A}_{\text{circ}}, \mathbf{B}, \mathbf{u}), \\ \mathbf{c}_{\mathbf{B}}^\top = (\mathbf{s}_0^\top \mid \mathbf{s}^\top \mid \mathbf{t}^\top) \mathbf{B} + \mathbf{e}_{\mathbf{B}}^\top, \mathbf{c}_{\text{msg}}^\top = \mathbf{s}_0^\top \mathbf{A}_0 + \mathbf{s}^\top (\mathbf{A}_C \mathbf{u} \otimes \mathbf{I}_m) + \mathbf{e}_{\text{msg}}^\top, \\ \mathbf{T} = \mathbf{A}_{\mathbf{t}} \mathbf{R} - \text{bits}(\mathbf{s}, \text{sd}) \otimes \mathbf{G}, \mathbf{D} = \mathbf{t}^\top (\mathbf{A}_{\text{path}} - (1, \text{bits}(\mathbf{T})) \otimes \mathbf{G}) + \mathbf{e}_{\mathbf{D}}^\top, \\ \{\mathbf{c}_{0,i}^\top = \mathbf{s}_0^\top (\mathbf{A}_0 \mathbf{r}_i) + \mathbf{e}_{0,i}^\top, \mathbf{c}_{\text{att},i}^\top = \mathbf{s}_{\mathbf{r}_i}^\top (\mathbf{A}_{\text{att}} - (1, \mathbf{x}_i^\top) \otimes \mathbf{G}) + \mathbf{e}_{\text{att},i}^\top, \mathbf{S}_{\mathbf{r}_i} + \mathbf{e}_{\mathbf{S},i}^\top, \mathbf{E}_{\mathbf{r}_i} + \mathbf{e}_{\mathbf{E},i}^\top, \mathbf{r}_i\}_{i \in [Q]} \end{array} \right)$$

Hyb₂. This hybrid is same as Hyb₁, except we compute

$$\mathbf{c}_{0,i}^\top := \mathbf{c}_{\text{msg}}^\top \cdot \mathbf{r}_i - \text{UEvalCX}(\mathbf{A}_{\text{att}}, \mathbf{c}_{\text{att},i}^\top, \mathbf{A}_{\text{circ}}, \mathbf{E}_{\mathbf{r}_i} + \mathbf{e}_{\mathbf{E},i}^\top, C, \mathbf{x}, \mathbf{S}'_{\mathbf{r}_i}) \mathbf{u} - \mathbf{s}_{\mathbf{r}_i}^\top \mathbf{G} \mathbf{u} + \mathbf{e}_{0,i}^\top$$

where $\mathbf{S}'_{\mathbf{r}_i} = \text{Rnd}(\mathbf{S}_{\mathbf{r}_i} + \mathbf{e}_{\mathbf{S},i}^\top)$. Note that since $\|\mathbf{e}_{\mathbf{S},i}\| \leq \|\mathbf{e}_{1,i}\| \leq 2^{2\lambda} \lambda^{\omega(1)} \leq \frac{1}{2^{\lambda}} q/p$, Claim 4.7 implies that $\text{Rnd}(\mathbf{S}_{\mathbf{r}_i} + \mathbf{e}_{\mathbf{S},i}^\top) = \text{Rnd}(\mathbf{S}_{\mathbf{r}_i})$ with overwhelming probability, where $\text{Rnd}(\mathbf{S}_{\mathbf{r}_i})$ is encoded in $\mathbf{E}_{\mathbf{r}_i}$.

We claim that Hyb₁ and Hyb₂ are statistically indistinguishable. To see this, note

that:

- we have $\|\mathbf{s}_{r_i}\| \leq B$ and $\|\mathbf{e}_{\text{att},i}\| \leq 2^{5\lambda}\sqrt{\lambda} \leq B$.
- Next, by the correctness of UEvalCX , we have

$$\begin{aligned}
& \mathbf{c}_{\text{msg}}^\top \cdot \mathbf{r}_i - \text{UEvalCX}(\mathbf{A}_{\text{att}}, \mathbf{c}_{\text{att},i}^\top, \mathbf{A}_{\text{circ}}, \mathbf{E}_{r_i} + \mathbf{e}_{\mathbf{E},i}^\top, C, \mathbf{x}, \mathbf{S}_{r_i} + \mathbf{e}_{\mathbf{S},i}^\top) \mathbf{u} - \mathbf{s}_{r_i}^\top \mathbf{G} \mathbf{u} \\
&= (\mathbf{s}_0^\top \mathbf{A}_0 + \mathbf{s}^\top (\mathbf{A}_C \mathbf{u} \otimes \mathbf{I}_m) + \mathbf{e}_{\text{msg}}^\top) \cdot \mathbf{r}_i - \mathbf{s}^\top (\mathbf{I}_{n+1} \otimes \mathbf{r}) (\mathbf{A}_C - C(\mathbf{x})\mathbf{G}) \mathbf{u} + \mathbf{e}_{C,\mathbf{x}}^\top \mathbf{u} - \mathbf{s}_{r_i}^\top \mathbf{G} \mathbf{u} \\
&= (\mathbf{s}_0^\top \mathbf{A}_0 + \mathbf{s}^\top (\mathbf{A}_C \mathbf{u} \otimes \mathbf{I}_m) + \mathbf{e}_{\text{msg}}^\top) \cdot \mathbf{r}_i - \mathbf{s}^\top (\mathbf{A}_C \mathbf{u} \otimes \mathbf{I}_m) (\mathbf{1} \otimes \mathbf{r}) + \mathbf{s}_{r_i}^\top \mathbf{G} \mathbf{u} + \mathbf{e}_{C,\mathbf{x}}^\top \mathbf{u} - \mathbf{s}_{r_i}^\top \mathbf{G} \mathbf{u} \\
&= \mathbf{s}_0^\top \mathbf{A}_0 \mathbf{r}_i + \mathbf{e}_{\text{msg}}^\top \cdot \mathbf{r}_i + \mathbf{e}_{C,\mathbf{x}}^\top \mathbf{u}
\end{aligned}$$

where $\|\mathbf{e}_{C,\mathbf{x}}\| \leq (m + 2)^{O(\log n \log \log q)} \cdot m$.
 $(\|\text{PRF}_1(\text{sd}, \mathbf{r}_i)^\top\| + \|(\mathbf{s}_r^\top (\mathbf{e}_s + \mathbf{err}_2))^\top\|) + \|\mathbf{e}_{\text{circ}}\| \cdot 2^{O(\log^5 \lambda)} \leq 2^{6\lambda}$, by our parameter setting.

- We have $\mathbf{e}_{0,i}^\top \approx_s \mathbf{e}_{\text{msg}}^\top \mathbf{r}_i + \mathbf{e}_{C,\mathbf{x}}^\top \mathbf{u} + \mathbf{e}_{0,i}^\top$ by noise flooding (Theorem 2.3) since $\|\mathbf{e}_{\text{msg}}\| \|\mathbf{r}_i\| + \|\mathbf{e}_{C,\mathbf{x}}\| \|\mathbf{u}\| \leq 2^{7\lambda} \lambda^{\omega(1)} = \chi_0$.

Thus it suffices to show pseudorandomness of the following, given aux_0 ,

$$\left(\begin{aligned} & \text{mpk} = (\bar{\mathbf{A}}_{\text{fhe}}, \mathbf{A}_0, \mathbf{A}_{\text{path}}, \mathbf{A}_{\text{att}}, \mathbf{A}_{\text{circ}}, \mathbf{B}, \mathbf{u}), \\ & \mathbf{c}_{\mathbf{B}}^\top = (\mathbf{s}_0^\top \mid \mathbf{s}^\top \mid \mathbf{t}^\top) \mathbf{B} + \mathbf{e}_{\mathbf{B}}^\top, \quad \mathbf{c}_{\text{msg}}^\top = \mathbf{s}_0^\top \mathbf{A}_0 + \mathbf{s}^\top (\mathbf{A}_C \mathbf{u} \otimes \mathbf{I}_m) + \mathbf{e}_{\text{msg}}^\top, \\ & \mathbf{T} = \mathbf{A}_t \mathbf{R} - \text{bits}(\mathbf{s}, \text{sd}) \otimes \mathbf{G}, \quad \mathbf{D} = \mathbf{t}^\top (\mathbf{A}_{\text{path}} - (\mathbf{1}, \text{bits}(\mathbf{T})) \otimes \mathbf{G}) + \mathbf{e}_{\mathbf{D}}^\top, \\ & \{(\mathbf{c}_{0,i}^\top)' = \mathbf{s}_{r_i}^\top \mathbf{G} \mathbf{u} + \mathbf{e}_{0,i}, \quad \mathbf{c}_{\text{att},i}^\top = \mathbf{s}_{r_i}^\top (\mathbf{A}_{\text{att}} - (\mathbf{1}, \mathbf{x}_i^\top) \otimes \mathbf{G}) + \mathbf{e}_{\text{att},i}^\top, \quad \mathbf{S}_{r_i} + \mathbf{e}_{\mathbf{S},i}^\top, \quad \mathbf{E}_{r_i} + \mathbf{e}_{\mathbf{E},i}^\top, \quad \mathbf{r}_i\}_{i \in [Q]} \end{aligned} \right)$$

Hyb₃. This hybrid is same as **Hyb₂**, except we sample $\mathbf{c}_{\mathbf{B}} \leftarrow \mathbb{Z}_q^{(n+1)(m+2)w}$, $\mathbf{c}_{\text{msg}} \leftarrow \mathbb{Z}_q^m$.

Note that the LWE secret \mathbf{s}_0 does not appear anywhere else. We have $\text{Hyb}_2 \approx_c \text{Hyb}_3$, via the LWE assumption.

Thus it suffices to show pseudorandomness of the following distribution, given

aux_0 ,

$$\left(\begin{array}{l} \text{mpk} = (\bar{\mathbf{A}}_{\text{fhe}}, \mathbf{A}_0, \mathbf{A}_{\text{path}}, \mathbf{A}_{\text{att}}, \mathbf{A}_{\text{circ}}, \mathbf{B}, \mathbf{u}), \mathbf{c}_B \leftarrow \mathbb{Z}_q^{(n+1)(m+2)w}, \mathbf{c}_{\text{msg}} \leftarrow \mathbb{Z}_q^m, \\ \mathbf{T} = \mathbf{A}_t \mathbf{R} - \text{bits}(\mathbf{s}, \text{sd}) \otimes \mathbf{G}, \mathbf{D} = \mathbf{t}^\top (\mathbf{A}_{\text{path}} - (1, \text{bits}(\mathbf{T})) \otimes \mathbf{G}) + \mathbf{e}_D^\top, \\ \{(\mathbf{c}_{0,i}^\top)'\} = \mathbf{s}_{r_i}^\top \mathbf{G} \mathbf{u} + \mathbf{e}_{0,i}, \mathbf{c}_{\text{att},i}^\top = \mathbf{s}_{r_i}^\top (\mathbf{A}_{\text{att}} - (1, \mathbf{x}_i^\top) \otimes \mathbf{G}) + \mathbf{e}_{\text{att},i}^\top \}_{i \in [Q]} \\ \{\mathbf{S}_{r_i} = \text{hct}_{\mathbf{s}_{r_i}}(\text{bits}(\mathbf{s}_{r_i}); \text{PRF}_2(\text{sd}, \mathbf{r}_i)) + \mathbf{e}_{\mathbf{S},i}^\top, \\ \mathbf{E}_{r_i} = \mathbf{s}_{r_i}^\top (\mathbf{A}_{\text{circ}} - (1, \text{bits}(\mathbf{S}_{r_i})) \otimes \mathbf{G}) + \text{PRF}_3(\text{sd}, \mathbf{r}_i) + \mathbf{e}_{\mathbf{E},i}^\top, \mathbf{r}_i\}_{i \in [Q]} \end{array} \right)$$

Hyb_4 . This hybrid is same as Hyb_3 , except we sample $\mathbf{A}_t \leftarrow \mathbb{Z}_q^{(n+1) \times m}$, $\mathbf{D} \leftarrow \mathbb{Z}_q^{1 \times (L_T+1)m}$.

We have $\text{Hyb}_3 \approx_c \text{Hyb}_4$, via the LWE assumption. We show that if there exists an adversary \mathcal{A} who can distinguish between the two hybrids with non-negligible advantage, then there is a reduction \mathcal{B} that breaks LWE security with non-negligible advantage. The reduction is as follows.

1. The LWE challenger sends $\mathbf{A}_{\text{lwe}} \in \mathbb{Z}_q^{n \times (L_T+2)m}$, $\mathbf{b} \in \mathbb{Z}_q^{(L_T+2)m}$ to \mathcal{B} .
2. \mathcal{B} parses $\mathbf{A}_{\text{lwe}} = \begin{pmatrix} \bar{\mathbf{A}}'_{\text{fhe}} & \mathbf{A}'_{\text{path}} \end{pmatrix}$, where $\bar{\mathbf{A}}'_{\text{fhe}} \in \mathbb{Z}_q^{n \times m}$, $\mathbf{A}'_{\text{path}} \in \mathbb{Z}_q^{n \times (L_T+1)m}$ and $\mathbf{b}^\top = [(\bar{\mathbf{b}}')^\top \ (\mathbf{b}'_1)^\top]$, where $\bar{\mathbf{b}}' \in \mathbb{Z}_q^m$, $\mathbf{b}'_1 \in \mathbb{Z}_q^{(L_T+1)m}$ and does the following.
 - Sets $\mathbf{A}_t = \begin{pmatrix} \bar{\mathbf{A}}'_{\text{fhe}} \\ (\bar{\mathbf{b}}')^\top \end{pmatrix}$.
 - Computes $\mathbf{T} = \mathbf{A}_t \mathbf{R} - \text{bits}(\mathbf{s}, \text{sd}) \otimes \mathbf{G}$.
 - Sets $\bar{\mathbf{A}}_{\text{path}} = \mathbf{A}'_{\text{path}} + (1, \text{bits}(\mathbf{T})) \otimes \bar{\mathbf{G}}$ and $\mathbf{A}_{\text{path}} = \begin{pmatrix} \bar{\mathbf{A}}_{\text{path}} \\ \mathbf{a}_{\text{path}}^\top \end{pmatrix}$, where $\mathbf{a}_{\text{path}} \leftarrow \mathbb{Z}_q^{(L_T+1)m}$.
 - Sends \mathbf{T} , $\mathbf{D} = (\mathbf{b}'_1)^\top - \left(\mathbf{a}_{\text{path}}^\top - (1, \text{bits}(\mathbf{T})) \otimes \iota_{n+1}^\top \otimes \mathbf{g} \right)$ to \mathcal{A} .
3. The adversary outputs a bit β' . \mathcal{B} forwards the bit β' to the LWE challenger.

We note that if the LWE challenger sent $\mathbf{b} = \mathbf{t} \mathbf{A}_{\text{lwe}} + \mathbf{e}_{\text{lwe}}$, then \mathcal{B} simulated Hyb_3 with \mathcal{A} else if LWE challenger sent random $\mathbf{b} \leftarrow \mathbb{Z}_q^{(L_T+2)m}$ then \mathcal{B} simulated Hyb_4 with \mathcal{A} .

To see the latter case, we note that if $\mathbf{b} \leftarrow \mathbb{Z}_q^{(L_T+2)m}$ then it implies $\bar{\mathbf{b}}' \leftarrow \mathbb{Z}_q^m$ and $\mathbf{b}'_1 \leftarrow \mathbb{Z}_q^{(L_T+1)m}$. Uniformity of $\bar{\mathbf{b}}' \leftarrow \mathbb{Z}_q^m$ implies $\mathbf{A}_t \leftarrow \mathbb{Z}_q^{(n+1) \times m}$. Randomness of \mathbf{b}'_1 implies randomness of $\mathbf{D} = (\mathbf{b}'_1)^\top - (\mathbf{a}_{\text{path}}^\top - (1, \text{bits}(\mathbf{T})) \otimes \iota_{n+1}^\top \otimes \mathbf{g})$.

Thus it suffices to show pseudorandomness of the following distribution, given aux_0 ,

$$\left(\begin{array}{l} \text{mpk} = (\bar{\mathbf{A}}_{\text{fhe}}, \mathbf{A}_0, \mathbf{A}_{\text{path}}, \mathbf{A}_{\text{att}}, \mathbf{A}_{\text{circ}}, \mathbf{B}, \mathbf{u}), \mathbf{c}_B \leftarrow \mathbb{Z}_q^{(n+1)(m+2)w}, \mathbf{c}_{\text{msg}} \leftarrow \mathbb{Z}_q^m, \\ \mathbf{T} = \mathbf{A}_t \mathbf{R} - \text{bits}(\mathbf{s}, \text{sd}) \otimes \mathbf{G}, \mathbf{D} \leftarrow \mathbb{Z}_q^{1 \times (L_T+1)m}, \\ \{(\mathbf{c}_{0,i}^\top)'\} = \mathbf{s}_{r_i}^\top \mathbf{G} \mathbf{u} + \mathbf{e}_{0,i}, \mathbf{c}_{\text{att},i}^\top = \mathbf{s}_{r_i}^\top (\mathbf{A}_{\text{att}} - (1, \mathbf{x}_i^\top) \otimes \mathbf{G}) + \mathbf{e}_{\text{att},i}^\top \}_{i \in [Q]} \\ \{\mathbf{S}_{r_i} = \text{hct}_{\mathbf{s}_{r_i}}(\text{bits}(\mathbf{s}_{r_i}); \text{PRF}_2(\text{sd}, \mathbf{r}_i)) + \mathbf{e}_{\mathbf{S},i}^\top, \\ \mathbf{E}_{r_i} = \mathbf{s}_{r_i}^\top (\mathbf{A}_{\text{circ}} - (1, \text{bits}(\mathbf{S}_{r_i})) \otimes \mathbf{G}) + \text{PRF}_3(\text{sd}, \mathbf{r}_i) + \mathbf{e}_{\mathbf{E},i}^\top, \mathbf{r}_i \}_{i \in [Q]} \end{array} \right)$$

Hyb₅. This hybrid is same as **Hyb₄**, except we sample $\mathbf{T} \leftarrow \mathbb{Z}_q^{(n+1) \times m_T}$. We have

$\text{Hyb}_4 \approx_s \text{Hyb}_5$ using leftover hash lemma. By leftover hash lemma (Theorem 2.4) we have that the statistical distance between $\mathbf{A}_t \mathbf{R}$ and a uniform matrix $U \leftarrow \mathbb{Z}_q^{(n+1) \times m(\lambda+L)}$ is negligible. This implies that the statistical distance between $\mathbf{A}_t \mathbf{R} - \text{bits}(\mathbf{s}, \text{sd}) \otimes \mathbf{G}$ and $\mathbf{T} \leftarrow \mathbb{Z}_q^{(n+1) \times m_T}$ is negligible. Thus it suffices to show pseudorandomness of the following distribution, given aux_0 ,

$$\left(\begin{array}{l} \text{mpk} = (\bar{\mathbf{A}}_{\text{fhe}}, \mathbf{A}_0, \mathbf{A}_{\text{path}}, \mathbf{A}_{\text{att}}, \mathbf{A}_{\text{circ}}, \mathbf{B}, \mathbf{u}), \mathbf{c}_B \leftarrow \mathbb{Z}_q^{(n+1)(m+2)w}, \mathbf{c}_{\text{msg}} \leftarrow \mathbb{Z}_q^m, \\ \mathbf{T} \leftarrow \mathbb{Z}_q^{(n+1) \times m_T}, \mathbf{D} \leftarrow \mathbb{Z}_q^{1 \times (L_T+1)m}, \\ \{(\mathbf{c}_{0,i}^\top)'\} = \mathbf{s}_{r_i}^\top \mathbf{G} \mathbf{u} + \mathbf{e}_{0,i}, \mathbf{c}_{\text{att},i}^\top = \mathbf{s}_{r_i}^\top (\mathbf{A}_{\text{att}} - (1, \mathbf{x}_i^\top) \otimes \mathbf{G}) + \mathbf{e}_{\text{att},i}^\top \}_{i \in [Q]} \\ \{\mathbf{S}_{r_i} = \text{hct}_{\mathbf{s}_{r_i}}(\text{bits}(\mathbf{s}_{r_i}); \text{PRF}_2(\text{sd}, \mathbf{r}_i)) + \mathbf{e}_{\mathbf{S},i}^\top, \\ \mathbf{E}_{r_i} = \mathbf{s}_{r_i}^\top (\mathbf{A}_{\text{circ}} - (1, \text{bits}(\mathbf{S}_{r_i})) \otimes \mathbf{G}) + \text{PRF}_3(\text{sd}, \mathbf{r}_i) + \mathbf{e}_{\mathbf{E},i}^\top, \mathbf{r}_i \}_{i \in [Q]} \end{array} \right)$$

Hyb₆: In this hybrid, we change all the PRF values computed using sd to random.

We claim that an adversary who can distinguish between **Hyb₅** and **Hyb₆** can be

used to break PRF security⁸. Intuitively, this is because, the PRF seed \mathbf{sd} is now no longer used in the distribution.

In more detail, the reduction queries the $\text{PRF} = (\text{PRF}_1, \text{PRF}_2, \text{PRF}_3)$ challenger with inputs \mathbf{r}_i for $i \in [Q]$ and obtains real or random outputs $\mathbf{y}_{1,i}^\top \in [-\sigma', \sigma']^m$, $\mathbf{y}_{2,i} \in \{0, 1\}^{m^2 n \lceil \log q \rceil}$, $\mathbf{y}_{3,i}^\top \in [-\sigma', \sigma']^{(L_S+1)m}$. It embeds these into the construction of $\mathbf{a}_{\mathbf{r}_i}, \mathbf{S}_{\mathbf{r}_i}, \mathbf{E}_{\mathbf{r}_i}$ as follows:

- Computes $\mathbf{a}_{\mathbf{r}_i}^\top = \bar{\mathbf{s}}_{\mathbf{r}_i}^\top \bar{\mathbf{A}}_{\text{fhe}} + \mathbf{y}_{1,i}$.
- Computes $\mathbf{S}_{\mathbf{r}_i} = \begin{pmatrix} \bar{\mathbf{A}}_{\text{fhe}} \\ \mathbf{a}_{\mathbf{r}_i}^\top \end{pmatrix} \mathbf{y}_{2,i} - \text{bits}(\mathbf{s}_{\mathbf{r}_i}) \otimes \mathbf{G}$, and $\tilde{\mathbf{S}}_{\mathbf{r}_i} = \text{Rnd}(\mathbf{S}_{\mathbf{r}_i})$.
- Computes $\mathbf{E}_{\mathbf{r}_i} = \mathbf{s}_{\mathbf{r}_i}^\top (\mathbf{A}_{\text{circ}} - (1, \text{bits}(\tilde{\mathbf{S}}_{\mathbf{r}_i})) \otimes \mathbf{G}) + \mathbf{y}_{3,i} + \mathbf{e}_{\mathbf{E},i}^\top$.

Thus it suffices to show pseudorandomness of the following distribution, given aux_0 ,

$$\left(\begin{array}{l} \text{mpk} = (\bar{\mathbf{A}}_{\text{fhe}}, \mathbf{A}_0, \mathbf{A}_{\text{path}}, \mathbf{A}_{\text{att}}, \mathbf{A}_{\text{circ}}, \mathbf{B}, \mathbf{u}), \mathbf{c}_{\mathbf{B}} \leftarrow \mathbb{Z}_q^{(n+1)(m+2)w}, \mathbf{c}_{\text{msg}} \leftarrow \mathbb{Z}_q^m, \\ \mathbf{T} \leftarrow \mathbb{Z}_q^{(n+1) \times m_T}, \mathbf{D} \leftarrow \mathbb{Z}_q^{1 \times (L_T+1)m}, \\ \{(\mathbf{c}_{0,i}^\top)' = \mathbf{s}_{\mathbf{r}_i}^\top \mathbf{G} \mathbf{u} + \mathbf{e}_{0,i}, \mathbf{c}_{\text{att},i}^\top = \mathbf{s}_{\mathbf{r}_i}^\top (\mathbf{A}_{\text{att}} - (1, \mathbf{x}_i^\top) \otimes \mathbf{G}) + \mathbf{e}_{\text{att},i}^\top\}_{i \in [Q]} \\ \{\mathbf{S}_{\mathbf{r}_i} = \text{hct}_{\mathbf{s}_{\mathbf{r}_i}}(\text{bits}(\mathbf{s}_{\mathbf{r}_i})) + \mathbf{e}_{\mathbf{S},i}^\top, \mathbf{E}_{\mathbf{r}_i} = \mathbf{s}_{\mathbf{r}_i}^\top (\mathbf{A}_{\text{circ}} - (1, \text{bits}(\mathbf{S}_{\mathbf{r}_i})) \otimes \mathbf{G}) + \mathbf{e}_{\text{circ},i}^\top + \mathbf{e}_{\mathbf{E},i}^\top, \mathbf{r}_i\}_{i \in [Q]} \end{array} \right)$$

where $\mathbf{e}_{\text{circ},i} \leftarrow [-\sigma', \sigma']^{(L_S+1)m}$.

Hyb₇: In this hybrid, we apply circular tensor LWE (Assumption 4.13). We claim

$\text{Hyb}_6 \approx_c \text{Hyb}_7$.

To see this, first we observe the following using $\mathbf{s}_{\mathbf{r}_i}^\top = (\bar{\mathbf{s}}_{\mathbf{r}_i}^\top, -1)$

$$\begin{aligned} - \mathbf{E}_{\mathbf{r}_i} &= \mathbf{s}_{\mathbf{r}_i}^\top (\mathbf{A}_{\text{circ}} - (1, \text{bits}(\mathbf{S}_{\mathbf{r}_i})) \otimes \mathbf{G}) + \mathbf{e}_{\text{circ},i}^\top + \mathbf{e}_{\mathbf{E},i}^\top \\ &= \begin{pmatrix} \bar{\mathbf{s}}_{\mathbf{r}_i}^\top (\bar{\mathbf{A}}_{\text{circ}} - (1, \text{bits}(\mathbf{S}_{\mathbf{r}_i})) \otimes \bar{\mathbf{G}}) + \mathbf{e}_{\text{circ},i}^\top \end{pmatrix} - \\ &\quad \left(\mathbf{a}_{\text{circ}}^\top - (1, \text{bits}(\mathbf{S}_{\mathbf{r}_i})) \otimes \iota_{n+1}^\top \otimes \mathbf{g} \right) + \mathbf{e}_{\mathbf{E},i}^\top. \\ - \mathbf{c}_{\text{att},i}^\top &= \mathbf{s}_{\mathbf{r}_i}^\top (\mathbf{A}_{\text{att}} - (1, \mathbf{x}_i^\top) \otimes \mathbf{G}) + \mathbf{e}_{\text{att},i}^\top \\ &= \left(\bar{\mathbf{s}}_{\mathbf{r}_i}^\top (\bar{\mathbf{A}}_{\text{att}} - (1, \mathbf{x}_i^\top) \otimes \bar{\mathbf{G}}) + \mathbf{e}_{\text{att},i}^\top \right) - \left(\mathbf{a}_{\text{att}}^\top - (1, \mathbf{x}^\top) \otimes \iota_{n+1}^\top \otimes \mathbf{g} \right). \end{aligned}$$

⁸We assume that a single PRF challenger has both PRF_1 and PRF_2 .

Next we note that

- We can change \mathbf{S}_r to random if $\text{hct}(\mathbf{A}_{\mathbf{s}_{r_i}}, \text{bits}(\mathbf{s}_{r_i}))$ is indistinguishable from random.
- We can change \mathbf{E}_r to random if $\bar{\mathbf{s}}_{r_i}^\top (\bar{\mathbf{A}}_{\text{circ}} - (1, \text{bits}(\mathbf{S}_i)) \otimes \bar{\mathbf{G}}) + \mathbf{e}_{\text{circ},i}^\top$ and \mathbf{S}_{r_i} are indistinguishable from random.
- We can change $\mathbf{c}_{\text{att},i}^\top$ to random if $\bar{\mathbf{s}}_{r_i}^\top (\bar{\mathbf{A}}_{\text{att}} - (1, \mathbf{x}_i^\top) \otimes \bar{\mathbf{G}}) + \mathbf{e}_{\text{att},i}^\top$ is indistinguishable from random.

Invoking circular tensor LWE assumption with respect to secret $\mathbf{s}_{r_i}, \mathbf{r}_i$, we achieve randomness of the latter terms in the above.

Thus it suffices to show pseudorandomness of the following distribution, given aux_0 ,

$$\left(\begin{array}{l} \text{mpk} = (\bar{\mathbf{A}}_{\text{fhe}}, \mathbf{A}_0, \mathbf{A}_{\text{path}}, \mathbf{A}_{\text{att}}, \mathbf{A}_{\text{circ}}, \mathbf{B}, \mathbf{u}), \mathbf{c}_B \leftarrow \mathbb{Z}_q^{(n+1)(m+2)w}, \mathbf{c}_{\text{msg}} \leftarrow \mathbb{Z}_q^m, \\ \mathbf{T} \leftarrow \mathbb{Z}_q^{(n+1) \times m_T}, \mathbf{D} \leftarrow \mathbb{Z}_q^{1 \times (L_T+1)m}, \\ \{(\mathbf{c}_{0,i}^\top)'\} \leftarrow \mathbb{Z}_q, \mathbf{c}_{\text{att},i}^\top \leftarrow \mathbb{Z}_q^{(L+1)m}, \mathbf{S}_{r_i} \leftarrow \mathbb{Z}_q^{(n+1) \times m_S}, \mathbf{E}_{r_i} \leftarrow \mathbb{Z}_q^{1 \times (L_S+1)m}, \mathbf{r}_i \}_{i \in [Q]} \end{array} \right)$$

which completes the proof. ■

Theorem 4.15. *Under the LWE assumption, evasive LWE assumption (Assumption 2.6) for the sampler class \mathcal{SC} induced by the sampler as defined in Equation (4.7), and circular tensor assumption (Assumption 4.13), there exists a very selectively secure CP-ABE scheme supporting unbounded depth circuits $\{C : \{0, 1\}^L \rightarrow \{0, 1\}\}$ and one bit message with efficiency $|\text{mpk}| = \text{poly}(\lambda, L)$, $|\text{sk}| = \text{poly}(\lambda, L)$, $|\text{ct}| = \text{poly}(\lambda)$.*

4.5 GENERIC COMPILER: ABE FOR TURING MACHINES AND NL

For conciseness, we first provide a generic construction of ABE that combines many instance of ABE together in Section 4.5.1. This compiler is adapted from [16], who gave it for FE in the bounded key setting. We then construct ABE for TM in Section 4.5.2 and ABE for NL in Section 4.5.3 by instantiating the generic construction.

4.5.1 Generalized Bundling of Functionality

Consider an ABE scheme $\text{ABE} = (\text{ABE.Setup}, \text{ABE.KeyGen}, \text{ABE.Enc}, \text{ABE.Dec})$ for a parameter $\text{prm} = 1^i$, a relation $R_i : \mathcal{X}_i \times \mathcal{Y}_i \rightarrow \{0, 1\} \cup \{\perp\}$ for all $i \in \mathbb{N}$ and a message space \mathcal{M} . Using such ABE, we will construct a new ABE with ciphertext attribute space \mathcal{A} , key attribute space \mathcal{B} and message space \mathcal{M} . We assume that there exist efficiently computable maps $\mathcal{S} : \mathbb{N} \rightarrow 2^{\mathbb{N}}$ and $\mathcal{T} : \mathbb{N} \rightarrow 2^{\mathbb{N}}$ such that $\max \mathcal{S}(n)$ and $\max \mathcal{T}(n)$ can be bounded by some fixed polynomial in n . We also assume that there exist maps f with domain \mathcal{A} and g with domain \mathcal{B} such that

$$f(x) \in \prod_{i \in \mathcal{S}(|x|)} \mathcal{X}_i \quad \text{and} \quad g(y) \in \prod_{i \in \mathcal{T}(|y|)} \mathcal{Y}_i,$$

where $|x|$ and $|y|$ are the lengths of x and y as binary strings. Namely, f and g are maps such that

$$f : \mathcal{A} \ni x \mapsto \{f(x)_i \in \mathcal{X}_i\}_{i \in \mathcal{S}(|x|)}, \quad g : \mathcal{B} \ni y \mapsto \{g(y)_i \in \mathcal{Y}_i\}_{i \in \mathcal{T}(|y|)}.$$

Here, we require that the length of $|f(x)|_i$ and $|g(y)|_i$ can be computed from the length of $|x|$ alone and they do not depend on the actual value of x . In this setting, we can construct an ABE scheme $\text{Bd-ABE} = (\text{Setup}, \text{KeyGen}, \text{Enc}, \text{Dec})$ for a two input function $R^{\text{bndl}} : \mathcal{A} \times \mathcal{B} \rightarrow \{0, 1\}^*$ defined in the following

$$R^{\text{bndl}}(x, y) = \begin{cases} 0, & \text{if } R_i(f(x)_i, g(y)_i) = 0 \text{ for all } i \in \mathcal{S}(|x|) \cap \mathcal{T}(|y|), \\ 1, & \text{if } R_i(f(x)_i, g(y)_i) = 1 \text{ for all } i \in \mathcal{S}(|x|) \cap \mathcal{T}(|y|). \end{cases} \quad (4.9)$$

where $f(x)_i \in \mathcal{X}_i$ and $g(y)_i \in \mathcal{Y}_i$ are the i -th entries of $f(x)$ and $g(y)$, respectively.

Ingredients. We now describe the underlying building blocks used to obtain our ABE construction:

1. An ABE scheme $\text{ABE} = (\text{ABE.Setup}, \text{ABE.KeyGen}, \text{ABE.Enc}, \text{ABE.Dec})$ for a parameter $\text{prm} = 1^i$, a relation $R_i : \mathcal{X}_i \times \mathcal{Y}_i \rightarrow \{0, 1\} \cup \{\perp\}$ for $i \in \mathbb{N}$ and a message space \mathcal{M} .

2. A garbled circuit scheme $GC = (GC.Garble, GC.Eval)$. We assume that a label is represented by a binary string and denote its length by $L(\lambda, |C|)$, where C is the circuit being garbled. We can instantiate it by Yao's garbled circuit [173], which can be based on any one-way function.
3. An IBE scheme $IBE = (IBE.Setup, IBE.Enc, IBE.KeyGen, IBE.Dec)$ with IND-CPA security whose identity space and message space are $\{0, 1\}^*$. We assume that the key generation algorithm is deterministic. This is without loss of generality, since we can use PRF to derandomize the key generation algorithm. We can instantiate IBE from various standard assumptions including LWE [6, 71], CDH, and Factoring [83].

Construction. Here we provide the description of the construction of Bd-ABE = (Setup, KeyGen, Enc, Dec) for R^{bndl} above.

Setup(1^λ) \rightarrow (mpk, msk). On input the security parameter λ , do the following.

1. Run $(IBE.mpk, IBE.msk) \leftarrow IBE.Setup(1^\lambda)$.
2. Output the master key pair as $(mpk, msk) := (IBE.mpk, IBE.msk)$.

KeyGen(msk, y) $\rightarrow sk_y$. On input master secret key $msk = IBE.msk$, a key attribute $y \in B$, do the following.

1. Compute $\mathcal{T}(|y|) \subseteq \mathbb{N}$, where $|y|$ is the length of y as a binary string.
2. Run $(ABE.mpk_i, ABE.msk_i) \leftarrow ABE.Setup(1^\lambda, 1^i)$ for $i \in \mathcal{T}(|y|)$.
3. Compute $g(y) = \{g(y)_i \in \mathcal{Y}_i\}_{i \in \mathcal{T}(|y|)}$.
4. For $i \in \mathcal{T}(|y|)$, compute

$$ABE.sk_i \leftarrow ABE.KeyGen(ABE.msk_i, g(y)_i).$$

5. Let $\ell_i := |ABE.mpk_i|$. For all $i \in \mathcal{T}(|y|)$ and $j \in \ell_i$, generate a secret key as

$$IBE.sk_{i,j} \leftarrow IBE.KeyGen(IBE.msk, (i, j, ABE.mpk_{i,j}))$$

where $ABE.mpk_{i,j}$ is the j -th bit of $ABE.mpk_i \in \{0, 1\}$ as a binary string.

6. Output

$$\mathbf{sk}_y = \left(\mathcal{T}(|y|), \left\{ \text{ABE.sk}_i, g(y)_i, \left\{ \text{IBE.sk}_{i,j} \right\}_{j \in [\ell_i]} \right\}_{i \in \mathcal{T}(|y|)} \right). \quad (4.10)$$

$\text{Enc}(\text{mpk}, x, \mu) \rightarrow \text{ct}_x$. On input the encryption key $\text{mpk} = \text{IBE.mpk}$, a ciphertext attribute $x \in \mathcal{A}$, and a message $\mu \in \mathcal{M}$, do the following.

1. Compute $\mathcal{S}(|x|) \subset \mathbb{N}$, where $|x|$ is the length of x as a binary string.
2. Compute $f(x) = \{f(x)_i\}_{i \in \mathcal{S}(|x|)}$.
3. Do the following for all $i \in \mathcal{S}(|x|)$.
 - a) Compute the length ℓ_i of ABE.mpk_i . This is done without knowing ABE.mpk_i .
 - b) Sample a randomness r_i for the encryption algorithm $\text{ABE.Enc}(\text{ABE.mpk}_i, f(x)_i, \mu; r_i)$. This is done without knowing ABE.mpk_i .
 - c) Define a circuit

$$E_i(\cdot) := \text{ABE.Enc}(\cdot, f(x)_i, \mu; r_i)$$

that takes as input a string $\text{str} \in \{0, 1\}^{\ell_i}$ and outputs $\text{ABE.Enc}(\text{str}, f(x)_i, \mu; r_i)$, where str is interpreted as a master public key of the ABE.

- d) Generate a garbled circuit

$$\{\text{lab}_{i,j,b}\}_{j \in [\ell_i], b \in \{0,1\}} \leftarrow \text{GC.Garble}(1^\lambda, E_i).$$

- e) For all $j \in [\ell_i]$ and $b \in \{0, 1\}$, compute

$$\text{IBE.ct}_{i,j,b} \leftarrow \text{IBE.Enc}(\text{IBE.mpk}, (i, j, b), \text{lab}_{i,j,b}).$$

4. Output

$$\text{ct}_x = \left(\mathcal{S}(|x|), \{f(x)_i\}_{i \in \mathcal{S}(|x|)}, \{\text{IBE.ct}_{i,j,b}\}_{i \in \mathcal{S}(|x|), j \in [\ell_i], b \in \{0,1\}} \right). \quad (4.11)$$

$\text{Dec}(\text{sk}_y, y, \text{ct}_x, x) \rightarrow \mu' \setminus \perp$. On input a secret key sk_y , key attribute y , a ciphertext ct_x , and ciphertext attribute x , do the following.

1. Parse the secret key sk_y as Eq. (4.10) and the ciphertext ct_x as Eq. (4.11).
2. For all $i \in \mathcal{S}(|x|) \cap \mathcal{T}(|y|)$ do the following.
 - a) Retrieve $ABE.mpk_i$ from the $ABE.sk_i$.
 - b) For all $j \in [\ell_i]$ compute $lab'_{i,j} := IBE.Dec(IBE.sk_{i,j}, ABE.mpk_{i,j}, IBE.ct_{i,j}, ABE.mpk_{i,j})$.
 - c) Compute $c_i := GC.Eval(\{lab'_{i,j}\}_{j \in [\ell_i]})$.
 - d) Compute $z_i := ABE.Dec(ABE.sk_i, g(y)_i, c_i, f(x)_i)$.
3. Output $\{z_i\}_{i \in \mathcal{S}(|x|) \cap \mathcal{T}(|y|)}$.

Correctness. We now show that the above construction is correct via the following theorem.

Theorem 4.16. *Suppose ABE, GC and IBE schemes are correct. Then the Bd-ABE scheme is correct.*

Proof. We observe that $lab'_{i,j} = lab_{i,j,ABE.mpk_{i,j}}$ holds for all $lab'_{i,j}$ recovered in Step 2b of the decryption algorithm by the correctness of IBE, since the ciphertext and secret key are both generated with respect to the identity $(i, j, ABE.mpk_{i,j})$. Then, by the correctness of GC, we have $c_i = ABE.Enc(ABE.mpk_i, f(x)_i, \mu; r_i)$ for all $i \in \mathcal{S}(|x|) \cap \mathcal{T}(|y|)$ recovered in Step 2c of the decryption algorithm. Finally, if $R^{bndl}(x, y) = 0$, then by the definition of R^{bndl} , we have $R_i(f(x)_i, g(y)_i) = 0$ for all $i \in \mathcal{S}(|x|) \cap \mathcal{T}(|y|)$ and thus by the correctness of ABE, we get $z_i = \mu$ for all z_i recovered in Step 2d of the decryption algorithm as desired. ■

Security. We prove the security of the Bd-ABE scheme via the following theorem.

Theorem 4.17. *Suppose GC is a secure garbled circuit scheme and IBE is IND-CPA secure identity-based encryption scheme. Then, assuming ABE is Sel-IND secure (Definition 2.5), so is Bd-ABE. Furthermore, if ABE is VerSel-IND secure*

(Definition 2.7), so is Bd-ABE⁹.

Proof. Here we prove the statement for Sel-IND security. The proof proceeds via a sequence of hybrid games between the challenger and a PPT adversary \mathcal{A} .

Hyb₀. This is the real world with $\beta = 0$, i.e., the challenge ciphertext is computed using the message μ_0 . We write the complete game here to set up the notations and easy reference in later hybrids.

1. \mathcal{A} outputs the challenge ciphertext attribute $x \in \mathcal{A}$.
2. The challenger generates $(\text{IBE.mpk}, \text{IBE.msk}) \leftarrow \text{IBE.Setup}(1^\lambda)$, sets $\text{mpk} = \text{IBE.mpk}$ and sends it to \mathcal{A} .
3. **Key Queries:** The adversary can make key queries, before and after challenge query, in an arbitrary order. For each key query $y \in \mathcal{B}$, the challenger does the following.
 - It computes the mapping $\mathcal{T}(|y|)$ and generates $(\text{ABE.mpk}_i, \text{ABE.msk}_i) \leftarrow \text{ABE.Setup}(1^\lambda, 1^i)$ for $i \in \mathcal{T}(|y|)$.
 - It computes the mapping $g(y)$ and generates $\text{ABE.sk}_i \leftarrow \text{ABE.KeyGen}(\text{ABE.msk}_i, g(y)_i)$ for all $i \in \mathcal{T}(|y|)$.
 - It sets $\ell_i := |\text{ABE.mpk}_i|$ and computes $\text{IBE.sk}_{i,j} \leftarrow \text{IBE.KeyGen}(\text{IBE.msk}, (i, j, \text{ABE.mpk}_{i,j}))$ for all $i \in \mathcal{T}(|y|)$ and $j \in \ell_i$.
 - It returns $\text{sk}_y = (\mathcal{T}(|y|), g(y), \{\text{ABE.sk}_i\}_{i \in \mathcal{T}(|y|)}, \{\text{IBE.sk}_{i,j}\}_{i \in \mathcal{T}(|y|), j \in \ell_i})$.
4. **Challenge Query:** \mathcal{A} outputs a pair of equal length messages $(\mu_0, \mu_1) \in \mathcal{M}^2$. The challenger does the following.
 - It computes $\mathcal{S}(|x|)$ and $f(x) = \{f(x)_i\}_{i \in \mathcal{S}(|x|)}$.
 - For all $i \in \mathcal{S}(|x|)$,
 - It samples the randomness r_i and defines the circuit $E_i(\cdot) := \text{ABE.Enc}(\cdot, f(x)_i, \mu_0; r_i)$ as in the construction.

⁹In the latter case, the adversary can see the public key before declaring all the key queries.

- It generates $\{\text{lab}_{i,j,b}\}_{j \in [\ell_i], b \in \{0,1\}} \leftarrow \text{GC.Garble}(1^\lambda, E_i)$.
- It computes $\text{IBE.ct}_{i,j,b} \leftarrow \text{IBE.Enc}(\text{IBE.mpk}, (i, j, b), \text{lab}_{i,j,b})$.
- It returns $\text{ct}_x = (\mathcal{S}(|x|), f(x), \{\text{IBE.ct}_{i,j,b}\}_{i \in \mathcal{S}(|x|), j \in [\ell_i], b \in \{0,1\}})$.

5. \mathcal{A} outputs a guess bit β' .

Hyb₁. In this hybrid, we change the way challenge query is answered. In particular, the challenger computes $\text{IBE.ct}_{i,j,b} \leftarrow \text{IBE.Enc}(\text{IBE.mpk}, (i, j, b), \text{lab}_{i,j,\text{ABE.mpk}_{i,j}})$, for all $i \in \mathcal{S}(|x|)$, $j \in [\ell_i]$, and $b \in \{0, 1\}$, where $\text{ABE.mpk}_{i,j}$ is the j -th bit of ABE.mpk_i .

Note that we encrypt the same label for both $b = 0$ and $b = 1$.

Hyb₂. In this hybrid, we further change the way challenge query is answered. In particular, we change the way $\text{lab}_{i,j,\text{ABE.mpk}_{i,j}}$ is generated. The challenger does the following for all $i \in \mathcal{S}(|x|)$,

- It computes $\text{ABE.ct}_i \leftarrow \text{ABE.Enc}(\text{ABE.mpk}_i, f(x)_i, \mu_0)$.
- It generates $\{\text{lab}_{i,j}\}_{j \in [\ell_i]} \leftarrow \text{GC.Sim}(1^\lambda, 1^{\ell_i}, 1^{|E_i(\cdot)|}, \text{ABE.ct}_i)$.
- It sets $\text{lab}_{i,j,\text{ABE.mpk}_{i,j}} = \{\text{lab}_{i,j}\}$ for all $j \in [\ell_i]$.

Hyb₃. In this hybrid, we further change the way challenge query is answered. In particular, the challenger computes $\text{ABE.ct}_i \leftarrow \text{ABE.Enc}(\text{ABE.mpk}_i, f(x)_i, \mu_1)$ for all $i \in \mathcal{S}(|x|)$.

The following hybrids are unwinding of the preceding hybrids.

Hyb₄. In this hybrid, we change the way challenge query is answered. In particular, we change the way $\text{lab}_{i,j,\text{ABE.mpk}_{i,j}}$ is generated. The challenger does the following for all $i \in \mathcal{S}(|x|)$.

- It samples an encryption randomness r_i and define $E_i(\cdot) = \text{ABE.Enc}(\cdot, f(x)_i, \mu_1; r_i)$.
- It generates $\{\text{lab}_{i,j,b}\}_{j \in [\ell_i], b \in \{0,1\}} \leftarrow \text{GC.Garble}(1^\lambda, E_i)$.
- It sets $\text{lab}_{i,j,\text{ABE.mpk}_{i,j}} = \{\text{lab}_{i,j,b}\}$ for all $j \in [\ell_i]$, where $b = \text{ABE.mpk}_{i,j}$.

Hyb₅. In this hybrid, we further change the way challenge query is answered. In particular, the challenger computes $\text{IBE.ct}_{i,j,b} \leftarrow \text{IBE.Enc}(\text{IBE.mpk}, (i, j, b), \text{lab}_{i,j,b})$, for all $i \in \mathcal{S}(|x|)$, $j \in [\ell_i]$, and $b \in \{0, 1\}$.

This is the real world with $\beta = 1$.

Indistinguishability of hybrids We now show that the consecutive hybrids are indistinguishable.

Claim 4.18. Assume that IBE is IND-CPA secure. Then $\text{Hyb}_0 \approx_c \text{Hyb}_1$.

Proof. We show that if \mathcal{A} can distinguish between Hyb_0 and Hyb_1 with non-negligible advantage ϵ , then there exists a PPT adversary \mathcal{B} against the IND-CPA security of IBE scheme with advantage ϵ . The reduction is as follows.

1. The IBE challenger generates $(\text{IBE.mpk}, \text{IBE.msk}) \leftarrow \text{IBE.Setup}(1^\lambda)$ and samples a bit $\hat{\beta} \leftarrow \{0, 1\}$. It sends IBE.mpk to \mathcal{B} .
2. \mathcal{B} invokes \mathcal{A} . \mathcal{A} outputs the challenge ciphertext attribute $x \in A$. \mathcal{B} sets $\text{mpk} = \text{IBE.mpk}$ and forwards it to \mathcal{A} .
3. **Key Queries:** For a key query $y \in B$, \mathcal{B} does the following.
 - It computes the mapping $\mathcal{T}(|y|)$ and generates $(\text{ABE.mpk}_i, \text{ABE.msk}_i) \leftarrow \text{ABE.Setup}(1^\lambda, 1^i)$ for $i \in \mathcal{T}(|y|)$.
 - It computes the mapping $g(y)$ and generates $\text{ABE.sk}_i \leftarrow \text{ABE.KeyGen}(\text{ABE.msk}_i, g(y)_i)$ for all $i \in \mathcal{T}(|y|)$.
 - It sets $\ell_i := |\text{ABE.mpk}_i|$ and sends secret key query for identity $(i, j, \text{ABE.mpk}_{i,j})$ for all $i \in \mathcal{T}(|y|)$ and $j \in \ell_i$ to the IBE challenger. The

challenger returns $\{\text{IBE.sk}_{i,j}\}_{i \in \mathcal{T}(|y|), j \in \ell_i}$.

- It returns $\text{sk}_y = (\mathcal{T}(|y|), g(y), \{\text{ABE.sk}_i\}_{i \in \mathcal{T}(|y|)}, \{\text{IBE.sk}_{i,j}\}_{i \in \mathcal{T}(|y|), j \in \ell_i})$ to \mathcal{A} .

4. **Challenge Query:** \mathcal{A} outputs a pair of equal length messages (μ_0, μ_1) . \mathcal{B} does the following.

- It sets $\mu = \mu_0$ and computes $\{\text{lab}_{i,j,b}\}_{j \in [\ell_i], b \in \{0,1\}}$, for all $i \in \mathcal{S}(|x|)$, as in the honest encryption algorithm.
- It computes $\{\text{IBE.ct}_{i,j,b}\}_{i \in \mathcal{S}(|x|), j \in [\ell_i], b \in \{0,1\}}$ as follows.
 - If $b = \text{ABE.mpk}_{i,j}$, it honestly encrypts $\text{lab}_{i,j,b}$ to obtain $\text{IBE.ct}_{i,j,b}$.
 - If $b = 1 - \text{ABE.mpk}_{i,j}$, \mathcal{B} submits identity (i, j, b) and messages $(\text{lab}_{i,j,b}, \text{lab}_{i,j,1-b})$ to the IBE challenger. The challenger encrypts

$$\text{IBE.ct}_{i,j,b} \leftarrow \begin{cases} \text{IBE.Enc}(\text{IBE.mpk}, (i, j, b), \text{lab}_{i,j,b}), & \text{if } \hat{\beta} = 0 \\ \text{IBE.Enc}(\text{IBE.mpk}, (i, j, b), \text{lab}_{i,j,1-b}), & \text{if } \hat{\beta} = 1 \end{cases}$$

and returns $\text{IBE.ct}_{i,j,b}$ to \mathcal{B} .

Note that the same label $\text{lab}_{i,j,\text{ABE.mpk}_{i,j}}$ is encrypted for identities $(i, j, 0)$ and $(i, j, 1)$ if $\beta = 1$.

- It returns $\text{ct}_x = (\mathcal{S}(|x|), f(x), \{\text{IBE.ct}_{i,j,b}\}_{i \in \mathcal{S}(|x|), j \in [\ell_i], b \in \{0,1\}})$ to \mathcal{A} .

5. \mathcal{A} outputs a guess bit β' . \mathcal{B} forwards β' to the IBE challenger.

We observe that if the IBE challenger samples $\hat{\beta} = 0$, then \mathcal{B} simulated Hyb_0 , else Hyb_1 with \mathcal{A} . Hence, advantage of $\mathcal{B} = |\Pr(\beta' = 1 | \hat{\beta} = 0) - \Pr(\beta' = 1 | \hat{\beta} = 1)| = |\Pr(\beta' = 1 | \text{Hyb}_0) - \Pr(\beta' = 1 | \text{Hyb}_1)| = \epsilon$ (by assumption).

Admissibility of \mathcal{B} . Observe that \mathcal{B} submits identities of the form $(i, j, \text{ABE.mpk}_{i,j})$ for secret key queries and identities of the form $(i, j, 1 - \text{ABE.mpk}_{i,j})$ for encryption queries. So, \mathcal{B} does not make a secret key query for an identity that is also submitted to the IBE challenger for an encryption query. This establishes the admissibility of \mathcal{B} . ■

Claim 4.19. Assume that GC is secure. Then $\text{Hyb}_1 \approx_c \text{Hyb}_2$.

Proof. We show that if \mathcal{A} can distinguish between Hyb_1 and Hyb_2 with non-negligible advantage ϵ , then there exists a PPT adversary \mathcal{B} against the security of GC scheme with advantage ϵ . The reduction is as follows.

1. The GC challenger samples a bit $\hat{\beta} \leftarrow \{0, 1\}$ and starts the game with \mathcal{B} .
2. \mathcal{B} invokes \mathcal{A} . \mathcal{A} outputs the challenge ciphertext attribute $x \in \mathcal{A}$.
3. \mathcal{B} generates $(\text{IBE.mpk}, \text{IBE.msk}) \leftarrow \text{IBE.Setup}(1^\lambda)$, sets $\text{mpk} = \text{IBE.mpk}$ and forwards it to \mathcal{A} .
4. **Key Queries:** For a key query $y \in \mathcal{B}$, \mathcal{B} computes $\text{sk}_y = (\mathcal{T}(|y|), g(y), \{\text{ABE.sk}_i\}_{i \in \mathcal{T}(|y|)}, \{\text{IBE.sk}_{i,j}\}_{i \in \mathcal{T}(|y|), j \in \ell_i})$ as in Hyb_0 .
5. **Challenge Query:** \mathcal{A} outputs a pair of equal length messages (μ_0, μ_1) . \mathcal{B} does the following for all $i \in \mathcal{S}(|x|)$.
 - It samples r_i and defines a circuit $E_i(\cdot) := \text{ABE.Enc}(\cdot, f(x)_i, \mu_0; r_i)$.
 - It submits the circuit $E_i(\cdot)$ and ABE.mpk_i to the GC challenger. The challenger does the following.
 - If $\hat{\beta} = 0$, it computes $\{\text{lab}'_{i,j,b}\}_{j \in [\ell_i], b \in \{0,1\}} \leftarrow \text{GC.Garble}(1^\lambda, E_i)$ and sets $\text{lab}_{i,j,\text{ABE.mpk}_{i,j}} = \text{lab}'_{i,j,\text{ABE.mpk}_{i,j}}$, where $\text{ABE.mpk}_{i,j}$ is the j -th bit of ABE.mpk_i .
 - If $\hat{\beta} = 1$, it computes $\{\text{lab}'_{i,j}\}_{j \in [\ell_i]} \leftarrow \text{GC.Sim}(1^\lambda, 1^{\ell_i}, 1^{|E_i(\cdot)|}, E_i(\text{ABE.mpk}_i))$ and sets $\text{lab}_{i,j,\text{ABE.mpk}_{i,j}} = \text{lab}'_{i,j}$.
 - It returns $\text{lab}_{i,j,\text{ABE.mpk}_{i,j}}$ to \mathcal{B} .
 - It computes $\text{IBE.ct}_{i,j,b} \leftarrow \text{IBE.Enc}(\text{IBE.mpk}, (i, j, b), \text{lab}_{i,j,\text{ABE.mpk}_{i,j}})$, for all $i \in \mathcal{S}(|x|)$, $j \in [\ell_i]$, and $b \in \{0, 1\}$.
 - It returns $\text{ct}_x = (\mathcal{S}(|x|), f(x), \{\text{IBE.ct}_{i,j,b}\}_{i \in \mathcal{S}(|x|), j \in [\ell_i], b \in \{0,1\}})$ to \mathcal{A} .
6. \mathcal{A} outputs a guess bit β' . \mathcal{B} forwards β' to the GC challenger.

We observe that if the GC challenger samples $\hat{\beta} = 0$, then \mathcal{B} simulated the distribution Hyb_1 , else Hyb_2 with \mathcal{A} . Hence, advantage of $\mathcal{B} = |\Pr(\beta' = 1 | \hat{\beta} = 0) - \Pr(\beta' = 1 | \hat{\beta} = 1)| = |\Pr(\beta' = 1 | \text{Hyb}_1) - \Pr(\beta' = 1 | \text{Hyb}_2)| = \epsilon$ (by assumption). ■

Claim 4.20. Assume that ABE is Sel-IND secure. Then $\text{Hyb}_2 \approx_c \text{Hyb}_3$.

Proof. We show that if \mathcal{A} can distinguish between Hyb_2 and Hyb_3 with non-negligible advantage ϵ , then there exists a PPT adversary \mathcal{B} against the security of ABE scheme with advantage ϵ . In this game, the ABE challenger is the challenger corresponding to all the i -th instance of ABE such that $i \in \mathcal{S}(|x|)$. For $i \notin \mathcal{S}(|x|)$, the reduction \mathcal{B} itself generates such i -th instance ABE. The reduction is as follows.

1. The ABE challenger samples a bit $\hat{\beta} \leftarrow \{0, 1\}$ and starts the game with \mathcal{B} .
2. \mathcal{B} invokes \mathcal{A} . \mathcal{A} outputs the challenge ciphertext attribute $x \in \mathcal{A}$.
3. \mathcal{B} generates $(\text{IBE.mpk}, \text{IBE.msk}) \leftarrow \text{IBE.Setup}(1^\lambda)$, sets $\text{mpk} = \text{IBE.mpk}$ and forwards it to \mathcal{A} . \mathcal{B} also computes $\mathcal{S}(|x|)$ and $f(x) = \{f(x)_i\}_{i \in \mathcal{S}(|x|)}$.
4. \mathcal{B} sends the challenge ciphertext attributes $f(x)_i$ for the i -th instance of ABE, where $i \in \mathcal{S}(|x|)$, to the ABE challenger. The ABE challenger generates $(\text{ABE.mpk}_i, \text{ABE.msk}_i) \leftarrow \text{ABE.Setup}(1^\lambda, 1^i)$ for $i \in \mathcal{S}(|x|)$ and returns $\{\text{ABE.mpk}_i\}_{i \in \mathcal{S}(|x|)}$ to \mathcal{B} .
5. **Key Queries:** For a key query $y \in \mathcal{B}$, \mathcal{B} does the following.
 - It computes the mapping $\mathcal{T}(|y|)$ and for all $i \in \mathcal{T}(|y|) \setminus \mathcal{S}(|x|)$, it generates $(\text{ABE.mpk}_i, \text{ABE.msk}_i) \leftarrow \text{ABE.Setup}(1^\lambda, 1^i)$.
 - It computes the mapping $g(y) = \{g(y)_i\}_{i \in \mathcal{T}(|y|)}$. For all $i \in \mathcal{T}(|y|) \cap \mathcal{S}(|x|)$, it sends a key query $g(y)_i$ for the i -th instance of ABE to its challenger and gets back ABE.sk_i . For $i \in \mathcal{T}(|y|) \setminus \mathcal{S}(|x|)$, \mathcal{B} computes ABE.sk_i itself.
 - It sets $\ell_i := |\text{ABE.mpk}_i|$ and computes $\text{IBE.sk}_{i,j} \leftarrow \text{IBE.KeyGen}(\text{IBE.msk}, (i, j, \text{ABE.mpk}_{i,j}))$ for all $i \in \mathcal{T}(|y|)$ and $j \in \ell_i$.
 - It returns $\text{sk}_y = (\mathcal{T}(|y|), g(y), \{\text{ABE.sk}_i\}_{i \in \mathcal{T}(|y|)}, \{\text{IBE.sk}_{i,j}\}_{i \in \mathcal{T}(|y|), j \in \ell_i})$ to \mathcal{A} .
6. **Challenge Query:** \mathcal{A} outputs a pair of equal length messages (μ_0, μ_1) . \mathcal{B} does the following.
 - For all $i \in \mathcal{S}(|x|)$,
 - It sends challenge ciphertext query μ_0, μ_1 for the i -th instance of ABE to the challenger. The challenger computes and returns $\text{ABE.ct}_i \leftarrow \text{ABE.Enc}(\text{ABE.mpk}_i, f(x)_i, \mu_{\hat{\beta}})$.
 - It generates $\{\text{lab}_{i,j}\}_{j \in [\ell_i]} \leftarrow \text{GC.Sim}(1^\lambda, 1^{\ell_i}, 1^{|E_i(\cdot)|}, \text{ABE.ct}_i)$ and

computes $\text{IBE.ct}_{i,j,b} \leftarrow \text{IBE.Enc}(\text{IBE.mpk}, (i, j, b), \text{lab}_{i,j})$, for all $j \in [\ell_i]$, and $b \in \{0, 1\}$.

- It returns $\text{ct}_x = (\mathcal{S}(|x|), f(x), \{\text{IBE.ct}_{i,j,b}\}_{i \in \mathcal{S}(|x|), j \in [\ell_i], b \in \{0,1\}})$ to \mathcal{A} .

7. \mathcal{A} outputs a guess bit β' . \mathcal{B} forwards β' to the ABE challenger.

We observe that if the ABE challenger samples $\hat{\beta} = 0$, then \mathcal{B} simulated the distribution Hyb_2 , else Hyb_3 with \mathcal{A} . Hence, advantage of $\mathcal{B} = |\Pr(\beta' = 1 | \hat{\beta} = 0) - \Pr(\beta' = 1 | \hat{\beta} = 1)| = |\Pr(\beta' = 1 | \text{Hyb}_1) - \Pr(\beta' = 1 | \text{Hyb}_2)| = \epsilon$.

Admissibility of \mathcal{B} . Observe that \mathcal{B} issues ciphertext queries for the challenge attribute of type $\{f(x)_i\}_{i \in \mathcal{S}(|x|)}$ and key queries of the form $\{g(y)_i\}_{i \in \mathcal{T}(|y|) \cap \mathcal{S}(|x|)}$. Also, by the admissibility of \mathcal{A} , it can only issue key queries $y \in \mathcal{B}$ such that $R^{\text{bndl}}(x, y) = 1$ for the challenge ciphertext attribute $x \in \mathcal{A}$. Then by the definition of R^{bndl} , we have $R_i(f(x)_i, g(y)_i) = 1$ for all $i \in \mathcal{S}(|x|) \cap \mathcal{T}(|y|)$. Thus all the key queries made by \mathcal{B} satisfies $R_i(f(x)_i, g(y)_i) = 1$. This establishes the admissibility of \mathcal{B} . ■

The rest of the hybrids, Hyb_4 and Hyb_5 , are simply unwinding the previous hybrids and their proofs of indistinguishability are same as their corresponding counterparts in the first set of hybrids and hence, omitted. ■

4.5.2 ABE for Turing Machines

Here, we provide the construction of ABE for Turing machines. Namely, a ciphertext is associated with $(x, 1^t)$ and a secret key is for a Turing machine M , and the decryption results to 1 if the machine accepts the input within t steps and 0 otherwise. To construct such a scheme, we start with constructing two schemes with partial functionality and then combine them. The one scheme takes care of the case where $|(x, 1^t)| > |M|$, while the other takes care of the case where $|(x, 1^t)| \leq |M|$. The idea for the construction is very similar to FE for TM in [16].

Circuit $U_{i,x,t}$

Hardwired constants: x, t .

On input $M = (Q, \delta, F) \in \{0, 1\}^i$, proceed as follows:

1. Parse the input $M \in \{0, 1\}^i$ as a description of a Turing machine.
2. Run M on input x for t steps.
3. Output 1 if the state is in F (accept) and 0 otherwise.

Figure 4.2: Circuit $U_{i,x,t}$.

The Case of $|(x, 1^t)| > |M|$

We first show that by applying the conversion in Section 4.5.1 to the CP-ABE scheme for unbounded depth circuits in Section 4.4, we can obtain an ABE scheme for Turing machines for the case where $|(x, 1^t)| > |M|$. Formally, we construct an ABE for $R^> : A \times B \rightarrow \{0, 1\}$, where $A = \{0, 1\}^*$, B is the set of all Turing machines, and

$$R^>((x, 1^t), M) = \begin{cases} 1 & \text{(if } M \text{ accepts } x \text{ in } t \text{ steps) } \wedge (|(x, 1^t)| > |M|) \\ 0 & \text{otherwise.} \end{cases}.$$

To apply the conversion, we use various instances of the unbounded depth and size cpABE for $\text{prm} = 1^i$, $R_i : \mathcal{X}_i \times \mathcal{Y}_i \rightarrow \{0, 1\} \cup \{\perp\}$ where \mathcal{X}_i is the set of circuits with input length i , and output length 1, $\mathcal{Y}_i = \{0, 1\}^i$, and $R_i(C, x) = C(x)$. We then set \mathcal{S} , \mathcal{T} , f , and g as

$$\mathcal{S}(i) = \{1, 2, \dots, i-1\}, \quad \mathcal{T}(i) = \{i\}, \quad f(x, 1^t) = \{U_{i,x,t}(\cdot)\}_{i \in [|x, 1^t|]-1}, \quad g(M) = M$$

where $U_{i,x,t}(\cdot)$ is defined as Figure 4.2. The circuit (in particular, Step 2 of the computation) is padded so that the circuit size only depends on $|(x, 1^t)|$. Note that $U_{i,x,t}$ is in \mathcal{X}_i even for x and t with unbounded length, since \mathcal{X}_i contains circuits of unbounded size. Then, by inspection, we can observe that for $\mathcal{X}_i, \mathcal{Y}_i, \mathcal{S}, \mathcal{T}, f, g, A, B$ defined as above, R^{bndl} defined as Equation (4.9) is equivalent to $R^>$ except for the case of $|(x, 1^t)| \leq |M|$. In this case, we have $\mathcal{S}(|x|) \cap \mathcal{T}(|y|) = \emptyset$. To handle this, we add the extra step to the

decryption algorithm of the former where it outputs \perp if $\mathcal{S}(|x|) \cap \mathcal{T}(|y|) = \emptyset$. Since the original scheme is VerSel-IND secure, so is the resulting scheme by Theorem 4.17.

The Case of $|(x, 1^t)| \leq |M|$

We next show that by applying the conversion in Section 4.5.1 to the KP-ABE scheme from [43], we can obtain an ABE scheme for TM for the case where $|(x, 1^t)| \leq |M|$. Formally, we construct an ABE for $R^\leq : A \times B \rightarrow \{0, 1\} \cup \{\perp\}$, where $A = \{0, 1\}^*$, B is the set of all non-deterministic Turing machines, and

$$R^\leq((x, 1^t), M) = \begin{cases} 1 & \text{(if } M \text{ accepts } x \text{ within } t \text{ steps) } \wedge (|(x, 1^t)| \leq |M|) \\ 0 & \text{otherwise.} \end{cases}$$

To apply the conversion, we use various instances of kpABE from [43] for $\text{prm} = 1^i$, $R_i : \mathcal{X}_i \times \mathcal{Y}_i \rightarrow \{0, 1\}$ where $\mathcal{X}_i = \{0, 1\}^i$ and \mathcal{Y}_i is the set of circuits with input length i , depth $i \cdot \lambda$, and output length 1, and $R_i(x, C) = C(x)$. We then set \mathcal{S} , \mathcal{T} , f , and g as

$$\mathcal{S}(i) = i, \quad \mathcal{T}(i) = \{1, 2, \dots, i\}, \quad f(x, 1^t) = (x, 1^t), \quad g(M) = \{U_{i,M}(\cdot)\}_{i \in [M]},$$

where $U_{i,M}(\cdot)$ is defined as Figure 4.3. Here, we check that $U_{i,M}(\cdot)$ is in \mathcal{Y}_i . Recall that even though \mathcal{Y}_i supports circuits with unbounded size, it has a bound on the depth of the circuit. We therefore argue that the depth of $U_{i,M}(\cdot)$ does not exceed $i\lambda$, even for unbounded size $|M|$. We evaluate the depth of Step 2 of the circuit, since this is the only non-trivial step. By Lemma 4.1, this step can be implemented by a circuit with depth

$$t \cdot \text{poly}(\log |x|, \log t, \log |M|) \leq i \cdot \text{poly}(\log \lambda) \leq i \cdot \lambda$$

Then, by inspection, we can observe that for $\mathcal{X}_i, \mathcal{Y}_i, \mathcal{S}, \mathcal{T}, f, g, A, B$ defined as above, R^{bndl} defined as Equation (4.9) is equivalent to R^\leq except for the case of $|(x, 1^t)| > |M|$. In this case, we have $\mathcal{S}(|x|) \cap \mathcal{T}(|y|) = \emptyset$. To handle this, we add the extra step to the decryption algorithm of the former where it outputs \perp if $\mathcal{S}(|x|) \cap \mathcal{T}(|y|) = \emptyset$. Since the original scheme is Sel-IND secure, so is the resulting scheme by Theorem 4.17.

Circuit $U_{i,M}$

Hardwired constants: Description of a Turing Machine M .

On input $y \in \{0, 1\}^i$, proceed as follows:

1. Parse the input $y \in \{0, 1\}^i$ as $(x, 1^t)$.
2. Otherwise, run M on input x for t steps.
3. Output 1 if the state is in F (accept) and 0 otherwise.

Figure 4.3: Circuit $U_{i,M}$.

Putting the Pieces Together

Here, we combine the two schemes we considered so far to obtain the full-fledged scheme.

We set $A = \{0, 1\}^*$ and B to be the set of all Turing machines. We also set $R_1 = R^>$, $R_2 = R^{\leq}$. $\mathcal{X}_i = A$, $\mathcal{Y}_i = B$ for $i = 1, 2$. We have already constructed schemes for R_1 and R_2 and now combine them. To do so, we set \mathcal{S} , \mathcal{T} , f , and g as

$$\mathcal{S}(i) = \{1, 2\}, \quad \mathcal{T}(i) = \{1, 2\}, \quad f(x, 1^t) = \{(x, 1^t), (x, 1^t)\}, \quad g(M) = \{M, M\}$$

We observe that for $\mathcal{X}_i, \mathcal{Y}_i, \mathcal{S}, \mathcal{T}, f, g, A, B$ defined as above, R^{bndl} defined as Equation (4.9) is

$$R^{\text{bndl}}((x, 1^t), M) = \begin{cases} (1, 0) & \text{(if } M \text{ accepts } x \text{ within } t \text{ steps) } \wedge (|(x, 1^t)| > |M|) \\ (0, 1) & \text{(if } M \text{ accepts } x \text{ within } t \text{ steps) } \wedge (|(x, 1^t)| \leq |M|) \\ (0, 0) & \text{otherwise.} \end{cases}$$

To obtain the ABE scheme for TM, we add an extra step for the decryption algorithm of the ABE scheme for R^{bndl} obtained above, where we output the message recovered if the decryption succeeds for either left or right slot and \perp otherwise. To sum up, we have the following theorem

Theorem 4.21. *Assume a VerSel-IND secure cpABE scheme that supports circuits with unbounded depth and a Sel-IND secure kpABE scheme that supports bounded*

depth circuits. Then there exists a ABE for Turing machines, presented in Section 4.5.2, satisfying VerSel-IND security.

Instantiating the VerSel-IND secure cpABE scheme from Section 4.4 and Sel-IND secure kpABE scheme from [43], we get the following corollary.

Corollary 4.21.1. *Under the LWE assumption, evasive LWE assumption (Assumption 2.6) for the sampler class \mathcal{SC} induced by the sampler as defined in Equation (4.7), and circular tensor LWE assumption (Assumption 4.13), there exists a very selectively secure ABE for TM with $|\text{mpk}| = \text{poly}(\lambda)$, $|\text{sk}| = \text{poly}(|M|, \lambda)$, $|\text{ct}| = \text{poly}(\lambda, |x|, t)$. where the Turing machine M runs on input x for time step t .*

4.5.3 ABE for NL

Here, we provide the construction of ABE for NL. Namely, a ciphertext encrypts a string $(x, 1^t, 1^{2^s})$, where x is a string, t is the time bound, s is the space bound for the computation and a secret key is associated with a non-deterministic Turing machine M . The decryption is possible if M accepts x within t steps and the space used for the computation does not exceed s . The idea for the construction is very similar to FE for NL in [16]. We consider two schemes that complement each other and then combine them.

Transition Matrix. Before describing the schemes, we need some preparations. Similarly to [143], we represent the computation of $M(x)$ as a multiplication of matrices. To do so, let us enumerate all the possible internal configurations that may appear when we run $M = (Q, \delta, F)$ on input x with the space for the computation being bounded by s . As an internal configuration, we have $|x|$ and s choices for the input and work tape pointers, respectively. We also have $|Q|$ choices for the possible state, and 2^s possible choices for the contents of the work tape. Therefore, we have

$$N := s2^s \cdot |x| \cdot |Q|$$

possible internal configurations. We associate each $i \in [N]$ with such configuration and represent the configuration by a vector \mathbf{e}_i , which is a unit vector whose entries are all 0 except for the i -th entry that is set to be 1. We also define the matrix $\text{Mat}(M, x, s)$ as

$$\text{Mat}(M, x, s)_{i,j} := \begin{cases} 1 & \text{if the configuration } i \text{ can reach } j \text{ in one step by } \delta \\ 0 & \text{otherwise} \end{cases}$$

where $\text{Mat}(M, x, s)_{i,j}$ is the (i, j) -th entry of the matrix.

Furthermore, we consider a special matrix multiplication over $\{0, 1\}$, where the multiplication $\mathbf{A} \cdot \mathbf{B} \in \{0, 1\}^{N_1 \times N_3}$ of two matrices $\mathbf{A} \in \{0, 1\}^{N_1 \times N_2}$ and $\mathbf{B} \in \{0, 1\}^{N_2 \times N_3}$ is defined as

$$(\mathbf{A} \cdot \mathbf{B})_{i,j} = \bigvee_{k \in [N_2]} (\mathbf{B}_{i,k} \wedge \mathbf{C}_{k,j})$$

where we denote the (i, j) -th entry of a matrix \mathbf{D} by $\mathbf{D}_{i,j}$ above. The multiplication can be defined for any size of matrices and in particular, also defined for computations involving vectors. We then define \mathbf{e}_{stt} as the vector corresponding to the initial state of the computation and \mathbf{u}_{acc} as

$$\mathbf{u}_{\text{acc}} = \sum_{i \in [N] : i \text{ encodes accepting state}} \mathbf{e}_i.$$

Then, we observe that

$$\mathbf{e}_{\text{stt}}^\top \cdot (\text{Mat}(M, x, s)_{i,j})^t \cdot \mathbf{u}_{\text{acc}} = \begin{cases} 1 & \text{if } M \text{ accepts } x \text{ within } t \text{ steps and space } s \\ 0 & \text{otherwise} \end{cases}$$

holds from the property of the transition matrix, where we use the special multiplication above.

The Case of $|(x, 1^t, 1^{2^s})| > |M|$

We first show that by applying the conversion in Section 4.5.1 to the CP-ABE scheme for circuits in [169], we can obtain an ABE scheme for NL for the case where $|(x, 1^t, 1^{2^s})| > |M|$. Formally, we construct an ABE for $R^> : A \times B \rightarrow \{0, 1\}$, where $A = \{0, 1\}^*$, B is

the set of all non-deterministic Turing machines, and

$$R^>((x, 1^t, 1^{2^s}), M) = \begin{cases} 1 & \text{(if } M \text{ accepts } x \text{ within } t \text{ steps and space } s) \wedge (|(x, 1^t, 1^{2^s})| > |M|) \\ 0 & \text{otherwise.} \end{cases}.$$

To apply the conversion, we use various instances of cpABE from [169] for $\text{prm} = 1^i$, $R_i : \mathcal{X}_i \times \mathcal{Y}_i \rightarrow \{0, 1\} \cup \{\perp\}$ where \mathcal{X}_i is the set of circuits with input length i , depth λ , and output length 1, $\mathcal{Y}_i = \{0, 1\}^i$, and $R_i(C, x) = C(x)$. We then set \mathcal{S} , \mathcal{T} , f , and g as

$$\mathcal{S}(i) = \{1, 2, \dots, i-1\}, \quad \mathcal{T}(i) = \{i\}, \quad f(x, 1^t, 1^{2^s}) = \{U_{i,x,t,s}(\cdot)\}_{i \in [|x, 1^t, 1^{2^s}|-1]}, \quad g(M) = M$$

where $U_{i,x,t,s}(\cdot)$ is defined as Figure 4.4. The circuit may be padded so that it does not leak more information about the hardwired constants (x, t, s) beyond $|(x, 1^t, 1^{2^s})|$.

We now show that f is a valid map. Namely, $U_{i,x,t,s}$ is in \mathcal{X}_i even for x and t with unbounded length. Recall that \mathcal{X}_i has a bound on the depth of the circuits it supports, even though it does not have such a bound on the size. We argue that the depth of $U_{i,x,t,s}$ is bounded by λ . To do so, we evaluate the depth of each computation step of $U_{i,x,t,s}$. First, Step 1 of the computation is trivial. Step 2 can be implemented by a circuit of depth $\text{poly}(\log |x|, \log s, \log |M|)$ by Lemma 4.2. Step 3 can be executed by $O(\log t)$ multiplication of matrices of size $N \times N$, which can be implemented with depth $O(\log N)$. Therefore, this step can be computed with depth $O(\log N \log t)$. Step 4 can also be implemented with depth $O(\log N)$. Therefore, the entire computation can be implemented with depth

$$\text{poly}(\log s, \log |x|, \log t, \log N, \log |M|) \leq \text{poly}(\log \lambda) \leq \lambda$$

as desired. Then, by inspection, we can observe that for $\mathcal{X}_i, \mathcal{Y}_i, \mathcal{S}, \mathcal{T}, f, g, A, B$ defined as above, R^{bndl} defined as Equation (4.9) is equivalent to $R^>$ except for the case of $|(x, 1^t, 1^{2^s})| \leq |M|$. In this case, we have $\mathcal{S}(|x|) \cap \mathcal{T}(|y|) = \emptyset$. To handle this, we add the extra step to the decryption algorithm of the former where it outputs \perp if $\mathcal{S}(|x|) \cap \mathcal{T}(|y|) = \emptyset$. This gives us the construction of ABE for $R^>$. Since the original

Circuit $U_{i,x,t,s}$

Hardwired constants: x, t, s .

On input $M \in \{0, 1\}^i$, proceed as follows:

1. Parse the input $M = (Q, \delta, F)$ as a description of a Turing machine.
2. Compute $\text{Mat}(M, x, s)_{j,k}$ for all $j, k \in [N]$ in parallel, where $N = s2^s \cdot |x| \cdot |Q|$.
3. Compute $\mathbf{A} := \text{Mat}(M, x, s)^t$.
4. Compute and output $\mathbf{e}_{\text{stt}}^\top \cdot \mathbf{A} \cdot \mathbf{u}_{\text{acc}}$.

Figure 4.4: Circuit $U_{i,x,t,s}$.

scheme is VerSel-IND secure, so is the resulting scheme by Theorem 4.17.

The Case of $|(x, 1^t, 1^{2^s})| \leq |M|$

We next show that by applying the conversion in Section 4.5.1 to the KP-ABE scheme from [43], we can obtain an ABE scheme for NL for the case where $|(x, 1^t, 1^{2^s})| \leq |M|$. Formally, we construct an ABE for $R^\leq : A \times B \rightarrow \{0, 1\} \cup \{\perp\}$, where $A = \{0, 1\}^*$, B is the set of all non-deterministic Turing machines, and

$$R^\leq((x, 1^t, 1^{2^s}), M) = \begin{cases} 1 & \text{(if } M \text{ accepts } x \text{ within } t \text{ steps and space } s) \wedge (|(x, 1^t, 1^{2^s})| \leq |M|) \\ 0 & \text{otherwise.} \end{cases}.$$

To apply the conversion, we use various instances of **kpABE** from [43] for $\text{prm} = 1^i$, $R_i : \mathcal{X}_i \times \mathcal{Y}_i \rightarrow \{0, 1\}$ where $\mathcal{X}_i = \{0, 1\}^i$ and \mathcal{Y}_i is the set of circuits with input length i , depth λ , and output length 1, and $R_i(x, C) = C(x)$. We then set $\mathcal{S}, \mathcal{T}, f$, and g as

$$\mathcal{S}(i) = i, \quad \mathcal{T}(i) = \{1, 2, \dots, i\}, \quad f(x, 1^t, 1^{2^s}) = (x, 1^t, 1^{2^s}), \quad g(M) = \{U_{i,M}(\cdot)\}_{i \in [M]},$$

where $U_{i,M}(\cdot)$ is defined as Figure 4.5.

We now show that g is a valid map. Namely, $U_{i,M}$ is in \mathcal{Y}_i even for M with unbounded size. In particular, we have to show that the depth of the circuit is bounded by λ .

Circuit $U_{i,M}$

Hardwired constants: Description of a Turing Machine M .

On input $y \in \{0, 1\}^i$, proceed as follows:

1. Parse the input $y \in \{0, 1\}^i$ as $(x, 1^t, 1^{2^s})$.
2. Compute $\text{Mat}(M, x, s)_{j,k}$ for all $j, k \in [N]$ in parallel, where $N = s2^s \cdot |x| \cdot |Q|$.
3. Compute $\mathbf{A} := \text{Mat}(M, x, s)^t$.
4. Compute and output $\mathbf{e}_{\text{stt}}^\top \cdot \mathbf{A} \cdot \mathbf{u}_{\text{acc}}$.

Figure 4.5: Circuit $U_{i,M}$.

This can be shown by the same argument as for $U_{i,x,t,s}$, since both circuits compute $\mathbf{e}_{\text{stt}}^\top \cdot \text{Mat}(M, x, s)^t \cdot \mathbf{u}_{\text{acc}}$ from (x, t, s, M) in exactly the same way. Then, by inspection, we can observe that for $\mathcal{X}_i, \mathcal{Y}_i, \mathcal{S}, \mathcal{T}, f, g, \mathbf{A}, \mathbf{B}$ defined as above, R^{bndl} defined as Equation (4.9) is equivalent to R^\leq except for the case of $|(x, 1^t, 1^{2^s})| > |M|$. However, this case can be handled by modifying the decryption algorithm as in Section 4.5.3. This gives us the construction of ABE for R^\leq . Since the original scheme is Sel-IND secure, so is the resulting scheme by Theorem 4.17.

Putting the Pieces Together

Here, we combine the two schemes we considered so far to obtain the full-fledged scheme. We set $\mathbf{A} = \{0, 1\}^*$ and \mathbf{B} to be the set of all Turing machines. We also set $R_1 = R^>$, $R_2 = R^\leq$. $\mathcal{X}_i = \mathbf{A}$, $\mathcal{Y}_i = \mathbf{B}$ for $i = 1, 2$. We have already constructed schemes for R_1 and R_2 and now combine them. To do so, we set $\mathcal{S}, \mathcal{T}, f$, and g as

$$\mathcal{S}(i) = \{1, 2\}, \quad \mathcal{T}(i) = \{1, 2\}, \quad f(x, 1^t) = \{(x, 1^t, 1^{2^s}), (x, 1^t, 1^{2^s})\}, \quad g(M) = \{M, M\}$$

We observe that for $\mathcal{X}_i, \mathcal{Y}_i, \mathcal{S}, \mathcal{T}, f, g, A, B$ defined as above, R^{bndl} defined as Equation (4.9) is

$$R^{\text{bndl}}((x, 1^t, 1^{2^s}), M) = \begin{cases} (1, 0) & \text{(if } M \text{ accepts } x \text{ within } t \text{ steps and space } s) \wedge (|(x, 1^t, 1^{2^s})| > |M|) \\ (0, 1) & \text{(if } M \text{ accepts } x \text{ within } t \text{ steps and space } s) \wedge (|(x, 1^t, 1^{2^s})| \leq |M|) \\ (0, 0) & \text{otherwise.} \end{cases}$$

To obtain the ABE scheme for NL, we add an extra step for the decryption algorithm of the ABE scheme for R^{bndl} obtained above, where we output the message recovered if the decryption succeeds for either left or right slot and \perp otherwise. To sum up, we have the following theorem:

Theorem 4.22. *Assume a VerSel-IND secure cpABE and Sel-IND secure kpABE scheme that supports circuits with bounded depth. Then the ABE scheme for NL presented in Section 4.5.3 satisfies VerSel-IND security.*

Instantiating the VerSel-IND secure cpABE scheme from [169] and Sel-IND secure kpABE scheme from [43], we get the following corollary.

Corollary 4.22.1. *Under the (public-coin) evasive LWE assumption (Assumption 2.6) and tensor LWE assumption (Assumption 4.13), there exists a very selectively secure ABE for NL with $|\text{mpk}| = \text{poly}(\lambda)$, $|\text{sk}| = \text{poly}(|M|, \lambda)$, $|\text{ct}| = \text{poly}(\lambda, |x|, t, 2^s)$ where the machine M runs on input x for time step t and takes space s .*

CHAPTER 5

COMPACT PSEUDORANDOM FUNCTIONAL ENCRYPTION FROM EVASIVE LWE

5.1 INTRODUCTION

Functional encryption (FE) [157, 46] generalizes public key encryption to support computation of non-trivial functions on encrypted data, beyond “all or nothing” access. More formally, in FE, a ciphertext is associated with a vector \mathbf{x} , a secret key is associated with a circuit f and decryption enables recovery of $f(\mathbf{x})$ and nothing else. Aside from its direct relevance to real world applications for computing on encrypted data, FE has proved to be a powerful tool in the theory of cryptography, and can be used to build a number of advanced primitives. The most prominent amongst these is Indistinguishability Obfuscation (iO), which is considered essentially “crypto-complete” by virtue of its ability to instantiate almost every known cryptographic primitive [34, 90, 129].

For FE to succeed in being bootstrapped all the way to iO, it must satisfy a strong efficiency property known as *compactness* – at a high level, this posits that the size of the ciphertext should be sublinear in the size of the circuits being supported by the scheme. There has been substantial research effort in the community for instantiating FE (or directly iO) from well-understood assumptions, leading to a sequence of exciting results [128, 129, 5, 20, 171, 96, 82, 129, 130, 154]. The breakthrough work of Jain, Lin and Sahai [129] finally obtained the first construction of compact FE for \mathbf{P} from standard assumptions. This has been subsequently improved by [130, 154]. However, all these works rely quite crucially on pairings which is dissatisfying. Not only are pairings quantum insecure, it is also desirable to have alternate pathways for constructing such important primitives as FE and iO, even for restricted classes of functions. In the realm of conjectured quantum safety, there exist several candidates from lattices but their

security is either based on heuristics, or their underlying assumptions have been broken [5, 20, 171, 96, 82, 120, 126]. Thus, an outstanding open question in the area is:

Can we construct compact Functional Encryption for any nontrivial functionality from simple lattice assumptions?

Attribute Based Encryption. A special case of Functional Encryption is the notion of Attribute Based Encryption (ABE) [157, 116]. ABE is similar to FE but with a crucial difference – in ABE the computation is performed on *public* attributes encoded in the ciphertext, while the burden of privacy is only on an unchanging message. In more detail, the ciphertext encodes a public *attribute* \mathbf{x} together with a secret message m , the secret key is generated for a public function f , and decryption outputs m if and only if $f(\mathbf{x}) = 1$. Security is similar to FE, except that \mathbf{x} need not be hidden. This is formalized in an indistinguishability style game which asks that an adversary should be unable to distinguish between an encryption of (m_0, \mathbf{x}) and (m_1, \mathbf{x}) , even given secret keys for functions f_i so long as $f_i(\mathbf{x}) = 0$ for all i . ABE comes in two avatars – “key-policy” where the function f is encoded in the secret key, or “ciphertext-policy” where it is encoded in the ciphertext. These are denoted by **kpABE** and **cpABE** respectively. An interesting generalization of ABE is the so-called “Predicate Encryption” (PE) [133] where the attribute \mathbf{x} is also hidden but only against an adversary that does not receive any decrypting key, namely $f_i(\mathbf{x}) = 0$ for all f_i queried by the adversary.

Prior Work in ABE. There has been significant progress in constructing ABE for circuits over the last several years [116, 107, 43, 26, 169, 122] from well-understood assumptions. Here, there has been asymmetry between the key and ciphertext policy variants – while **kpABE** can be constructed from the standard Learning With Errors (LWE) assumption, it’s **cpABE** counterpart additionally requires a relatively new, strong assumption called *Evasive LWE* [169, 163].

Recently, an important focus area in ABE research has been to achieve asymptotic optimality for both **kpABE** and **cpABE** schemes [170, 127]. The very recent work of Wee [170] constructs **kpABE** for *bounded* depth circuits using a new assumption called L -succinct **LWE**, which is weaker than evasive **LWE**. Assuming compact **FE**, [127] achieve optimality even for *unbounded* depth circuits, but this assumption necessitates the reliance on pairings as discussed above. We summarize the state of the art in Table 5.1. Thus, an outstanding open question in ABE research is:

*Can we construct **kp/cp ABE** for \mathbf{P} with optimal parameters from lattices?*

Reference	KP/CP	$ \text{mpk} $	$ \text{sk} $	$ \text{ct} $	Depth	Assumptions
[127]	KP	$\mathcal{O}(1)$	$\mathcal{O}(1)$	$\mathcal{O}(1)$	Unbdd	Compact FE for \mathbf{P}
[127]	CP	$\mathcal{O}(1)$	$\mathcal{O}(1)$	$\mathcal{O}(1)$	Unbdd	Compact FE for \mathbf{P}
[122]	KP	$\mathcal{O}(L)$	$\mathcal{O}(1)$	$\mathcal{O}(L)$	Unbdd	evasive circular LWE + circular LWE
[13]	CP	$\mathcal{O}(L)$	$\mathcal{O}(L)$	$\mathcal{O}(1)$	Unbdd	LWE + evasive LWE + circular tensor LWE
[170]	KP	$L^2 \cdot \mathcal{O}(d)$	$\mathcal{O}(d)$	$\mathcal{O}(d)$	Bdd	L -succinct LWE
[123]	CP	$\mathcal{O}(d)$	$L \cdot \mathcal{O}(d)$	$L \cdot \mathcal{O}(d)$	Bdd	LWE + evasive learning with structured errors
This Work	KP	$\mathcal{O}(1)$	$\mathcal{O}(1)$	$\mathcal{O}(1)$	Unbdd	LWE + evasive LWE
This Work	CP	$\mathcal{O}(1)$	$\mathcal{O}(1)$	$\mathcal{O}(1)$	Unbdd	LWE + evasive LWE

Table 5.1: State of the Art: Attribute Based Encryption

5.1.1 Our Results

In this work, we make progress on both the above questions and achieve the following:

(i) we provide the first construction of compact **FE** from private-coin evasive **LWE** (and **LWE**) for a nontrivial class of *pseudorandom* functionalities, (ii) we use our new **FE** to construct **kpABE** and **cpABE** achieving *optimal* parameters for *unbounded depth* circuits, also from evasive **LWE** and **LWE**. We summarize our results as informal theorems below.

The first theorem shows that we can obtain partially hiding pseudorandom **FE** (**PHprFE**) with optimal parameter sizes. In **PHprFE**, an input \mathbf{x} to the function f is divided into a public part \mathbf{x}_{pub} and private part \mathbf{x}_{priv} . While we consider the same correctness requirement as usual **FE**, we consider a relaxed security notion where the public part is allowed to leak to the adversary. The advantage of allowing part of the input to be public

is that it leads to shorter ciphertexts whose size does not depend on the length of \mathbf{x}_{pub} , when we follow the convention of ignoring the length of the public part when measuring ciphertext size.

Theorem 5.1 (Partially Hiding prFE for Unbounded Depth). *Assuming LWE and evasive LWE assumptions, there exists a partially hiding pseudorandom FE scheme, for function class $\mathcal{F} = \{f : \{0, 1\}^{L_{\text{pub}}} \times \{0, 1\}^{L_{\text{priv}}} \rightarrow \{0, 1\}\}$, that satisfies very selective security (more accurately, security as per Definition 5.20) whose master public key and the secret key sizes are fixed polynomial $\text{poly}(\lambda)$. Furthermore, the size of the ciphertext is $L_{\text{priv}} + \text{poly}(\lambda)$.*

Here, very selective security refers to the security notion where the adversary has to choose its challenge query and key queries before seeing the public parameter.

Note that in particular, this implies a plain FE scheme for pseudorandom functionalities with optimal parameters (by setting the public part of the input to \perp). We then leverage the above FE to construct Attribute Based and Predicate Encryption (ABE and PE respectively) for unbounded depth circuits with optimal parameters, in both the KP and CP setting. In more detail, we prove the following.

Theorem 5.2 (Optimal KP-ABE). *[Theorem 5.30] Under the LWE and Evasive LWE assumptions, there exists very selectively secure KP-ABE scheme supporting circuits $\{C : \{0, 1\}^\ell \rightarrow \{0, 1\}\}$ with unbounded depth and message space $\{0, 1\}^\lambda$ with*

$$|\text{mpk}| = \text{poly}(\lambda), |\text{sk}_C| = \text{poly}(\lambda), |\text{ct}| = \text{poly}(\lambda).$$

Theorem 5.3 (Optimal KP-PE). *[Theorem 5.32] Under the LWE and Evasive LWE assumptions, there exists very selectively secure KP-PE scheme supporting circuits $\{C : \{0, 1\}^\ell \rightarrow \{0, 1\}\}$ with unbounded depth and message space $\{0, 1\}^\lambda$ with*

$$|\text{mpk}| = \text{poly}(\lambda), |\text{sk}_C| = \text{poly}(\lambda), |\text{ct}| = \text{poly}(\lambda) + |\mathbf{x}|$$

where $\mathbf{x} \in \{0, 1\}^\ell$.

Theorem 5.4 (Optimal CP-ABE). *[Theorem 5.34] Under the LWE and Evasive LWE*

assumptions, there exists very selectively secure CP-ABE scheme supporting circuits with unbounded depth $\{C : \{0, 1\}^\ell \rightarrow \{0, 1\}\}$ and message space $\{0, 1\}^\lambda$ with

$$|\text{mpk}| = \text{poly}(\lambda), \quad |\text{sk}_x| = \text{poly}(\lambda), \quad |\text{ct}| = \text{poly}(\lambda)$$

where $x \in \{0, 1\}^\ell$.

We observe that since our schemes support message space $\{0, 1\}^\lambda$, they can be used to support arbitrary message space by using the hybrid encryption framework. We note that our unbounded CP-ABE scheme Theorem 5.34 and unbounded KP-ABE scheme Theorem 5.30 can be used to instantiate ABE for Turing Machines ([13]). We obtain the following corollary, which improves both the assumptions and the parameters of [13].

Corollary 5.4.1 (KP-ABE for TM). *[Corollary 5.35.1] Under the LWE and evasive LWE assumptions, there exists a very selectively secure ABE for TM with*

$$|\text{mpk}| = \text{poly}(\lambda), \quad |\text{sk}| = |M| \cdot \text{poly}(\lambda), \quad |\text{ct}| = |x| \cdot t \cdot \text{poly}(\lambda)$$

where the Turing machine M runs on input x for time step t .

In contrast, [13] uses LWE, evasive LWE and circular tensor LWE and achieves $|\text{mpk}| = \text{poly}(\lambda)$, $|\text{sk}| = \text{poly}(|M|, \lambda)$, $|\text{ct}| = \text{poly}(\lambda, |x|, t)$.

5.1.2 Additional Prior Work

The notion of iO for pseudorandom functionalities was considered implicitly by the work of Mathialagan, Peters and Vaikuntanathan [146] where they used subexponential LWE and evasive LWE to construct adaptively sound zero-knowledge SNARKs for UP. In a previous version of their work [147], which was in private circulation and shared with us, this notion was defined explicitly and leveraged to obtain unlevelled fully homomorphic encryption. However, an explicit construction of iO for pseudorandom functionalities was not provided. Moreover, FE or MIFE for any class of functionalities was not considered.

Our Companion Paper. In a companion work [14], we bootstrap our compact FE for pseudorandom functionalities to the multi-input setting. This yields the first multi-input FE for pseudorandom functionalities from LWE and (a stronger variant of) private-coin evasive LWE . Using the techniques of [29], this multi-input FE can be naturally deployed to construct the first iO for pseudorandom functionalities. We then leverage our pseudorandom multi-input FE and iO for applications – we refer the reader to [14] for details.

In the present work, our construction of ABE with optimal parameters (Section 5.5) uses a restricted iO (called pPRIO) as a building block which is developed in the companion work [14]. This creates a dependence between the two works. The rationale for our choice is primarily to have the most concise presentation and avoid duplication of content. We chose to split the papers across the axis of single input and multi-input since this seems most natural. The notion of pPRIO uses MIFE for constant arity, so fits naturally in the companion work. By *not* using it in the present work, we would achieve ABE schemes with sub-optimal parameters which nevertheless outperform the state of the art. We would then improve these schemes in the companion work, which was undesirable. Another alternative was to provide the multi-input compiler for constant arity in the present work and generalize this to polynomial arity (from a stronger assumption) in the companion work. However this would lead to duplicating the MIFE construction in both works. Therefore, we believe that using the pPRIO black-box as a tool in the ABE application developed here is the best option.

5.1.3 Recent Attacks on Evasive LWE and Repercussions.

Subsequent to the first online appearance of the present work, some counter-examples for evasive LWE were developed. We discuss these and their impact on the present work below. To begin, we recall the definition of evasive LWE .

The evasive LWE assumption [169, 163] roughly says that if

$$(\mathbf{B}, \mathbf{P}, \mathbf{s}^\top \mathbf{B} + \mathbf{e}_\mathbf{B}, \mathbf{s}^\top \mathbf{P} + \mathbf{e}_\mathbf{P}, \text{aux}) \approx_c (\mathbf{B}, \mathbf{P}, \$, \$, \text{aux})$$

where $\$$ represents random, then

$$(\mathbf{B}, \mathbf{P}, \mathbf{s}^\top \mathbf{B} + \mathbf{e}_\mathbf{B}, \mathbf{B}^{-1}(\mathbf{P}), \text{aux}) \approx_c (\mathbf{B}, \mathbf{P}, \$, \mathbf{B}^{-1}(\mathbf{P}), \text{aux})$$

Above $\mathbf{B}^{-1}(\mathbf{P})$ refers to a low norm matrix, say \mathbf{K} , such that $\mathbf{BK} = \mathbf{P} \bmod q$. It is clear that given $\mathbf{B}^{-1}(\mathbf{P})$ and $\mathbf{s}^\top \mathbf{B} + \mathbf{e}_\mathbf{B}$, the adversary can compute $\mathbf{s}^\top \mathbf{P} + \mathbf{e}_\mathbf{P}$. Evasive LWE intuitively says that the adversary cannot exploit $\mathbf{B}^{-1}(\mathbf{P})$ in any other way. The rationale behind the assumption, discussed at length in [169], is that the final value obtained by the adversary after decryption, represented above by $\mathbf{s}^\top \mathbf{P} + \mathbf{e}_\mathbf{P}$, is large and completely avoids the so-called “zeroizing” regime, namely the situation where the attacker obtains low norm polynomial equations over the integers and can solve these to obtain harmful leakage. Thus, evasive LWE should allow us to construct schemes where the decryption yields a pseudorandom output, as required by the pre-condition. Indeed, this is the motivation for the name “evasive” LWE – it should only support construction of evasive functionalities where decryption of challenge ciphertexts is not allowed, such as ABE, but not FE/iO.

Evasive LWE has been studied in two main regimes, namely “public-coin” and “private-coin”, where the former means that the randomness used to sample \mathbf{P} and auxiliary information aux , is made available to the adversary, and the latter means that this information needs to be hidden. The private-coin version was known to have some contrived counter-examples since the work of [164], but this was not considered too problematic as it relied on highly unnatural auxiliary information which contained obfuscations that would output secrets given $\mathbf{B}^{-1}(\mathbf{P})$ but not otherwise. No attacks were known in the public-coin setting used by Wee’s original formulation or its extensions, such as the circular evasive LWE by Hsieh, Lin and Luo [122]. Thus, evasive LWE has been seen as a meaningful “middle point” in the land between LWE on one hand, and

lattice assumptions used for iO on the other. The community has tried to make progress on long standing “evasive” problems in the world of lattice based cryptography using this assumption, while simultaneously trying to improve the assumption (see for instance [170]).

New Attacks. In a very exciting, very recent development subsequent to the first online posting of the present work, some surprising new counter-examples against evasive LWE became known [19, 121, 85]. In a nutshell, these attacks show that the intuition that evasive LWE evades the zeroizing regime is not always true, even in the public-coin setting. However, by taking suitable precautions, security can be recovered as discussed below.

Malicious Sampler Attacks. The works of [19, 121, 85] demonstrated that the general formulation of evasive LWE, which allows for arbitrary malicious samplers, is *false as stated*. In particular, as related to the present work, the works of [19, 121] showed (via essentially the same attack) that by carefully crafting a contrived circuit to implement a PRF which is used (non black-box) in our scheme, the pre-condition can still be argued true while the post-condition can be shown false. Additionally, besides attacking private coin versions of Evasive LWE that are prevalent in the literature, the attacks by [19] also affect some public-coin variants, including the “circular, small-secret evasive LWE” by [122].

We view these attacks as an important step forward in our understanding of evasive LWE. Note that evasive LWE should be seen as a family of assumptions parametrized by the description of the sampler and choice of error distributions, which, if invoked in full generality, is now known to be false, even in the public-coin setting. However, as discussed in [19], the original intuition by Wee [169] and Tsabary [163] about the security of evasive LWE can be recovered by taking precautions to identify and respect a

“safe zone” for evasive LWE – thus, by refining/restraining the formulation of evasive LWE , even the original construction (presented in the first online posting of this work) is secure. In addition, we modify our original construction of prFE to implement a modulus reduction step, which negates the effect of choosing contrived circuits in the construction – for this modified construction, we do not even know of any attack using malicious samplers. We also remark that in the real world, the circuit representation of functions is chosen by the key generator who is an honest party (it holds the master secret key), and we view the counter-examples emerging from such circuit choices as indicating the limits of the assumption rather than the security of the existing constructions.

Contrived Functionality Attacks. The work of [65] and [19] show that there exists a contrived “self-referential” functionality for which pseudorandom functional encryption or pseudorandom obfuscation cannot exist. As discussed in [19], this result may be seen as analogous to the impossibilities known for the random oracle model [68] or virtual black box (VBB) obfuscation [34]. In more detail, despite impossibilities known for ROM and VBB obfuscation [68, 34], the meaningfulness of ROM for practical security, and of VBB obfuscation for restricted functionalities [167, 69] is accepted widely. The pseudorandom functionalities that are useful for our applications, such as computing blind garbled circuits or FE ciphertexts, are quite natural and do not fall prey to such attacks.

Attacks by withholding information about \mathbf{B} or \mathbf{P} . The elegant work of [66] presents attacks against classes of evasive LWE such that either \mathbf{B} or \mathbf{P} are not known to the adversary. In our case, both \mathbf{B} and \mathbf{P} are known to the adversary – indeed \mathbf{P} is publicly computable using auxiliary information.

5.1.4 Technical Overview

In this section, we present the core ideas that we develop in this work. Below \underline{X} denotes a noisy version of X where the exact value of noise is not important.

KP-ABE for Unbounded Depth by HLL. The seminal work of Boneh et al. [43] developed algorithms for evaluating arithmetic functions on the ciphertext as well as the public key of an ABE scheme, which form the cornerstone of several subsequent constructions. Their core technique is as follows: given an input $\mathbf{x} \in \{0, 1\}^\ell$, and a matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times \ell m}$, one can homomorphically evaluate a circuit $f : \{0, 1\}^\ell \rightarrow \{0, 1\}$ on an “input encoding” matrix of form $\mathbf{A} - \mathbf{x} \otimes \mathbf{G}$ by multiplying on the right by a low norm matrix $\mathbf{H}_{\mathbf{A},f,\mathbf{x}}$ to obtain the term $\mathbf{A}_f - f(\mathbf{x})\mathbf{G}$. Here, \mathbf{G} is a special gadget matrix defined as follows. Let $\mathbf{g} = [1, 2, 2^2, \dots, 2^{\log q}]^\top$ and $\mathbf{G} = \mathbf{I} \otimes \mathbf{g}^\top$. In key evaluation, one can homomorphically evaluate a circuit $f : \{0, 1\}^\ell \rightarrow \{0, 1\}$ on \mathbf{A} to obtain $\mathbf{A}_f = \mathbf{A} \cdot \mathbf{H}_{\mathbf{A},f}$ for some low norm matrix $\mathbf{H}_{\mathbf{A},f}$. In ciphertext evaluation, given an attribute \mathbf{x} and corresponding attribute encoding of the form $\underline{\mathbf{s}^\top(\mathbf{A} - \mathbf{x} \otimes \mathbf{G})}$, which we refer to as BGG^+ encoding, right multiplication by $\mathbf{H}_{\mathbf{A},f,\mathbf{x}}$ yields $\underline{\mathbf{s}^\top(\mathbf{A}_f - f(\mathbf{x})\mathbf{G})}$ without substantially blowing up the noise in the encoding since $\mathbf{H}_{\mathbf{A},f,\mathbf{x}}$ is low norm. Skipping several details, since the key generator can compute $\mathbf{A}_f = \mathbf{A} \cdot \mathbf{H}_{\mathbf{A},f}$, it can provide a matching key which allows the decryptor to cancel out the masking term $\mathbf{s}^\top \mathbf{A}_f$ and proceed with decryption. We refer to $\mathbf{H}_{\mathbf{A},f}$ and $\mathbf{H}_{\mathbf{A},f,\mathbf{x}}$ as the PK and CT evaluation matrices respectively.

Providing Advice for Bootstrapping. The essential barrier in supporting circuits of unbounded depth for homomorphic computation is that the norm of the matrix $\mathbf{H}_{\mathbf{A},f,\mathbf{x}}$ grows exponentially with the depth of the circuit being computed, causing the the noise in the ciphertext encoding to blow out of control after some number of evaluations. The same barrier was encountered while constructing fully homomorphic encryption (FHE), to resolve which, Gentry proposed the beautiful “bootstrapping” idea. Roughly speaking, bootstrapping suggests performing homomorphic evaluation of the decryption circuit on

the large-noise ciphertext – this has the effect of internally throwing away the large noise accumulated in the ciphertext and “refreshing” it with smaller, more manageable noise. Since this procedure can be performed every time the ciphertext has accumulated large noise, the evaluator can keep on going! However, to execute this approach, one needs to assume that the scheme satisfies circular security, namely it should be safe to provide a ciphertext encrypting the scheme’s own secret key, as this is required for homomorphic decryption described above.

While it has been long known how to use circular security in the context of unbounded depth FHE [97], its utility in the context of ABE was uncovered only very recently, in an elegant work by Hsieh, Lin and Luo [122] (HLL). In more detail, HLL supports unbounded homomorphism in ABE via two steps: (i) noise removal, and (ii) bootstrapping. The noise removal step is via modulus reduction à la BV/BGV [62, 55] which destroys the algebraic structure of the BGG^+ encoding required for further evaluation, while the bootstrapping step makes use of a circular encoding, or “advice”, which enables to restore the structure of the BGG^+ encoding, making it suitable for further evaluation. In more detail, the HLL advice has the following structure:

$$\mathbf{S} = \text{hct}_s(\mathbf{s}), \quad \mathbf{E} = \underline{\mathbf{s}^\top (\mathbf{A}_{\text{circ}} - \mathbf{S} \otimes \mathbf{G})} \quad (5.1)$$

Above, $\text{hct}_s(\cdot)$ is an FHE ciphertext decryptable by secret key \mathbf{s} denoted in the subscript – thus \mathbf{S} is a circular FHE ciphertext, and \mathbf{E} is a BGG^+ encoding with attribute \mathbf{S} and *re-using* the FHE secret \mathbf{s} as the LWE secret! The trick of reusing the FHE secret as the LWE secret in the BGG^+ encoding of attribute $\text{hct}_s(\cdot)$, was introduced by Brakerski et al. [60] and can lead to “automatic decryption” of the FHE ciphertext, as described next. Recall that in the GSW FHE scheme [99], the secret key is \mathbf{s}^\top , a ciphertext for message \mathbf{y}^\top is a matrix \mathbf{C} and decryption computes $\mathbf{s}^\top \mathbf{C}$ to recover (a noisy version of) \mathbf{y}^\top . Brakerski et al. [60] suggested “vectorizing” the BGG^+ ciphertext evaluation procedure so that homomorphic evaluation on the encoding produces a term of the form $\underline{\mathbf{s}^\top (\mathbf{A}_f - \text{hct}_s(\mathbf{y}^\top))}$ (i.e. without \mathbf{G}). Now, the inner product of \mathbf{s} and $\text{hct}_s(\mathbf{y}^\top)$ causes

FHE decryption to occur automatically and we obtain the encoding $\underline{s^\top \mathbf{A}_f + \mathbf{y}^\top}$, where the noise in the encoding is low. It turns out that this term is exactly what is needed to restore the structure of the ill-formed encoding obtained by step (i) of HLL – this allows to create a low noise BGG^+ encoding which can be used to evaluate further. Put together, HLL prove the following:

Theorem 5.5. *Assuming circular evasive LWE and LWE , there exists a very selectively secure kpABE scheme for circuits of unbounded depth and attribute length ℓ with $|\text{mpk}| = \text{poly}(\lambda, \ell)$, $|\text{sk}_C| = \text{poly}(\lambda)$, $|\text{ct}| = \text{poly}(\lambda, \ell)$.*

Above *circular evasive LWE* is a new assumption introduced by HLL which “mixes” circularity into the recently introduced evasive LWE assumption [169, 163].

Randomizing Advice for CP-ABE. Recently, Agrawal, Kumari and Yamada [13] (AKY) built upon the construction by HLL to obtain the first ABE for Turing machines from lattice assumptions. A key technical contribution of the AKY construction is a way to randomize the advice provided in the HLL ciphertext, making it suitable for integration with Wee’s bounded depth CP-ABE. These techniques led to the first CP-ABE for unbounded depth circuits, which they further leveraged to construct a (KP-)ABE for Turing machines. In more detail, the AKY transformation requires computation of the following randomized HLL terms:

$$\mathbf{S}_r = \text{hct}_{s_r}(\mathbf{s}_r), \quad \mathbf{E}_r = \underline{s_r^\top (\mathbf{A}_{\text{circ}} - \mathbf{S}_r \otimes \mathbf{G})} \quad (5.2)$$

Above, $\mathbf{s}_r = \mathbf{s}^\top (\mathbf{I} \otimes \mathbf{r})$ where \mathbf{r} is chosen by the key generator while \mathbf{s} is chosen by the encryptor.

Evidently, neither party can provide the encodings directly, and while randomizing the message inside an FHE ciphertext from \mathbf{s} to \mathbf{s}_r is easy given knowledge of \mathbf{r} , randomizing the secret key of FHE ciphertext is much more challenging. To get around this difficulty,

AKY suggest that the structure of the advice provided by the encryptor be changed, so that the true power of FHE – which is to transform encoded messages rather than underlying secret keys – be further leveraged. Thus, they provide:

$$\mathbf{T} = \text{hct}_{\mathbf{t}}(\mathbf{s}, \mathbf{sd}), \quad \mathbf{D} = \underline{\mathbf{t}^{\top}(\mathbf{A}_1 - (1, \text{bits}(\mathbf{T})) \otimes \mathbf{G})}$$

where \mathbf{sd} is a PRF seed, \mathbf{t} is the secret of a fresh FHE scheme and \mathbf{A}_1 is a public matrix of appropriate dimensions. Now, one can homomorphically evaluate on the encoding \mathbf{D} in *bounded* depth, using knowledge of \mathbf{T} , to obtain

$$\underline{\mathbf{t}^{\top} \mathbf{A}'_{\mathbf{r}} + \mathbf{t}^{\top} \text{hct}_{\mathbf{t}}(\mathbf{S}_{\mathbf{r}}, \mathbf{E}_{\mathbf{r}})} = \underline{\mathbf{t}^{\top} \mathbf{A}'_{\mathbf{r}} + (\mathbf{S}_{\mathbf{r}}, \mathbf{E}_{\mathbf{r}})}$$

where $\mathbf{A}'_{\mathbf{r}}$ is some \mathbf{r} dependent matrix and the equality follows by automatic decryption. To get rid of the masking term $\mathbf{t}^{\top} \mathbf{A}'_{\mathbf{r}}$, the encryptor additionally provides $\underline{\mathbf{t}^{\top} \mathbf{C}}$ for some fixed matrix \mathbf{C} and the key generator provides $\mathbf{C}^{-1}(\mathbf{A}'_{\mathbf{r}})$ where $\mathbf{C}^{-1}(\mathbf{A}'_{\mathbf{r}})$ is not a true matrix inverse but rather a low norm matrix so that $\mathbf{C} \cdot \mathbf{C}^{-1}(\mathbf{A}'_{\mathbf{r}}) = \mathbf{A}'_{\mathbf{r}}$. Together these allow the decryptor to compute the term $\underline{\mathbf{t}^{\top} \mathbf{A}'_{\mathbf{r}}}$ and cancel it out from the encoding above, to recover $(\mathbf{S}_{\mathbf{r}}, \mathbf{E}_{\mathbf{r}})$ in the clear.

The security of the above construction, relies on evasive LWE (aside from other assumptions), and depends crucially on the fact that the computed terms $(\mathbf{S}_{\mathbf{r}}, \mathbf{E}_{\mathbf{r}})$ are pseudorandom. Digging deeper into the AKY proof, the term $\underline{\mathbf{s}^{\top} \mathbf{P}}$ in Evasive LWE can be essentially simplified to the advice terms $(\mathbf{S}_{\mathbf{r}}, \mathbf{E}_{\mathbf{r}})$ which, therefore need to be pseudorandom for invoking the assumption. Put together, AKY show the following:

Theorem 5.6. [13][Thm 5.6] *Under LWE, circular tensor LWE and evasive LWE, there exists a very selectively secure ABE for TM with $|\text{mpk}| = \text{poly}(\lambda)$, $|\text{sk}| = \text{poly}(\lambda, |M|)$, $|\text{ct}| = \text{poly}(\lambda, |\mathbf{x}|, t)$.*

Above M is the Turing machine, \mathbf{x} is the input and t is the worst case running time of M on any input. Note that while the AKY construction does not need the circular evasive LWE assumption used by HLL, they require a circular *tensor* assumption which is a new

assumption that they introduce.

Compact Functional Encryption for Pseudorandom Functionalities. Our starting point is the observation that the techniques developed in AKY are quite a bit more general and can be leveraged to compute functionalities beyond the randomized HLL advice they were developed for. Taking a step back, let us analyze what their technique enables: the encryptor provides an FHE ciphertext \mathbf{T} of a message (say \mathbf{x}), and a BGG^+ encoding of attribute \mathbf{T} with the FHE secret doubling up as the encoding randomness. Homomorphic evaluation of any function f coupled with automatic decryption allows to recover a masked version of $f(\mathbf{x})$ and evasive LWE allows to cancel the mask. Thus, this technique seems to enable computation of any function f on the input \mathbf{x} , while keeping it hidden! Intuitively, security follows from evasive LWE as long as the output of the functionality is pseudorandom, such as their $(\mathbf{S}_r, \mathbf{E}_r)$, but more generally the output of any pseudorandom function, such as a PRF.

We show that the above intuition can be formalized to yield the first compact FE for pseudorandom functionalities, namely, functionalities where the output is (pseudo)random for any given input that is seen by the adversary in the security game. We sketch our construction for prFE below. In the following, $f : \{0, 1\}^L \rightarrow \{0, 1\}^\ell$ has the property that the output of f is pseudorandom for every input seen by the adversary. We also use a $\text{PRF} : \{0, 1\}^\lambda \times \{0, 1\}^\lambda \rightarrow [-q/4, q/4]$. The usage of PRF is introduced for the security reasons which we will highlight later.

- The setup algorithm samples matrices \mathbf{A}_{att} and $(\mathbf{B}, \mathbf{B}^{-1})$ of appropriate dimensions and outputs $\text{mpk} := (\mathbf{A}_{\text{att}}, \mathbf{B})$ and $\text{msk} := \mathbf{B}^{-1}$. Here, \mathbf{B}^{-1} is the trapdoor for \mathbf{B} which allows to compute short preimages $\mathbf{B}^{-1}(\mathbf{U})$ for any target matrix \mathbf{U} .
- The encryptor on input \mathbf{x} first samples a GSW secret key \mathbf{s} and a PRF seed $\text{sd} \leftarrow \{0, 1\}^\lambda$. It then computes a GSW ciphertext, $\mathbf{X} = \text{hct}_{\mathbf{s}}(\mathbf{x}, \text{sd})$, using public key $\mathbf{A}_{\text{fhe}} = (\bar{\mathbf{A}}_{\text{fhe}} \quad \bar{\mathbf{s}}^\top \bar{\mathbf{A}}_{\text{fhe}} + \mathbf{e}_{\text{fhe}}^\top)$ and randomness \mathbf{R} , – followed by a BGG^+ encoding of \mathbf{X} using randomness \mathbf{s} as $\mathbf{c}_{\text{att}}^\top := \mathbf{s}^\top (\mathbf{A}_{\text{att}} - \mathbf{X} \otimes \mathbf{G}) + \mathbf{e}_{\text{att}}^\top$. It additionally computes $\mathbf{c}_{\mathbf{B}}^\top := \mathbf{s}^\top \mathbf{B} + \mathbf{e}_{\mathbf{B}}^\top$ and outputs the ciphertext $\text{ct} = (\mathbf{c}_{\mathbf{B}}, \mathbf{c}_{\text{att}}, \mathbf{X})$.

- The key generator on input $\text{msk} = \mathbf{B}^{-1}$ and function f does the following.
 - (a) Samples a nonce $\mathbf{r} \leftarrow \{0, 1\}^\lambda$ and defines function $F[f, \mathbf{r}]$, with f and \mathbf{r} hardwired, as

$$F[f, \mathbf{r}](\mathbf{x}, \text{sd}) = f(\mathbf{x}) \lfloor q/2 \rfloor + \text{PRF}(\text{sd}, \mathbf{r}).$$

It then computes the FHE evaluation circuit VEval_F w.r.t. the function $F[f, \mathbf{r}]$ (this can be computed using the knowledge of $F[f, \mathbf{r}]$). Note that the circuit VEval_F can be used to compute on a GSW ciphertext encoding an input, say \mathbf{y} , to recover a GSW ciphertext encoding $F[f, \mathbf{r}](\mathbf{y})$.

- (b) Next, it computes the matrix $\mathbf{H}_{\mathbf{A}_{\text{att}}}^F$ for the circuit VEval_F using the public matrix \mathbf{A}_{att} . Recall that the matrix $\mathbf{H}_{\mathbf{A}_{\text{att}}}^F$ and $\mathbf{H}_{\mathbf{A}_{\text{att}}, \mathbf{X}}^F$ (which can be computed given VEval_F , \mathbf{A}_{att} and \mathbf{X}) will satisfy the relation

$$(\mathbf{A}_{\text{att}} - \mathbf{X} \otimes \mathbf{G})\mathbf{H}_{\mathbf{A}_{\text{att}}, \mathbf{X}}^F = \mathbf{A}_{\text{att}}\mathbf{H}_{\mathbf{A}_{\text{att}}}^F - \text{VEval}_F(\mathbf{X}).$$

- (c) It sets $\mathbf{A}_F = \mathbf{A}_{\text{att}} \cdot \mathbf{H}_{\mathbf{A}_{\text{att}}}^F$, samples $\mathbf{K} \leftarrow \mathbf{B}^{-1}(\mathbf{A}_F)$ and outputs $\text{sk}_f = (\mathbf{K}, \mathbf{r})$.

- The decryption on input $\text{sk}_f = (\mathbf{K}, \mathbf{r})$ and $\text{ct} = (\mathbf{c}_B, \mathbf{c}_{\text{att}}, \mathbf{X})$ work as follows.

- (a) It first computes the matrix $\mathbf{H}_{\mathbf{A}_{\text{att}}, \mathbf{X}}^F$ for the circuit VEval_F using \mathbf{A}_{att} and \mathbf{X} .
- (b) Next, it computes $\mathbf{z} := \mathbf{c}_B^\top \cdot \mathbf{K} - \mathbf{c}_{\text{att}}^\top \cdot \mathbf{H}_{\mathbf{A}_{\text{att}}, \mathbf{X}}^F$, rounds \mathbf{z} co-ordinate wise and output the most significant bits.

To see the correctness of our scheme, we note that

$$\mathbf{c}_{\text{att}}^\top \cdot \mathbf{H}_{\mathbf{A}_{\text{att}}, \mathbf{X}}^F \approx \mathbf{s}^\top \mathbf{A}_{\text{att}} \mathbf{H}_{\mathbf{A}_{\text{att}}}^F - \mathbf{s}^\top (\text{hct}_s(F(\mathbf{x}, \text{sd}))) \approx \mathbf{s}^\top \mathbf{A}_F - F(\mathbf{x}, \text{sd}),$$

where the second approximate equality follows by automatic decryption. Now to remove the masking term " $\mathbf{s}^\top \mathbf{A}_F$ " we compute $\mathbf{c}_B^\top \cdot \mathbf{K} \approx \mathbf{s}^\top \mathbf{A}_F$ and thus $\mathbf{z} \approx \mathbf{s}^\top \mathbf{A}_F - \mathbf{s}^\top \mathbf{A}_F + F(\mathbf{x}, \text{sd}) = f(\mathbf{x}) \lfloor q/2 \rfloor + \text{PRF}(\text{sd}, \mathbf{r})$. Now, rounding gives us bits of $f(\mathbf{x})$ as long as $|\text{PRF}(\text{sd}, \mathbf{r})| \leq q/4$. The exact decryption error is

$$\mathbf{e}_B^\top \mathbf{K} + \text{PRF}(\text{sd}, \mathbf{r}) - (\mathbf{e}_{\text{fhe}}^\top \mathbf{R}_F + \mathbf{e}_{\text{att}}^\top \mathbf{H}_{\mathbf{A}_{\text{att}}, \mathbf{X}}^F)$$

where $\text{VEval}_F(\text{bits}(\mathbf{X})) = \mathbf{A}_{\text{fhe}} \mathbf{R}_F - (\mathbf{0} \ F[f, \mathbf{r}](\mathbf{x}, \text{sd}))^\top$.

Our construction supports functions of bounded polynomial depth $\text{dep} = \text{poly}(\lambda)$ and

has the following efficiency

$$|\text{mpk}| = L \cdot \text{poly}(\text{dep}, \lambda), \quad |\text{sk}_f| = \ell \cdot \text{poly}(\text{dep}, \lambda), \quad |\text{ct}| = L \cdot \text{poly}(\text{dep}, \lambda).$$

As seen above, our construction achieves compactness.

Security. Intuitively, our notion of security says that so long as the *output* of the functionality is pseudorandom, the *ciphertext* is pseudorandom, given all the additional information available to the adversary. We denote this security notion as prCT security. In more detail, let Samp be a PPT algorithm that on input 1^λ , outputs

$$(f_1, \dots, f_{Q_{\text{key}}}, x_1, \dots, x_{Q_{\text{msg}}}, \text{aux} \in \{0, 1\}^*)$$

where Q_{key} is the number of key queries, Q_{msg} is the number of message queries. We say that a prFE scheme is secure if

$$\left(\text{mpk}, \text{aux}, f_1, \dots, f_{Q_{\text{key}}}, \left\{ \text{Enc}(\text{mpk}, x_j) \right\}_{j \in [Q_{\text{msg}}]}, \text{sk}_{f_1}, \dots, \text{sk}_{f_{Q_{\text{key}}}} \right) \approx_c \left(\text{mpk}, \text{aux}, f_1, \dots, f_{Q_{\text{key}}}, \left\{ \delta_j \leftarrow \mathcal{CT} \right\}_{j \in [Q_{\text{msg}}]}, \text{sk}_{f_1}, \dots, \text{sk}_{f_{Q_{\text{key}}}} \right) \quad (5.3)$$

$$\text{given } \left(\text{aux}, f_1, \dots, f_{Q_{\text{key}}}, \{f_i(x_j)\}_{i \in [Q_{\text{key}}], j \in [Q_{\text{msg}}]} \right) \approx_c \left(\text{aux}, f_1, \dots, f_{Q_{\text{key}}}, \{\Delta_{i,j}\}_{i \in [Q_{\text{key}}], j \in [Q_{\text{msg}}]} \right) \quad (5.4)$$

where $(\text{mpk}, \text{msk}) \leftarrow \text{Setup}(1^\lambda)$, $\text{sk}_{f_i} \leftarrow \text{KeyGen}(\text{msk}, f_i)$ for $i \in [Q_{\text{key}}]$, \mathcal{CT} is the ciphertext space and $\Delta_{i,j} \leftarrow \{0, 1\}^\ell$ for $i \in [Q_{\text{key}}], j \in [Q_{\text{msg}}]$.

The careful reader may have noticed that the above definition has a multi-challenge flavour, even though the construction is in the public-key setting. This peculiarity arises because single-challenge security does not generically imply multi-challenge security for our definition. To see this, recall the standard hybrid argument to prove multi-challenge security from single-challenge security: the proof follows a sequence of hybrids, where we simulate some of the ciphertexts honestly, while trying to change

a particular honest ciphertext to be random. Now, to generate honest ciphertexts, we need to know the corresponding plaintexts. However, this could ruin the precondition for invoking single-challenge security for the target ciphertext, since knowing some inputs may ruin the pseudorandomness of outputs, if the inputs are correlated to each other.

We observe that the above definition is incomparable to the standard indistinguishability (IND) security definition that is used in FE schemes [46]. Indeed, IND style security does not appear meaningful for our functionality. Recall that in the IND game, the adversary must submit two challenge messages m_0 and m_1 and request keys for functions f_i such that $f_i(m_0) = f_i(m_1)$. However since the output $f_i(m_b)$ is pseudorandom for $b \in \{0, 1\}$, the event $f_i(m_0) = f_i(m_1)$ occurs only with negligible probability¹. However, our definition has a simulation security flavour and is very handy for applications, as we will see.

We provide some high level intuition for our proof of security for prFE. For simplicity, we focus here on the single challenge setting and refer the reader to the main body for the detailed proof in the multi-challenge setting. The proof begins by invoking evasive LWE with an appropriate sampler – this allows to reduce the reasoning to the distribution of the pre-condition, which replaces the term $\mathbf{K} = \mathbf{B}^{-1}(\mathbf{A}_F)$ with $\mathbf{c}_B^T \cdot \mathbf{K} = \underline{\mathbf{s}^T \mathbf{A}_F}$. Now, as we see in Equation (5.1.4),

$$\mathbf{c}_{\text{att}}^T \cdot \mathbf{H}_{\mathbf{A}_{\text{att}}, \mathbf{X}}^F = \underline{\mathbf{s}^T \mathbf{A}_F - F(\mathbf{x}, \text{sd})} = \underline{\mathbf{s}^T \mathbf{A}_F - f(\mathbf{x}) \lfloor q/2 \rfloor - \text{PRF}(\text{sd}, \mathbf{r})}.$$

This allows to simplify these two terms to $\underline{\mathbf{s}^T \mathbf{A}_F}$ and $\underline{f(\mathbf{x}) \lfloor q/2 \rfloor + \text{PRF}(\text{sd}, \mathbf{r})}$, where the latter term is pseudorandom, hence simulatable and can be ignored hereafter. The terms that remain can now be handled by relying on LWE using standard techniques [169, 122, 13]. Please see Section 5.3 for the detailed proof.

¹A generalization of IND security which requires $f_i(m_b)$ to be pseudorandom and hence computationally indistinguishable for any b might have been more suitable.

Handling Malicious Samplers. As discussed above, subsequent to the initial posting of this work on eprint, [19, 121, 85] demonstrated that the general formulation of evasive LWE is false as stated. As suggested in [19], there are multiple ways to restrain malicious samplers to prevent such attacks. The simplest one is to leverage the fact that the secret key is computed by the key generator who is an honest party (it holds the master secret key) in the real world, and can ensure that the circuit representation of any function f as well as the PRF can be made canonical by using the universal circuit or a garbled circuit representation. Nevertheless, here, we present an alternate fix to the scheme which uses modulus reduction to “throw away” the accumulated error after FHE evaluation, replacing it with rounding error which is no longer correlated with the PRF seed, even for a contrived circuit chosen by the adversary. To begin, we provide a high level outline of the attack.

Attack by [19, 121] The attacks proposed by [19, 121] on the (private-coin) Evasive LWE (Assumption 2.6) for *any* sampler breaks our initial construction. At a very high level, the attack, building upon clever ideas by [120], shows a way to create a correlation between the error term resulting from FHE evaluation (and automatic decryption) with the PRF output by using a contrived circuit to implement the PRF. In more detail, the adversary, given \mathbf{c}_B , \mathbf{c}_{att} , \mathbf{X} , and \mathbf{K} , computes $\mathbf{c}_B^\top \cdot \mathbf{K} - \mathbf{c}_{att}^\top \cdot \mathbf{H}_{A_{att}, \mathbf{X}}^F$. Simplifying, she obtains

$$f(\mathbf{x}) \lfloor q/2 \rfloor + \mathbf{e}_B^\top \mathbf{K} + \text{PRF}(\text{sd}, \mathbf{r}) - (\mathbf{e}_{fhe}^\top \mathbf{R}_F + \mathbf{e}_{att}^\top \mathbf{H}_{A_{att}, \mathbf{X}}^F)$$

By correctness, the adversary recovers $f(\mathbf{x})$ and can therefore strip it away to obtain $\text{PRF}(\text{sd}, \mathbf{r}) + \mathbf{e}_B^\top \mathbf{K} - \mathbf{e}_{fhe}^\top \mathbf{R}_F - \mathbf{e}_{att}^\top \mathbf{H}_{A_{att}, \mathbf{X}}^F$, as described in Section 5.1.4. Now, in the proof, the error term $\mathbf{e}_B^\top \mathbf{K}$ is replaced by i.i.d error \mathbf{e}_P and is used to break any correlation between $\text{PRF}(\text{sd}, \mathbf{r})$ and $\mathbf{e}_{fhe}^\top \mathbf{R}_F - \mathbf{e}_{att}^\top \mathbf{H}_{A_{att}, \mathbf{X}}^F$. This allows us to prove the pre-condition based on just plain LWE.

However, in the real world, $\mathbf{e}_B^\top \mathbf{K}$ cannot be used to break the dependence between the

above two terms. Then, by choosing the circuit implementing F in some contrived way, the authors set it up so that $\text{PRF}(\text{sd}, \mathbf{r})$ and $\mathbf{e}_{\text{fhe}}^\top \mathbf{R}_F$ are correlated, and in particular cancel each other modulo 2. Note that these terms are small, and do not wraparound modulo q , so computing mod 2 is well defined. Now we are left with $\mathbf{e}_B^\top \mathbf{K} + \mathbf{e}_{\text{att}}^\top \mathbf{H}_{\text{Att}, X}^F$ – but these are linear equations with known coefficients, in the error terms of the original encodings. Using sufficiently many equations, the adversary can easily recover the error terms. On the other hand, had the term $\mathbf{e}_B^\top \mathbf{K}$ been truly random, such a system of equations would not admit any solution. This leads to a distinguishing strategy.

Fixing the Scheme. We describe an approach that helps us break the problematic correlation even if the circuit implementation is chosen in a contrived manner – our idea is to use modulus reduction get rid of the problematic error terms involving $\mathbf{e}_{\text{fhe}}^\top \mathbf{R}_F$ so that the correlation is destroyed. Informally, we fix a rounding constant $M \in \mathbb{Z}$ such that $\|\mathbf{e}_{\text{fhe}}^\top \mathbf{R}_F\| < M$ which implies $\lfloor (\mathbf{e}_{\text{fhe}}^\top \mathbf{R}_F)/M \rfloor = 0$. This gets rid of the problematic error, replacing it with rounding error which is uncorrelated with the PRF seed. For concreteness, we elaborate the changes that must be made to our prFE construction to incorporate the above fix.

1. In setup algorithm, we output M as a part of mpk . The encryption algorithm remains the same.
2. The key generation algorithm has the following changes.
 - We parse $F[f, \mathbf{r}](\mathbf{x}, \text{sd}) = f(\mathbf{x}) \lfloor q/2 \rfloor + \text{PRF}(\text{sd}, \mathbf{r}) = M \cdot f_{\text{high}}(\mathbf{x}, \text{sd}) + f_{\text{low}}(\mathbf{x}, \text{sd})$, where $f_{\text{high}}(\mathbf{x}, \text{sd}) \in [0, q/M]^\ell$ and $f_{\text{low}}(\mathbf{x}, \text{sd}) \in [0, M-1]^\ell$. Next we define functions $F_{\text{high}} := M \cdot f_{\text{high}}$ and $F_{\text{low}} := M \cdot f_{\text{low}}$, which on input (\mathbf{x}, sd) outputs $M \cdot f_{\text{high}}(\mathbf{x}, \text{sd})$ and $M \cdot f_{\text{low}}(\mathbf{x}, \text{sd})$, respectively.
 - Next, it computes circuits $\text{VEval}_{\text{high}}$ and $\text{VEval}_{\text{low}}$ for the functions F_{high} and F_{low} , respectively and then uses these circuits to compute the matrices $\mathbf{H}_{\text{Att}}^{F_{\text{high}}}$ and $\mathbf{H}_{\text{Att}}^{F_{\text{low}}}$ (as described in the previous sketch).
 - It sets

$$\mathbf{A}_F = M \cdot \left\lfloor \frac{\mathbf{A}_{\text{att}} \cdot \mathbf{H}_{\text{Att}}^{F_{\text{high}}}}{M} \right\rfloor + \left\lfloor \frac{\mathbf{A}_{\text{att}} \cdot \mathbf{H}_{\text{Att}}^{F_{\text{low}}}}{M} \right\rfloor$$

and outputs $\mathbf{sk}_f = (\mathbf{K}, \mathbf{r})$ where $\mathbf{K} = \mathbf{B}^{-1}(\mathbf{A}_F)$.

3. The decryption algorithm is the same except that we compute \mathbf{z} differently as

$$\mathbf{z} := \mathbf{c}_B^\top \cdot \mathbf{K} - \left(M \cdot \left\lfloor \frac{\mathbf{c}_{\text{att}}^\top \cdot \mathbf{H}_{\mathbf{A}_{\text{att}}, \mathbf{X}}^{\text{Fhigh}}}{M} \right\rfloor + \left\lfloor \frac{\mathbf{c}_{\text{att}}^\top \cdot \mathbf{H}_{\mathbf{A}_{\text{att}}, \mathbf{X}}^{\text{Flow}}}{M} \right\rfloor \right)$$

We expand the correctness of the scheme to see how the above changes helps us get rid of the problematic noise terms while decryption. Observe that

$$\begin{aligned} \mathbf{c}_{\text{att}}^\top \cdot \mathbf{H}_{\mathbf{A}_{\text{att}}, \mathbf{X}}^{\text{Fhigh}} &= (\mathbf{s}^\top (\mathbf{A}_{\text{att}} - \text{bits}(1, \mathbf{X}) \otimes \mathbf{G}) + \mathbf{e}_{\text{att}}^\top) \mathbf{H}_{\mathbf{A}_{\text{att}}, \mathbf{X}}^{\text{Fhigh}} \\ &= \mathbf{s}^\top \mathbf{A}_{\text{high}} - \mathbf{F}_{\text{high}}(\mathbf{x}, \text{sd}) + \mathbf{e}_{\text{fhe}}^\top \mathbf{R}_{\text{high}} + \mathbf{e}_{\text{att}}^\top \mathbf{H}_{\mathbf{A}_{\text{att}}, \mathbf{X}}^{\text{Fhigh}} \\ \Rightarrow \left\lfloor \frac{\mathbf{c}_{\text{att}}^\top \cdot \mathbf{H}_{\mathbf{A}_{\text{att}}, \mathbf{X}}^{\text{Fhigh}}}{M} \right\rfloor &= \left\lfloor \frac{\mathbf{s}^\top \mathbf{A}_{\text{high}} - M \cdot f_{\text{high}}(\mathbf{x}, \text{sd}) + \mathbf{e}_{\text{fhe}}^\top \mathbf{R}_{\text{high}} + \mathbf{e}_{\text{att}}^\top \mathbf{H}_{\mathbf{A}_{\text{att}}, \mathbf{X}}^{\text{Fhigh}}}{M} \right\rfloor \\ &= \left\lfloor \frac{\mathbf{s}^\top \mathbf{A}_{\text{high}} + \mathbf{e}_{\text{fhe}}^\top \mathbf{R}_{\text{high}} + \mathbf{e}_{\text{att}}^\top \mathbf{H}_{\mathbf{A}_{\text{att}}, \mathbf{X}}^{\text{Fhigh}}}{M} \right\rfloor - f_{\text{high}}(\mathbf{x}, \text{sd}) \\ &= \mathbf{s}^\top \left\lfloor \frac{\mathbf{A}_{\text{high}}}{M} \right\rfloor - f_{\text{high}}(\mathbf{x}, \text{sd}) \text{ w.h.p.} \end{aligned}$$

where we set $\|\mathbf{e}_{\text{fhe}}^\top \mathbf{R}_F + \mathbf{e}_{\text{att}}^\top \mathbf{H}_{\mathbf{A}_{\text{att}}, \mathbf{X}}^{\text{Fhigh}}\| \ll M$ such that the last equation in the above will hold with high probability. Similarly, we get $\left\lfloor \frac{\mathbf{c}_{\text{att}}^\top \cdot \mathbf{H}_{\mathbf{A}_{\text{att}}, \mathbf{X}}^{\text{Flow}}}{M} \right\rfloor = \mathbf{s}^\top \left\lfloor \frac{\mathbf{A}_{\text{low}}}{M} \right\rfloor - f_{\text{low}}(\mathbf{x}, \text{sd})$ w.h.p. The rest of correctness follows from the same argument as in the previous sketch. Note that the final error obtained now is $\text{PRF}(\text{sd}) + \mathbf{err}$ where (please see Equation (5.17))

$$\begin{aligned} \mathbf{err} &= M \cdot \mathbf{e}_{\text{s,high}}^\top + \mathbf{e}_{\text{s,low}}^\top + M \cdot \mathbf{err}_{\text{high}} + \mathbf{err}_{\text{low}} \\ &= M \cdot \left(\mathbf{s}^\top \left\lfloor \frac{\mathbf{A}_{\text{high}}}{M} \right\rfloor - \left\lfloor \mathbf{s}^\top \frac{\mathbf{A}_{\text{high}}}{M} \right\rfloor \right) + \left(\mathbf{s}^\top \left\lfloor \frac{\mathbf{A}_{\text{low}}}{M} \right\rfloor - \left\lfloor \mathbf{s}^\top \frac{\mathbf{A}_{\text{low}}}{M} \right\rfloor \right) + M \cdot \mathbf{err}_{\text{high}} + \mathbf{err}_{\text{low}} \end{aligned}$$

where $\mathbf{err}_{\text{high}}, \mathbf{err}_{\text{low}} \in \{0, 1\}^\ell$ are rounding errors and matrices $\mathbf{A}_{\text{high}}, \mathbf{A}_{\text{low}}$ are publicly computable matrices. We conjecture flooding for appropriately defined sizes.

Relation to Circular Small-Secret Evasive LWE of HLL23. In Section 5.3.4, we show that our prFE construction can be based on a variant of the circular small-secret evasive LWE of [122] – the only difference between the assumption stated in [122] and the one

used in Section 5.3.4 is that the FHE encoding in [122] only encodes the secret \mathbf{s} , namely $\mathbf{S} = \text{hct}_{\mathbf{s}}(\mathbf{s})$, whereas in our case, this must additionally include \mathbf{x} , i.e. $\mathbf{S} = \text{hct}_{\mathbf{s}}(\mathbf{s}, \mathbf{x})$. This difference is created by the distinction between ABE and FE, since the attribute \mathbf{x} can be public in the former and must be hidden in the latter. Hence, in [122], \mathbf{x} can be output by the sampler directly, whereas in our case, it cannot.

Whether this fundamentally changes the assumption is a matter of opinion – in any event the [122] version of circular small-secret evasive LWE is also broken in [19] by a very similar attack as on the private coin evasive assumption – but we find the connection interesting since the assumption of [122] is considered public coin. We show in Section 5.3.4 that similar refinement of this assumption can also avoid known attacks.

Applications. While it is exciting to have a compact FE scheme for any nontrivial functionality from purely (conjectured) post-quantum assumptions, we show that our prFE is also surprisingly powerful, and yields important applications.

Removing Circularity from HLL. We demonstrate the utility of prFE by showing that it can be used to bootstrap a very weak kpABE scheme into a full fledged one. This enables us to improve the assumptions underlying prior works as discussed above.

In more detail, our weak kpABE scheme, denoted by $\mathbf{1ABE}$ is a *secret key* scheme which only supports a *single* ciphertext and *single* secret key query – this object is so simple that it can be constructed merely from one way functions. This is lifted using prFE to build a full fledged *public key* ABE scheme supporting *unbounded* ciphertexts and *unbounded* key queries. Our compiler does require $\mathbf{1ABE}$ to satisfy some structural properties:

1. **Decomposability:** The computation $\mathbf{1ABE}.\text{KeyGen}(C)$ can be decomposed into $\{\mathbf{1ABE}.\text{KeyGen}_i(C_i)\}_{i \in [C]}$ where C_i denotes the i -th gate of C and has fixed polynomial size. Here, the depth of $\mathbf{ABE}.\text{KeyGen}_i$ is fixed and independent of the parameters of C . Moreover, output of $\mathbf{1ABE}.\text{Enc}$ should be computable by a circuit of fixed depth, irrespective of the length of \mathbf{x} input to $\mathbf{1ABE}.\text{Enc}$.

2. Blindness: The 1ABE ciphertext and secret key should be pseudorandom when decryption is not allowed.

Given the above properties, the core idea is to use prFE to generate randomized versions of bits of 1ABE secret keys and ciphertexts using randomness generated jointly by the encryptor and the key generator. This is supported by prFE of fixed depth because of property 1 and respects the constraints imposed by prFE because of property 2. Thus we obtain a public key scheme which supports unbounded ciphertexts and keys. We outline our compiler below.

- The setup algorithm generates (prFE.msk, prFE.mpk) using prFE.Setup and outputs these as msk and mpk respectively.
- The encryption algorithm on input mpk, attribute \mathbf{x} and message μ computes a prFE ciphertext, prFE.ct, encoding input $(\mathbf{x}, \mu, \text{sd})$ where $\text{sd} \leftarrow \{0, 1\}^\lambda$ is a PRF seed.
- The keygen algorithm on input msk = prFE.msk and circuit C works as follows. It samples nonce $\mathbf{r} \leftarrow \{0, 1\}^\lambda$ and defines functions $F_{\text{key},i}[\mathbf{r}, C_i]$, with \mathbf{r} and i -th gate of C hardwired, for $i \in [|C|]$ and $F_{\text{ct}}[\mathbf{r}]$ as follows
 - (a) $F_{\text{key},i}[\mathbf{r}, C_i]$ on input $(\mathbf{x}, \mu, \text{sd})$, first computes 1ABE.msk using the randomness $\text{PRF}(\text{sd}, \mathbf{r})$, i.e.

$$1\text{ABE.msk} \leftarrow 1\text{ABE.Setup}(1^\lambda; \text{PRF}(\text{sd}, \mathbf{r}))$$
 and then outputs $1\text{ABE.sk}_{C_i} \leftarrow 1\text{ABE.KeyGen}_i(1\text{ABE.msk}, C_i)$.
 - (b) $F_{\text{ct}}[\mathbf{r}]$ on input $(\mathbf{x}, \mu, \text{sd})$, first computes 1ABE.msk as above and then outputs an 1ABE.ct encoding message μ w.r.t. attribute \mathbf{x} . It then computes prFE keys $\{\text{prFE.sk}_{\text{key},i}\}_{i \in [|C|]}$ and prFE.sk_{ct} corresponding to functions $\{F_{\text{key}}[\mathbf{r}, i]\}_{i \in [|C|]}$ and $F_{\text{ct}}[\mathbf{r}]$, respectively. It outputs $\text{sk}_C = (\{\text{prFE.sk}_{\text{key},i}\}_{i \in [|C|]}, \text{prFE.sk}_{\text{ct}})$.
- The decryption algorithm on input $\text{sk}_C = (\{\text{prFE.sk}_{\text{key},i}\}_{i \in [|C|]}, \text{prFE.sk}_{\text{ct}})$ and $\text{ct} = \text{prFE.ct}$ first runs the prFE decryption, using prFE keys and ciphertext prFE.ct, to compute

$$F_{\text{key},i}[\mathbf{r}, C_i](\mathbf{x}, \mu, \text{sd}) = 1\text{ABE.sk}_{C_i}, \quad F_{\text{ct}}[\mathbf{r}](\mathbf{x}, \mu, \text{sd}) = 1\text{ABE.ct}.$$

Finally it sets $1\text{ABE.sk}_C = (1\text{ABE.sk}_{C_1}, \dots, 1\text{ABE.sk}_{C_{|C|}})$ and outputs the decryption result as $1\text{ABE.Dec}(1\text{ABE.sk}_C, 1\text{ABE.ct})$.

Correctness follows from those of the prFE and 1ABE: By the correctness of prFE, the decryptor recovers the ciphertext and secret key pair of 1ABE and by the correctness

of 1ABE , one can recover the message μ when $C(\mathbf{x}) = 1$. To prove the security, it suffices to show that $\text{ct} = \text{prFE.ct}$ is pseudorandom. By the security of prFE , it suffices to show that the decryption results of the ciphertext using the secret keys are jointly pseudorandom. This follows from the security of 1ABE , since the decryption results are ciphertext and secret key pairs, where the decryption is not possible for each pair. Please see Section 5.B for further details.

Building 1ABE . It remains to instantiate 1ABE with the desired properties. Fortunately, these properties are relatively weak and easily satisfied. For instance, we can instantiate 1ABE simply by using blind garbled circuits [58] (Section 5.2.1) – here *blindness* is precisely the property that the garbled circuit and its labels should be pseudorandom, when the evaluation result of the garbled circuit using the given labels is random. Given a blind garbled circuit, the labels of the garbled circuit form the 1ABE ciphertext, the set of garbled gates form the 1ABE key, and decomposability follows from the structure of a garbled circuit, where we can garble each gate independently. Care needs to be taken to modify the circuit C in the secret key so that when $C(\mathbf{x}) = 0$, it outputs a random string rather than \perp so as to be compatible with our prFE . This yields the following theorem.

Theorem 5.7. *Assuming LWE and Evasive LWE , there exists a very selectively secure kpABE scheme for circuits of unbounded depth and attribute length ℓ with $|\text{mpk}| = \ell \cdot \text{poly}(\lambda)$, $|\text{sk}_C| = |C| \cdot \ell \cdot \text{poly}(\lambda)$, $|\text{ct}| = \ell \cdot \text{poly}(\lambda)$.*

Note that HLL relied on a new assumption called circular evasive LWE , which we replace by simple evasive LWE above albeit at the cost of larger parameters – in particular the secret key is large and scales with $|C|$.

Next, we show that by using the abstraction of Attribute Based Laconic Function Evaluation (abLFE) [153], we can match the parameters by HLL. Intuitively abLFE allows to compress a circuit C into a short digest, which is then used by an encryptor to compute a ciphertext ct for some attribute, message pair (\mathbf{x}, μ) . The decryptor, given C ,

ct can recover μ if and only if $C(\mathbf{x}) = 1$. For our compiler, we require an **abLFE** scheme where the encryption algorithm can be decomposed into an offline and an online phase. The offline encryption algorithm takes as input (\mathbf{x}, μ) and outputs ct_{off} and a private state st . The online encryption algorithm takes as input $(\text{st}, \text{digest})$ and outputs ct_{on} . This property is satisfied by the construction of [122].

At a high level, our **kpABE** uses the compression of the **abLFE** to shorten the secret key. The key generation computes a digest \hat{C} for the circuit C , then computes a **prFE** key for a circuit which outputs the online part of **abLFE** ciphertext. The encrypt algorithm computes the offline part of the **abLFE** ciphertext and the state st using input (\mathbf{x}, μ) . It then encrypts st using **prFE** encryption. Now, **prFE** decryption allows to recover the online part of the ciphertext, and **abLFE** decryption allows to recover μ if $C(\mathbf{x}) = 1$. We refer the reader to Section 5.B.3 for details. We prove the following theorem:

Theorem 5.8. *Under the circular LWE assumption and the Evasive LWE assumption, there exists a very selectively secure **kpABE** scheme for circuits of unbounded depth and attribute length ℓ with*

$$|\text{mpk}| = \text{poly}(\ell, \lambda), \quad |\text{sk}_C| = \text{poly}(\lambda), \quad |\text{ct}| = \text{poly}(\ell, \lambda).$$

Above, the parameters achieved match those of HLL and the assumptions are strictly weaker.

Constructing kpABE for Turing machines from Weaker Assumptions. Next, we turn our attention to **kpABE** for Turing machines. We show that using **prFE** and a **kpABE** for unbounded depth circuits, we can build a **cpABE** for unbounded depth circuits and further a **kpABE** for Turing machines with parameters matching AKY. The construction for **cpABE** is as follows.

- The setup algorithm generates $(\text{prFE.msk}, \text{prFE.mpk})$ using prFE.Setup and outputs these as cpABE.msk and cpABE.mpk respectively.

- The encryption algorithm, given circuit C and message μ , works as follows: It samples randomness R_{key} and computes $(\text{kpABE.mpk}, \text{kpABE.msk}) = \text{kpABE.Setup}(1^\lambda, 1^L; R_{\text{key}})$. Next, it computes a prFE.ct encoding $(R_{\text{key}}, \text{sd}, \mu)$, where sd is a PRF key, and a kpABE secret key for circuit C as $\text{kpABE.sk}_C \leftarrow \text{kpABE.KeyGen}(\text{kpABE.msk}, C)$. It outputs $\text{cpABE.ct} := (\text{prFE.ct}, \text{kpABE.mpk}, \text{kpABE.sk}_C)$.
- The key generation algorithm, given msk and attribute \mathbf{x} outputs a prFE key, prFE.sk_F , for the function $F[\mathbf{x}, \mathbf{r}]$ where $\mathbf{r} \leftarrow \{0, 1\}^\lambda$. It outputs $\text{cpABE.sk}_x := \text{prFE.sk}_F$.
The function $F[\mathbf{x}, \mathbf{r}]$ on input $(R_{\text{key}}, \text{sd}, \mu)$ first computes $(\text{kpABE.mpk}, \text{kpABE.msk})$ using the randomness R_{key} and then outputs a kpABE ciphertext kpABE.ct encoding μ w.r.t. attribute \mathbf{x} using randomness $\text{PRF}(\text{sd}, \mathbf{r})$.
- The decryption algorithm on input secret key prFE.sk_F and ciphertext $\text{cpABE.ct} := (\text{prFE.ct}, \text{kpABE.mpk}, \text{kpABE.sk}_C)$ first runs the prFE decryption to obtain $F[\mathbf{x}, \mathbf{r}](R_{\text{key}}, \text{sd}, \mu) = \text{kpABE.ct}$ and finally performs kpABE decryption using kpABE.ct and kpABE.sk_C .

Correctness follows from those of kpABE and prFE: The decryption of prFE ciphertext using the prFE secret key yields kpABE ciphertext encrypted under \mathbf{x} . This kpABE ciphertext can be decrypted using kpABE.sk_C when $C(\mathbf{x}) = 1$. To prove the security, we show prFE.ct is pseudorandom. By the security of prFE, it suffices to show that the decryption results of prFE.ct are pseudorandom. This follows from the security of kpABE, since the decryption results are kpABE ciphertexts which cannot be decrypted by kpABE.sk_C . We finally note that while the above overall proof strategy works when the underlying kpABE has pseudorandom ciphertext, we have to consider more general case where underlying kpABE does not necessarily have pseudorandom ciphertext. Therefore, the final construction as well as the security proof are slightly different. We refer the reader to Section 5.7 for details. Thus, we obtain the following theorem.

Theorem 5.9. *Under the circular LWE assumption and the Evasive LWE assumption, there exists a very selectively secure cpABE scheme for circuits of unbounded depth and attribute length ℓ with*

$$|\text{cpABE.mpk}| = \text{poly}(\lambda), |\text{cpABE.sk}_x| = \text{poly}(\ell, \lambda), |\text{cpABE.ct}| = \text{poly}(\lambda).$$

We note that the **cpABE** scheme instantiated as above replaces the reliance on circular tensor **LWE** and **LWE** assumptions used by **AKY** by simply circular **LWE**. It also achieves a shorter **mpk** as compared to that of **AKY** (which is $\text{poly}(\lambda, \ell)$), other parameters being the same.

AKY provided a compiler that uses **kpABE** for bounded depth circuits and **cpABE** for unbounded depth circuits to achieve **kpABE** for Turing machines. Plugging our new **cpABE** into this compiler, we obtain:

Corollary 5.9.1. *Under the circular **LWE** assumption and evasive **LWE** assumption, there exists a very selectively secure **ABE** for **TM** with the following parameters:*

$$|\text{mpk}| = \text{poly}(\lambda), \quad |\text{sk}| = \text{poly}(\lambda, |M|), \quad |\text{ct}| = \text{poly}(\lambda, |\mathbf{x}|, t).$$

In terms of parameters, the above theorem matches those of **AKY**. In terms of assumptions, it provides a strict improvement, as discussed above. Later, by instantiating the underlying **kpABE** to be one with optimal parameters (constructed below), we will obtain a **cpABE** scheme for unbounded depth circuits with optimal parameters – please see Section 5.7 for details.

The Quest for Optimal Parameters

Achieving optimal parameters is a central open question in the construction of **ABE** schemes, and one which has received much attention in the literature [127, 170]. Ideally, we wish to obtain optimal parameters for **kpABE** and **cpABE** for *unbounded* depth circuits and moreover, from the weakest assumptions possible. So far, we do not know how to step around the usage of evasive **LWE** either for unbounded depth **kpABE** or even for bounded depth **cpABE**, so an outstanding open question is:

*Can we construct **kp/cp ABE** for unbounded depth circuits with optimal parameters from **LWE** and evasive **LWE**?*

Below, we answer the above question in the affirmative. To achieve optimality, we will need a result from our companion paper [14], which we use black-box below. While reliance on the companion work is not needed to improve the assumptions and match the parameters from HLL and AKY as discussed above, we are aiming for optimality, which necessitates this dependence as discussed previously.

Wishful Thinking: Obfuscation to the Rescue? Since we do not wish to rely on any kind of circularity assumption, we return to our construction of kpABE from 1ABE (which we built using blind garbled circuits) and prFE as a starting point, and ask if we can compress the parameters. Observe that the kpABE secret key in this case contains prFE secret keys for outputting the randomized versions of the 1ABE ciphertext and secret key, i.e. $\text{sk}_C = \{\text{prFE.sk}_{\text{key},i}, \text{prFE.sk}_{\text{ct},j}\}_{\{i,j\}}$. As discussed above, this approach crucially relies on the decomposability of the secret key and ciphertext of the underlying 1ABE scheme, which allows it to randomize and output these “piece-wise” – gate by gate for the circuit and bit by bit for the input.

Suppose, as wishful thinking, one could have an obfuscation of a program which, given an index within the 1ABE ciphertext and secret key, could output the appropriate randomized garbled gate/label – this would allow to make the secret key and ciphertext size independent of the circuit and attribute size, and allow us to achieve optimal compression. Can such an approach be realized?

Let us examine the possibility. It is well known [29, 41] that compact FE for general circuits can be compiled into compact multi-input FE, which in turn can be compiled into iO for general circuits. Perhaps we can hope that such a compiler also be applied to our compact prFE to obtain multi-input prFE and iO for pseudorandom functionalities. Indeed this is true – in our companion work [14], we show that a similar compiler as [29, 41] can be made to work, albeit via a very different proof technique (since our security notion is itself quite different). However, even assuming the existence of iO for

pseudorandom functionalities, the problem is not solved because:

1. The FE to iO compiler suffers exponential loss in the reduction. Therefore, if we use iO, we do not get security from polynomial assumptions. Since we are finally constructing only ABE schemes, the reliance on exponential security feels too strong.
2. Even if we use iO, the problem of compression does not get solved because we require the program to “authenticate” the input and only provide the output for the legitimate input. As an example, it should not be possible to obtain a ciphertext component for $x_i = 0$ if $x_i = 1$. This would require hardwiring the input \mathbf{x} (similarly C) into the obfuscation which would bring us back to square one!

We resolve these issues by observing that the second issue can, surprisingly, be used to overcome the first. This leaves us with the second issue, to resolve which, we define a new notion which we call *laconic* poly-domain obfuscation for pseudorandom functionalities, which may be of independent interest.

Let us begin with trying to resolve the first issue: observe that the second issue tells us that we need obfuscation for a very special input domain – one that corresponds to the gates of our particular C or the bits of our particular \mathbf{x} , and that inputs outside this domain should not be accepted. This immediately has the effect of shrinking down the domain size from exponential to polynomial! While it is as-yet unclear how to force the domain to be restricted to the special polynomial sized set we require, it at least shows that supporting a polynomial sized domain suffices. We term an obfuscation for circuit with polynomial sized domain as “Poly-Domain Obfuscation”². We define iO for pseudorandom circuits which have polynomial sized domain (denoted by pPRIO) as follows.

1. $\text{Obf}(1^\lambda, C) \rightarrow \text{Obf}$. The obfuscation algorithm takes as input the security parameter λ and a circuit $C : [N] \rightarrow [M]$ with $\text{size}(C) \leq L$ for some arbitrary polynomial $L = L(\lambda)$. It outputs an obfuscation of the circuit Obf .
2. $\text{Eval}(\text{Obf}, x) \rightarrow y$. The evaluation algorithm takes as input an obfuscated circuit Obf and an input $x \in [N]$. It outputs $y \in [M]$.

²The informed reader may wonder about the connection with XIO [144] – note that in XIO, the size of the obfuscated circuit is allowed to be only slightly sublinear in the size of the truth table. In contrast, the efficiency requirement of our notion is the same as standard iO.

Security states that $\text{Obf}(1^\lambda, C)$ is pseudorandom if the truth table of C is pseudorandom. Please see Section 5.2.3 for a formal definition. Restricting the input space to be of polynomial size helps us to avoid the exponential loss of the transformation as discussed above. We now invoke a theorem from our companion work:

Theorem 5.10 ([14]). *Assuming LWE and evasive LWE assumptions, there exists a secure pPRIO scheme.*

We now turn to the second issue, which we resolve by defining a new primitive as described next.

Laconic Poly-Domain Obfuscation for Pseudorandom Functionalities. In pPRIO, we obfuscate circuits with a polynomial-sized input domain of the form $1, 2, \dots, N$. However, we must now find a way to have more control on what this polynomial set can be, so that we can implement some “authentication” of the inputs without making the size of the obfuscated circuit grow with the domain size, as discussed above. Towards this, we extend pPRIO to introduce the new notion of *laconic* pPRIO, where the input domain can be defined as $X := X_1, \dots, X_N$ for arbitrary strings $X_1, \dots, X_N \in \{0, 1\}^\ell$, with arbitrary length ℓ . The obfuscated circuit allows for the evaluation of inputs that are in the set $X := \{X_1, \dots, X_N\}$, but it does not allow evaluation for any inputs outside this set. This is exactly what we want! Looking ahead, the set X_1, \dots, X_N will be instantiated with the gates of C or the bits of \mathbf{x} in the final construction.

To define laconic pPRIO, we modify the pPRIO syntax so that, in order to obfuscate a circuit with a restricted input domain X , the obfuscator only needs a short digest of X , whose size does not depend on N , rather than the entire description of it. This results in a compact obfuscation whose size is independent of N . Now, the evaluation of an input belonging to X can be performed given the description of X in the clear. We present our definition next.

1. $\text{LDigest}(1^\lambda, X = \{X_i\}_{i \in [N]}) \rightarrow \text{dig}$. The digest algorithm takes as input the

security parameter λ and an input space X of the form $X = \{X_i \in \{0, 1\}^\ell\}_{i \in [N]}$ for some $\ell = \ell(\lambda)$ and $N \in \mathbb{N}$. We assume that X encodes the information of ℓ and N and one can retrieve them efficiently. It outputs a string **dig**.

2. $\text{LObfuscate}(1^\lambda, \text{dig}, E) \rightarrow \text{Lobf}$. The encode algorithm takes as input the security parameter λ , string **dig** and a circuit $E : \{0, 1\}^\ell \rightarrow \{0, 1\}^L$ whose size is S . It outputs a ciphertext **Lobf**.
3. $\text{LEval}(X, \text{Lobf}) \rightarrow Y$. The decode algorithm takes as input $X = \{X_i \in \{0, 1\}^\ell\}_{i \in [N]}$ and **Lobf**. It outputs $Y = \{Y_i \in \{0, 1\}^L\}_{i \in [N]}$.

For efficiency, we need that the size of **dig** = $\text{LDigest}(1^\lambda, X)$ should be bounded by $\text{poly}(\lambda, S)$. In particular, the size of **dig** should be independent of N . For security, we require (roughly) that:

$$\begin{aligned} \text{If } (X, E(X_1), \dots, E(X_N)) &\approx_c (X, \Delta_1 \leftarrow \{0, 1\}^L, \dots, \Delta_N \leftarrow \{0, 1\}^L) \\ \text{then } (X, \text{LObfuscate}(1^\lambda, \text{dig}, E)) &\approx_c (X, \delta \leftarrow \mathcal{O}) \end{aligned}$$

where $\text{dig} \leftarrow \text{LDigest}(1^\lambda, X = \{X_i\}_{i \in [N]})$ and \mathcal{O} denotes the output space of **LObfuscate** algorithm.

Constructing Laconic Poly-Domain Obfuscation. It remains to build laconic pPRIO.

To do so, we use pPRIO in addition to blind garbled circuits (bGC) and blind batch encryption (BBE). Since pPRIO and bGC have been introduced earlier, we recall the notion of BBE [58] here. Intuitively, BBE allows to hash an input \mathbf{x} so that encryption is performed against the hash of \mathbf{x} instead of \mathbf{x} together with some index $i \in [|\mathbf{x}|]$, and decryption allows to recover one of two messages depending on the i^{th} bit of \mathbf{x} ³.

In more detail, a BBE scheme $\text{BBE} = (\text{Setup}, \text{Gen}, \text{SingleEnc}, \text{SingleDec})$ is defined as follows: The setup algorithm on input security parameter λ and key length N outputs a common reference string **crs**. The **Gen** algorithm on input **crs** and a string $\mathbf{x} \in \{0, 1\}^N$ (to be used as secret key) outputs a hash value h to be used as a public key. The

³The informed reader may notice similarities with laconic oblivious transfer [73].

SingleEnc algorithm on input $(\text{crs}, h, i \in [N], (m_0, m_1) \in \mathcal{M}^2)$ outputs a (single) ciphertext ct . The SingleDec algorithm on input $(\text{crs}, \mathbf{x}, i \in [N], \text{ct})$ outputs a message $m \in \mathcal{M}$. For correctness, it is required that m_b is recovered if $x_i = b$ and h is the hash of \mathbf{x} . Security requires the other message to remain hidden. The blindness of the scheme states that if the encrypted message is random then the ciphertext is indistinguishable from random. We give a construction of a BBE scheme from LWE satisfying these properties in Section 5.A.

The basic intuition behind the construction of laconic pPRIO (denoted by LprIO) is as follows. To restrict the circuit evaluation to a special input domain X , we use the BBE algorithm Gen to generate the hash h of X and output this as part of the digest. The obfuscate algorithm, given as input a circuit E now provides a pPRIO obfuscation of another circuit which garbles E and generates BBE encryptions of its labels using the hash. The BBE ciphertext can be decrypted to recover the appropriate label, using which, the garbled circuit can be executed. In more detail:

1. The setup algorithm compresses the input domain $X = \{X_i \in \{0, 1\}^\ell\}_{i \in [N]}$ by hashing it using BBE.Gen into a short hash value h . It outputs $\text{dig} = h$.
2. The obfuscate algorithm LObfuscate($1^\ell, \text{dig}, E$) outputs a pPRIO obfuscation, pPRIO.Obf, of $C[\text{dig}, E]$.
The circuit $C[\text{dig}, E]$ on input $i \in [N]$ first computes bGC garbled circuit \tilde{E}_i and input labels $\text{lab}_{i,j,b}$ for $j \in [\ell], b \in \{0, 1\}$ – next, it encrypts the bGC labels $(\text{lab}_{i,j,0}, \text{lab}_{i,j,1})$ using BBE encryption and finally outputs \tilde{E}_i and BBE ciphertexts $\text{BBE.ct}_{i,j}$. The BBE ciphertexts are encrypted so that one can retrieve the labels for X_i for \tilde{E}_i if a decryptor knows X .
3. The LEval(X, Lobf) algorithm first runs pPRIO evaluation on pPRIO.Obf and $i \in [N]$ to obtain bGC garbled circuit \tilde{E}_i and BBE ciphertexts $\text{BBE.ct}_{i,j}$ for $j \in [\ell]$. Next it runs BBE decryption using X , index $(i - 1)\ell + j$ (this corresponds to j 'th bit of X_i) and $\text{BBE.ct}_{i,j}$ to get labels $\{\text{lab}_{i,j}\}_{j \in [\ell]}$ corresponding to X_i . Finally, it runs bGC evaluation using \tilde{E}_i and $\{\text{lab}_{i,j}\}_{j \in [\ell]}$ to obtain $E(X_i)$ for all $i \in [N]$.

Correctness is immediate given the decryption outline above. We then discuss security. Our goal is to show that pPRIO.Obf is pseudorandom assuming $E(X_1), \dots, E(X_N)$ are pseudorandom. By the security guarantee of pPRIO, it suffices to prove that the truth

table $\{C[\text{dig}, E](i) = (\tilde{E}_i, \{\text{BBE.ct}_{i,j}\}_j)\}_{i \in [N]}$ of the circuit $C[\text{dig}, E]$ is pseudorandom. To show this, we first replace the labels that are *not* corresponding to X_i encrypted inside $\text{BBE.ct}_{i,j}$ with random ones using the security of BBE. Then, by the security of bGC, we replace the garbled circuit \tilde{E}_i and labels corresponding to X_i with the simulated one, computed from $E(X_i)$. We then replace $\{E(X_i)\}_i$ with random strings, which can be done by invoking the assumption of the security definition. Then, by the blindness of bGC, \tilde{E}_i and labels can be replaced with random strings. Finally, by the blindness of BBE, we can replace $\{\text{BBE.ct}_{i,j}\}_j$ with random strings, since they encrypt random strings.

Unfortunately, the above proof strategy does *not* work as described. This is because we assumed in our description above that the output of the pPRIO is pseudorandom, which is required for the security of LprIO to hold. However, unfortunately, this is not the case. Rather, what we have is that the obfuscation of pPRIO can be divided into online and offline parts and that the online part of it is pseudorandom, where the offline part is not dependent on the circuit being obfuscated. Furthermore, the offline part can be reused across invocations. We can show similar property for LprIO, which will be sufficient for the upcoming application of LprIO to PHprFE. We refer the reader to Section 5.4 for details. We will get back to the issue when we discuss the security proof for our PHprFE.

Optimal Parameters at Last.

Using the techniques developed above, we construct a partially hiding compact FE for pseudorandom functionalities, denoted by PHprFE. The syntax of a PHprFE = (Setup, KeyGen, Enc, Dec) scheme is similar to that of a prFE scheme where the input to the encryption scheme now consists of a public part and a secret part. In more detail, the encryption algorithm takes as input $(\mathbf{x}_{\text{pub}}, \mathbf{x}_{\text{priv}})$ and outputs a ciphertext ct . The ciphertext ct can be decrypted using sk_f and \mathbf{x}_{pub} to recover $f(\mathbf{x}_{\text{pub}}, \mathbf{x}_{\text{priv}})$. The security of the scheme, at a high level, states that if $f(\mathbf{x}_{\text{pub}}, \mathbf{x}_{\text{priv}})$ is pseudorandom then

the ciphertext encoding $(\mathbf{x}_{\text{pub}}, \mathbf{x}_{\text{priv}})$ is indistinguishable from random. We note that a PHprFE scheme implies a prFE scheme if we set $\mathbf{x}_{\text{pub}} = \perp$. We provide a PHprFE construction that achieves optimal parameters and supports circuits of unbounded depth. This in turn will be used to construct optimal kpABE and cpABE for unbounded depth circuits.

Partially Hiding Pseudorandom FE. As discussed above, to get parameters independent of sizes of C and \mathbf{x}_{pub} , we compress them using LprIO scheme. Now, instead of letting the prFE keys directly output labels corresponding to the attribute \mathbf{x} or the garbled circuit corresponding to C , we have it output an obfuscation of circuits that compute these components, as described above. In more detail, we use LprIO, bGC and prFE scheme to construct a PHprFE scheme as follows.

1. The PHprFE setup algorithm generates $(\text{prFE.msk}, \text{prFE.mpk})$ using prFE.Setup and outputs these as msk and mpk respectively.
2. The PHprFE keygen algorithm, given as input a circuit C , does the following.
 - First it uses LprIO scheme to compress C , i.e. compute $\text{dig}_C \leftarrow \text{LDigest}(\{i, C_i\})$ where C_i represents the i^{th} gate of C .
 - Next it samples $\mathbf{r} \leftarrow \{0, 1\}^\lambda$ and gives prFE.sk for the circuit $F[\mathbf{r}, \text{dig}_C]$. The circuit $F[\mathbf{r}, \text{dig}_C]$ on input $(\text{dig}_{\mathbf{x}_{\text{pub}}}, \mathbf{x}_{\text{priv}})$ outputs the following :
 - bGC garbled labels corresponding to \mathbf{x}_{priv} .
 - Obfuscation, LprIO.Obf_1 , of a circuit that outputs bGC labels corresponding to \mathbf{x}_{pub} index by index. We only require short digest of \mathbf{x}_{pub} , $\text{dig}_{\mathbf{x}_{\text{pub}}}$, to compute LprIO.Obf_1 .
 - Obfuscation, LprIO.Obf_2 , of a circuit that outputs garbled gates of C , gate by gate. We only require short digest of C , dig_C , to compute LprIO.Obf_2 .
 - Outputs $\text{sk}_C = (\mathbf{r}, \text{prFE.sk})$.
3. The PHprFE encryptor on input mpk and $\mathbf{x} = (\mathbf{x}_{\text{pub}}, \mathbf{x}_{\text{priv}})$ does the following:

- Run $\text{dig}_{\mathbf{x}_{\text{pub}}} \leftarrow \text{LDigest}(1^\lambda, \{(i, \mathbf{x}_{\text{pub},i})\}_{i \in [L_{\text{pub}}]})$, where $\mathbf{x}_{\text{pub},i} \in \{0, 1\}$ is the i -th bit of \mathbf{x}_{pub} .

- Run $\text{prFE.ct} \leftarrow \text{prFE.Enc}(\text{prFE.mpk}, (\text{dig}_{\mathbf{x}_{\text{pub}}}, \mathbf{x}_{\text{priv}}))$.

We note that here the input size of prFE is independent of $|\mathbf{x}_{\text{pub}}|$.

4. The PHprFE decryptor does the following:

- It first runs the prFE decryption to obtain garbled labels corresponding to \mathbf{x}_{priv} and obfuscations LprIO.Obf_1 and LprIO.Obf_2 .
- Next it performs LprIO evaluations using $(\{i, \mathbf{x}_{\text{pub},i}\}_i, \text{LprIO.Obf}_1)$ and $(\{i, C_i\}_i, \text{LprIO.Obf}_2)$ to get the garbled labels corresponding to \mathbf{x}_{pub} and garbled circuit of C , respectively. Here $\mathbf{x}_{\text{pub},i}$ denotes the i -th bit of \mathbf{x}_{pub} and C_i denotes the i -th gate of C .
- Finally, it uses evaluation of the garbling scheme to evaluate the garbled circuit of C on the labels of $\mathbf{x} = (\mathbf{x}_{\text{pub}}, \mathbf{x}_{\text{priv}})$.

Correctness of the above scheme follows from the correctness of the underlying ingredients.

For security we want to show that the $\text{ct} = \text{prFE.ct}(\text{dig}_{\mathbf{x}_{\text{pub}}}, \mathbf{x}_{\text{priv}})$ is pseudorandom. The adversary additionally sees $\text{mpk}, \text{sk}_C = (\mathbf{r}, \text{prFE.sk})$ and we have the guarantee that $C(\mathbf{x}_{\text{pub}}, \mathbf{x}_{\text{priv}})$ is pseudorandom. We prove security via the following three steps.

- First, we invoke prFE security for function $F[\mathbf{r}, \text{dig}_C]$. By security guarantee of a prFE scheme, it suffices to show pseudorandomness of $F[\mathbf{r}, \text{dig}_C](\text{dig}_{\mathbf{x}_{\text{pub}}}, \mathbf{x}_{\text{priv}}) = (\text{lab}_{\mathbf{x}_{\text{priv}}}, \text{LprIO.Obf}_1, \text{LprIO.Obf}_2)$.
- Next, we invoke the security of LprIO scheme using which we wish to show the pseudorandomness of LprIO.Obf_1 and LprIO.Obf_2 . Recall that the security of LprIO scheme states that it suffices to show pseudorandomness of circuit's output on the compressed input domain. Thus, it suffices to show the pseudorandomness of $\text{lab}_{\mathbf{x}_{\text{pub}}}$ and \tilde{C} .

Now, we are left with showing the pseudorandomness of $\text{lab}_{\mathbf{x}_{\text{priv}}}$, $\text{lab}_{\mathbf{x}_{\text{pub}}}$ and \tilde{C} .

- Next, we use bGC simulator to output the garbled values. That is $(\tilde{C}, \text{lab}_{\mathbf{x}_{\text{priv}}}, \text{lab}_{\mathbf{x}_{\text{pub}}}) \leftarrow \text{bGC.Sim}(C(\mathbf{x}_{\text{pub}}, \mathbf{x}_{\text{priv}}))$. Now using the fact that $C(\mathbf{x}_{\text{pub}}, \mathbf{x}_{\text{priv}})$ is pseudorandom, we invoke the *blindness* of bGC scheme to substitute the simulated value with a random string. Hence the security.

While the above overview for the security proof conveys the main idea, it does not work

as it is, since we do not have the ideal LprIO whose obfuscated circuit is pseudorandom. Rather, we have an imperfect one, where only the online part of the obfuscated circuit is pseudorandom. This leads to a mismatch with the security guarantee of prFE , which requires the decryption result to be pseudorandom. To resolve the above issue, we change the construction so that the offline part of the LprIO obfuscation, which is not necessarily pseudorandom, is generated during the encryption and included into the ciphertext. This change is possible since the offline part of the obfuscation is not dependent on the circuit being obfuscated, which is unknown to the encryptor. Furthermore, the (pseudorandom) online part of the obfuscated circuit is generated during the decryption, which is aligned with the security guarantee of prFE . A subtlety that arises here is that this change requires strong security guarantee for LprIO , where pseudorandomness of the online part should hold even if its offline part is reused across many invocations. Luckily, we can prove such security for LprIO and thus the entire proof works. We refer the reader to Section 5.5 for details.

We reason about efficiency next. First we note that it suffices to use a bounded depth prFE scheme supporting circuits of depth $\text{poly}(\lambda, |\mathbf{x}_{\text{priv}}|)$. Next, we note that the size of prFE input and output is also bounded by $\text{poly}(\lambda, |\mathbf{x}_{\text{priv}}|)$. Thus we get a PHprFE scheme with $|\text{mpk}| = |\text{ct}| = |\text{sk}| = \text{poly}(\lambda, |\mathbf{x}_{\text{priv}}|)$. The sizes of the master public key, ciphertexts, and the secret keys of our construction above are all independent from the length of \mathbf{x}_{pub} and C . However, they still depend on the length of \mathbf{x}_{priv} . In Section 5.5.4 we show how to remove this dependency from the master public key and the secret keys using a SKE scheme. We get the following theorem.

Theorem 5.11. *Assuming LWE and evasive LWE assumptions, there exists a partially hiding pseudorandom FE scheme, for circuit class $\mathcal{C} = \{C : \{0, 1\}^{L_{\text{pub}}} \times \{0, 1\}^{L_{\text{priv}}} \rightarrow \{0, 1\}\}$, that satisfies reusable security (as per Definition 5.20) whose sizes of the master public key and the secret key are fixed polynomial $\text{poly}(\lambda)$. Furthermore, the size of the ciphertext is $L_{\text{priv}} + \text{poly}(\lambda)$.*

Optimal kpABE via PHprFE. Finally we are ready to show how to obtain an optimal kpABE scheme using an optimal PHprFE scheme. The construction is straightforward: we set the attribute as the public input of PHprFE scheme and the message as private input. Additionally we hardwire a nonce in the circuit C during keygen so that it outputs a pseudorandom value when $C(\mathbf{x}) = 0$, using some private input \mathbf{sd} . In more detail,

1. The setup generates $(\text{PHprFE.msk}, \text{PHprFE.mpk})$ using PHprFE.Setup and outputs these as \mathbf{msk} and \mathbf{mpk} respectively.
2. The key generator on input \mathbf{msk} and a circuit C does the following
 - Defines circuit $C[\mathbf{r}]$, for $\mathbf{r} \leftarrow \{0, 1\}^\lambda$, which on input $(\mathbf{x}, \mu, \mathbf{sd})$ outputs μ if $C(\mathbf{x}) = 1$ and $\text{PRF}(\mathbf{sd}, \mathbf{r})$ otherwise.
 - Computes a PHprFE.sk for circuit $C[\mathbf{r}]$.
 - Outputs $\mathbf{sk}_C := (\mathbf{r}, \text{PHprFE.sk})$.
3. The encryptor on input \mathbf{mpk} , \mathbf{x} , and μ , first samples a PRF seed $\mathbf{sd} \leftarrow \{0, 1\}^\lambda$, sets $\mathbf{x}_{\text{pub}} = \mathbf{x}$, $\mathbf{x}_{\text{priv}} = (\mu, \mathbf{sd})$ and outputs a PHprFE ciphertext $\text{PHprFE.ct} \leftarrow \text{PHprFE.Enc}(\text{PHprFE.mpk}, \mathbf{x}_{\text{pub}}, \mathbf{x}_{\text{priv}})$.
4. The decryptor on input the secret key and the ciphertext simply runs the PHprFE decryption.

The correctness of the scheme follows in a straightforward manner from the correctness of the PHprFE scheme. Note that the PHprFE decryption will output $C[\mathbf{r}](\mathbf{x}, \mu, \mathbf{sd}) = \mu$ for $C(\mathbf{x}) = 1$. Security of the scheme is implied by the security of the underlying PHprFE scheme and the PRF. To prove security, we need to show the pseudorandomness of $\text{ct} = \text{PHprFE.ct}$ when $C(\mathbf{x}) = 0$ – this follows by our usage of a PRF to generate the output, as discussed above. Finally, from the efficiency of a PHprFE scheme and noting that $|\mathbf{x}_{\text{priv}}| = 1 + \lambda$, we get the following theorem.

Theorem 5.12 (Optimal KP-ABE). *Under the LWE and Evasive LWE assumption, there exists very selectively secure KP-ABE scheme supporting circuits $\{C : \{0, 1\}^\ell \rightarrow \{0, 1\}\}$ with unbounded depth and single bit message space with*

$$|\mathbf{mpk}| = \text{poly}(\lambda), |\mathbf{sk}_C| = \text{poly}(\lambda), |\text{ct}| = \text{poly}(\lambda).$$

We also obtain a *predicate* encryption scheme for unbounded depth circuits with optimal parameters, where predicate encryption allows to hide the entire input of the ciphertext against restricted adversaries. The optimal **kpABE** implies an optimal **cpABE** for unbounded depth circuits as well as optimal **kpABE** for Turing machines as discussed above. For more details, please see Section 5.6.

Organization of the Chapter We define the preliminaries used in this work in Section 5.2. In Section 5.3 we define and construct the notion of (bounded-depth) compact functional encryption scheme for pseudorandom functionalities (**prFE**). In Section 5.4 we define the notion of laconic pseudorandom poly-domain obfuscation scheme (**LprIO**) and construct it using a blind garbling scheme, a blind batch encryption scheme (constructed in Section 5.A) and a **pPRIO** scheme (constructed in our companion paper [14]). In Section 5.5 we define the notion of partially-hiding **prFE** (denoted as **PHprFE**) and construct unbounded-depth **PHprFE** using a bounded depth **prFE** and a **LprIO**. In Section 5.6, we use unbounded-depth **PHprFE** to construct unbounded-depth **kpABE** and unbounded-depth **PE** with optimal parameters. In Section 5.7 we construct unbounded-depth **cpABE** with optimal parameters using our optimal **kpABE** and unbounded-depth **prFE** (as a special case of unbounded-depth **PHprFE**).

In Section 5.B we show an alternate pathway to construct an unbounded-depth **kpABE** scheme with sub-optimal parameters *without* using any building block from our companion paper [14] but still improving the state of the art.

5.2 PRELIMINARIES

In this section we introduce some additional preliminaries for this chapter. We refer to Chapter 2 for preliminaries on lattices, Section 4.2.7 for properties of GSW Homomorphic encryption and evaluation and Section 4.2.8 for properties of BGG^+ homomorphic evaluation procedures.

Definition 5.1 (Symmetric Key Encryption with Pseudorandom Ciphertext). A symmetric key encryption scheme for message space $\mathcal{M} = \{\mathcal{M}_\lambda\}_{\lambda \in [\mathbb{N}]}$ and key space $\mathcal{K} = \{\mathcal{K}_\lambda\}_{\lambda \in [\mathbb{N}]}$ and ciphertext space $\mathcal{CT}_{\text{SKE}} = \{\mathcal{CT}_{\text{SKE},\lambda}\}_{\lambda \in [\mathbb{N}]}$ has the following syntax:

$\text{Setup}(1^\lambda) \rightarrow \text{sk}$. The setup algorithm takes as input the security parameter λ and outputs a secret key sk .

$\text{Enc}(\text{sk}, m) \rightarrow \text{ct}$. The encryption algorithm takes as input the secret key sk and a message $m \in \mathcal{M}_\lambda$ and outputs a ciphertext ct .

$\text{Dec}(\text{sk}, \text{ct}) \rightarrow m'$. The decryption algorithm takes as input a secret key sk and a ciphertext ct and outputs a message $m' \in \mathcal{M}_\lambda$.

Correctness: A SKE scheme is said to be correct if there exists a negligible function $\text{negl}(\cdot)$ such that for all $\lambda \in \mathbb{N}$, for every message $m \in \mathcal{M}_\lambda$, we have

$$\Pr \left[\begin{array}{l} \text{sk} \leftarrow \text{Setup}(1^\lambda); \\ m' = m : \text{ct} \leftarrow \text{Enc}(\text{sk}, m); \\ m' = \text{Dec}(\text{sk}, \text{ct}). \end{array} \right] \geq 1 - \text{negl}(\lambda), \quad (5.5)$$

Security: A SKE scheme is said to have pseudorandom ciphertext if there exists a negligible function $\text{negl}(\cdot)$ such that for all $\lambda \in \mathbb{N}$, for every message $m \in \mathcal{M}_\lambda$, we have

$$\left| \Pr \left[\begin{array}{l} \beta' = \beta : \text{sk} \leftarrow \text{Setup}(1^\lambda); \\ \beta' \leftarrow \mathcal{A}^{\text{Enc}(\text{sk}, \cdot), \text{Enc}^\beta(\text{sk}, \cdot)}. \end{array} \right] - \frac{1}{2} \right| \leq \text{negl}(\lambda), \quad (5.6)$$

where the $\text{Enc}(\text{sk}, \cdot)$ oracle, on input a message m , returns $\text{Enc}(\text{sk}, m)$ and $\text{Enc}^\beta(\text{sk}, \cdot)$ oracle, on input a message m , returns ct_β , where $\text{ct}_0 \leftarrow \text{Enc}(\text{sk}, m)$ and $\text{ct}_1 \leftarrow \mathcal{CT}_{\text{SKE}}$.

5.2.1 Blind Garbled Circuit

Here we provide the definition of a garbling scheme for circuit class $\mathcal{C} = \{C : \{0, 1\}^{\ell_{\text{in}}} \rightarrow \{0, 1\}^{\ell_{\text{out}}}\}$. A garbling scheme for circuit class \mathcal{C} consists of three

algorithms (**Garble**, **Eval**, **Sim**) with the following syntax.

$\text{Garble}(1^\lambda, 1^{\ell_{\text{in}}}, 1^{\ell_{\text{out}}}, C) \rightarrow (\text{lab}, \tilde{C})$. The garbling algorithm takes as input the security parameter λ , the input length ℓ_{in} and output length ℓ_{out} for circuit C , the description of the circuit C , and a random value $\text{st} \in \{0, 1\}^\lambda$ and outputs the labels for input wire of the garbled circuit $\text{lab} = \{\text{lab}_{j,b}\}_{j \in [\ell_{\text{in}}], b \in \{0,1\}}$ where each $\text{lab}_{j,b} \in \{0, 1\}^\lambda$ and the garbled circuit \tilde{C} .

$\text{Eval}(1^\lambda, \tilde{C}, \text{lab}_{\mathbf{x}}) \rightarrow \mathbf{y}$. The evaluation algorithm takes as input the garbled circuit \tilde{C} and labels corresponding to an input $\mathbf{x} \in \{0, 1\}^{\ell_{\text{in}}}$, $\text{lab}_{\mathbf{x}} = \{\text{lab}_{i,x_i}\}_{i \in [\ell_{\text{in}}]}$ where x_i denotes the i -th bit of \mathbf{x} , and it outputs $\mathbf{y} \in \{0, 1\}^{\ell_{\text{out}}}$.

$\text{Sim}(1^\lambda, 1^{|C|}, 1^{\ell_{\text{in}}}, \mathbf{y})$ is a PPT algorithm that takes as input the security parameter, the description length of C , an input length ℓ_{in} and a string $\mathbf{y} \in \{0, 1\}^{\ell_{\text{out}}}$, and outputs a simulated garbled circuit \bar{C} and labels $\bar{\text{lab}}$.

A garbling scheme satisfies the following properties.

Definition 5.2 (Correctness). A garbling scheme is said to be correct if for any circuit $C \in \mathcal{C}$ and any input $x \in \{0, 1\}^{\ell_{\text{in}}}$, the following holds

$$\Pr [\mathbf{y} = C(\mathbf{x}) : (\text{lab}, \tilde{C}) \leftarrow \text{Garble}(1^\lambda, 1^{\ell_{\text{in}}}, 1^{\ell_{\text{out}}}, C); \mathbf{y} \leftarrow \text{Eval}(\tilde{C}, \text{lab}_{\mathbf{x}})] = 1.$$

Definition 5.3 (Simulation Security). A garbling scheme is said to satisfy simulation security if for any circuit $C \in \mathcal{C}$ and any input $\mathbf{x} \in \{0, 1\}^{\ell_{\text{in}}}$, the following holds

$$\{(\tilde{C}, \text{lab}_{\mathbf{x}}) \mid (\text{lab}, \tilde{C}) \leftarrow \text{Garble}(1^\lambda, 1^{\ell_{\text{in}}}, 1^{\ell_{\text{out}}}, C)\} \approx_c \{(\bar{C}, \bar{\text{lab}}) \mid (\bar{C}, \bar{\text{lab}}) \leftarrow \text{Sim}(1^\lambda, 1^{|C|}, 1^{\ell_{\text{in}}}, C(\mathbf{x}))\}$$

where $\text{lab} = \{\text{lab}_{j,b}\}_{j \in [\ell_{\text{in}}], b \in \{0,1\}}$ and $\text{lab}_{\mathbf{x}} = \{\text{lab}_{i,x_i}\}_{i \in [\ell_{\text{in}}]}$.

Definition 5.4 (Blindness). [58] A garbling scheme (**Garble**, **Eval**, **Sim**) is called blind if the distribution $\text{Sim}(1^\lambda, 1^{|C|}, 1^{\ell_{\text{in}}}, \mathbf{y})$ for $\mathbf{y} \leftarrow \{0, 1\}^{\ell_{\text{out}}}$, representing the output of the simulator on a completely uniform output, is indistinguishable from a completely uniform bit string. (Note that the distinguisher must not know the random output value that was

used for the simulation.)

Definition 5.5 (Decomposability). We note that the $\text{Garble}(1^\lambda, 1^{\ell_{\text{in}}}, 1^{\ell_{\text{out}}}, C)$ algorithm can be decomposed, using shared randomness st , as follows : (i) $\text{Garble}_i(1^\lambda, C_i; \text{st})$ for $i \in [|C|]$, where $\text{Garble}_i(1^\lambda, C_i)$ outputs the garbling of i -th gate of the circuit C (denoted by C_i) and (ii) $\text{Garble}_{\text{inp}}(1^\lambda, 1^{\ell_{\text{in}}}, 1^{\ell_{\text{out}}}; \text{st}) = \text{lab}$ which outputs $2 \cdot \ell_{\text{in}}$ labels.

Note that information of a single gate C_i of C can be represented by a binary string of length at most 4λ for example, since it suffices to encode its index, indices of its two incoming wires, and the truth table of the gate.

Theorem 5.13. [58] *Assume that one-way function exists. Then, there exists a blind garbled circuits scheme.*

5.2.2 Blind Batch Encryption

Here we provide the definition of a batch encryption scheme largely adapted from [58]. A batch encryption scheme with the message space $\mathcal{M} = \{\mathcal{M}_\lambda\}_\lambda$ consists of the following algorithms.

$\text{Setup}(1^\lambda, 1^N) \rightarrow \text{crs}$. The setup algorithm takes as input the security parameter λ and key length N , and outputs a common reference string crs .

$\text{Gen}(\text{crs}, \mathbf{x}) \rightarrow h$. The generation algorithm takes as input the common reference string crs and a secret key $\mathbf{x} \in \{0, 1\}^N$. It outputs a public key h .

$\text{SingleEnc}(\text{crs}, h, i, (m_0, m_1)) \rightarrow \text{ct}$. The encryption algorithm takes as input a common reference string crs , the public key h , an index $i \in [N]$, and a message $(m_0, m_1) \in \mathcal{M}^2$ and outputs a (single) ciphertext ct .

$\text{SingleDec}(\text{crs}, \mathbf{x}, i, \text{ct}) \rightarrow m$. The decryption algorithm takes as input a common reference string crs , the secret key \mathbf{x} , an index $i \in [N]$, and a (single) ciphertext ct and outputs a message $m \in \mathcal{M}$.

In [58], they define additional algorithms **Enc** and **Dec**, which can be defined using **SingleEnc** and **SingleDec** above. We omit the definition of these algorithms since they are not used in our paper.

Definition 5.6 (Correctness.). A batch encryption scheme is said to be correct if for any $\lambda, N \in \mathbb{N}$, secret key $\mathbf{x} \in \{0, 1\}^N$, $i \in [N]$, $(m_0, m_1) \in \mathcal{M}^2$, $\text{crs} \leftarrow \text{Setup}(1^\lambda, 1^N)$, and $h \leftarrow \text{Gen}(\text{crs}, \mathbf{x})$ it holds that

$$\Pr[m = m_{x_i} \mid m = \text{SingleDec}(\text{crs}, \mathbf{x}, i, \text{SingleEnc}(\text{crs}, h, i, (m_0, m_1)))] = 1,$$

where x_i is the i -th bit of \mathbf{x} .

Definition 5.7 (Succinctness). A batch encryption scheme is α -succinct if letting $\text{crs} \leftarrow \text{Setup}(1^\lambda, 1^N)$, $h = \text{Gen}(\text{crs}, \mathbf{x})$ for some $\mathbf{x} \in \{0, 1\}^N$, it holds that $|h| \leq \alpha N$. It is said fully succinct if $|h| \leq p(\lambda)$ for some fixed polynomial $p(\lambda)$.

Definition 5.8 (Security). A batch encryption scheme is said to be secure if for any PPT adversary \mathcal{A} , there exists a negligible function $\text{negl}(\cdot)$ such that the following holds

$$\Pr \left[\begin{array}{l} (1^N, \mathbf{x} \in \{0, 1\}^N, i \in [N]) \leftarrow \mathcal{A}(1^\lambda); \\ \text{crs} \leftarrow \text{Setup}(1^\lambda, 1^N); \\ \beta' = \beta : \left(\mathbf{m}^{(0)} = (m_0^{(0)}, m_1^{(0)}), \mathbf{m}^{(1)} = (m_0^{(1)}, m_1^{(1)}) \right) \leftarrow \mathcal{A}(\text{crs}); \\ h \leftarrow \text{Gen}(\text{crs}, \mathbf{x}), \beta \leftarrow \{0, 1\}, \text{ct}_\beta \leftarrow \text{SingleEnc}(\text{crs}, h, i, \mathbf{m}^{(\beta)}); \\ \beta' \leftarrow \mathcal{A}(\text{crs}, \text{ct}_\beta) \end{array} \right] \leq \frac{1}{2} + \text{negl}(\lambda)$$

where we require that $\mathbf{m}^{(0)}, \mathbf{m}^{(1)} \in \mathcal{M}^2$ and $m_{x_i}^{(0)} = m_{x_i}^{(1)}$.

We require somewhat stronger blindness condition than that defined in [58].

Definition 5.9 (Strong Blindness). A batch encryption scheme is said to satisfy blindness if for any PPT adversary \mathcal{A} , there exists a negligible function $\text{negl}(\cdot)$ such that the

following holds

$$\Pr \left[\begin{array}{l} \beta' = \beta : \\ (1^N, \mathbf{x} \in \{0, 1\}^N, i \in [N]) \leftarrow \mathcal{A}(1^\lambda); \\ \text{crs} \leftarrow \text{Setup}(1^\lambda, 1^N), h \leftarrow \text{Gen}(\text{crs}, \mathbf{x}); \\ \mathbf{m} \leftarrow \mathcal{M}^2, \beta \leftarrow \{0, 1\}; \\ \text{ct}_0 \leftarrow \text{SingleEnc}(\text{crs}, h, i, \mathbf{m}), \text{ct}_1 \leftarrow C\mathcal{T}; \\ \beta' \leftarrow \mathcal{A}(\text{crs}, \text{ct}_\beta) \end{array} \right] \leq \frac{1}{2} + \text{negl}(\lambda)$$

where $C\mathcal{T}$ is the ciphertext space of the scheme.

Remark 17. Here, we compare strong blindness defined above with the blindness defined in [58]. In [58], they divide a ciphertext ct into an “offline part” subct_1 and “online part” subct_2 , where subct_1 only depends on the encryption randomness and the CRS, whereas subct_2 may depend on h, i , and \mathbf{m} additionally. They then define blindness as the security notion that essentially requires that subct_2 is pseudorandom, whereas subct_1 may not be. The above security notion is stronger than theirs in that we require the entire ciphertext to be pseudorandom rather than part of it. In other words, we can see our definition as more stringent version of their blindness notion where we require subct_1 to be an empty string.

5.2.3 Poly-Input Obfuscation for Pseudorandom Functionalities

In this section we give the definition of poly-input indistinguishability obfuscation for pseudorandom functionalities (pPRIO), adapted from [14].

Syntax A pPRIO scheme consists of the following algorithms.

$\text{Obf}(1^\lambda, C) \rightarrow \text{Obf}$. The obfuscation algorithm takes as input the security parameter λ and a circuit $C : [N] \rightarrow [M]$ with $\text{size}(C) \leq L$ for some arbitrary polynomial $L = L(\lambda)$. It outputs an obfuscation of the circuit Obf .

We consider a definition where the Obf algorithm can be decomposed into the following two phases.

$\text{ObfOff}(1^\lambda, 1^L) \rightarrow (\text{obf}_{\text{off}}, \text{st})$. The offline obfuscation algorithm takes as input the security parameter λ and the circuit size bound L . It outputs obf_{off} and st .

$\text{ObfOn}(\text{st}, C) \rightarrow \text{obf}_{\text{on}}$. The online obfuscation algorithm takes as input the security parameter st and the circuit C and outputs obf_{on} .

The final output of Obf is $\text{Obf} = (\text{obf}_{\text{off}}, \text{obf}_{\text{on}})$.

$\text{Eval}(\text{Obf}, x) \rightarrow y$. The evaluation algorithm takes as input an obfuscated circuit Obf and an input $x \in [N]$. It outputs $y \in [M]$.

Next, we define the properties of a pPRIO scheme.

Definition 5.10 (Correctness). For all security parameters $\lambda \in \mathbb{N}$, for any $C : [N] \rightarrow [M]$, $L = L(\lambda)$ such that $\text{size}(C) \leq L$ and every input $x \in [N]$, we have that:

$$\Pr[\text{Eval}(\text{Obf}, x) = C(x) \mid \text{Obf} = (\text{obf}_{\text{off}}, \text{obf}_{\text{on}}), (\text{obf}_{\text{off}}, \text{st}) \leftarrow \text{ObfOff}(1^\lambda, 1^L), \text{obf}_{\text{on}} \leftarrow \text{ObfOn}(\text{st}, C)] = 1$$

where the probability is taken over the coin-tosses of the obfuscator Obf .

Definition 5.11 (Security). Let Samp be a PPT algorithm that on input 1^λ , outputs

$$(1^{N_1+N_2+\dots+N_Q}, 1^L, \text{aux}, C^1, \dots, C^Q), \quad \text{where } C^i : [N_i] \rightarrow [M_i], \text{ size}(C^i) \leq L$$

where we enforce Samp to output $1^{N_1+N_2+\dots+N_Q}$ to make sure that all N_i are bounded by $\text{poly}(\lambda)$. We say that a pPRIO scheme is secure with respect to the sampler class \mathcal{SC} if for every PPT sampler $\text{Samp} \in \mathcal{SC}$ the following holds.

$$\text{If } (\text{aux}, \{C^1(i)\}_{i \in [N_1]}, \dots, \{C^Q(i)\}_{i \in [N_Q]}) \approx_c (\text{aux}, \{\Delta_i^1\}_{i \in [N_1]}, \dots, \{\Delta_i^Q\}_{i \in [N_Q]}) \quad (5.7)$$

$$\text{then } (\text{aux}, \text{obf}_{\text{off}}, \text{obf}_{\text{on}}^1, \dots, \text{obf}_{\text{on}}^Q) \approx_c (\text{aux}, \text{obf}_{\text{off}}, \delta^1 \leftarrow C\mathcal{T}^1, \dots, \delta^Q \leftarrow C\mathcal{T}^Q), \quad (5.8)$$

where $\Delta_i^j \leftarrow [M_j]$ for $j \in [Q], i \in [N_j]$, $(\text{obf}_{\text{off}}, \text{st}) \leftarrow \text{ObfOff}(1^\lambda, 1^L)$, $\text{obf}_{\text{on}}^j \leftarrow \text{ObfOn}(\text{st}, C^j)$ for $j \in [Q]$, and $C\mathcal{T}^j$ denotes the set of binary strings of the same length as the output of $\text{obf}_{\text{on}}(\text{st}, C^j)$ algorithm.

Remark 18. It is shown in [19, 65] that there is no pPRIO scheme satisfying the above security for all general samplers. Therefore, when we use the security of pPRIO, we invoke the security *with respect to a specific sampler class* that is induced by the respective applications. For simplicity, we sometimes will treat as if there was pPRIO that is secure for all the samplers.

Theorem 5.14 ([14]). *Assuming LWE and evasive LWE assumptions, there exists a secure pPRIO scheme satisfying $|\text{Obf}_{\text{off}}| = \text{poly}(L, \lambda)$, $|\text{Obf}_{\text{on}}| = \text{poly}(L, \lambda)$ where $(\text{Obf}_{\text{off}}, \text{Obf}_{\text{on}}) \leftarrow \text{Obf}(1^\lambda, C)$ for circuit $C : [N] \rightarrow [M]$ whose size is bounded by $L = L(\lambda)$.*

5.3 FUNCTIONAL ENCRYPTION FOR PSEUDORANDOM FUNCTIONALITIES

5.3.1 Definition

In this section we give the definitions for functional encryption for pseudorandom functionalities. Consider a function family $\{\mathcal{F}_{\text{prm}} = \{f : \mathcal{X}_{\text{prm}} \rightarrow \mathcal{Y}_{\text{prm}}\}\}_{\text{prm}}$ for a parameter $\text{prm} = \text{prm}(\lambda)$. Each function $f \in \mathcal{F}_{\text{prm}}$ takes as input a string $x \in \mathcal{X}_{\text{prm}}$ and outputs $f(x) \in \mathcal{Y}_{\text{prm}}$.

Syntax. A functional encryption scheme prFE for pseudorandom functionalities \mathcal{F}_{prm} consists of four polynomial time algorithms (Setup , KeyGen , Enc , Dec) defined as follows.

$\text{Setup}(1^\lambda, \text{prm}) \rightarrow (\text{mpk}, \text{msk})$. The setup algorithm takes as input the security parameter λ and a parameter prm and outputs a master public key mpk and a master secret key msk ⁴.

$\text{KeyGen}(\text{msk}, f) \rightarrow \text{sk}_f$. The key generation algorithm takes as input the master

⁴We assume w.l.o.g that msk includes mpk .

secret key msk and a function $f \in \mathcal{F}_{\text{prm}}$ and it outputs a functional secret key sk_f .

$\text{Enc}(\text{mpk}, x) \rightarrow \text{ct}$. The encryption algorithm takes as input the master public key mpk and an input $x \in \mathcal{X}_{\text{prm}}$ and outputs a ciphertext $\text{ct} \in \mathcal{CT}$, where \mathcal{CT} is the ciphertext space.

$\text{Dec}(\text{mpk}, \text{sk}_f, f, \text{ct}) \rightarrow y$. The decryption algorithm takes as input the master public key mpk , secret key sk_f , function f and a ciphertext ct and outputs $y \in \mathcal{Y}_{\text{prm}}$.

Definition 5.12 (Correctness). A prFE scheme is said to satisfy *perfect* correctness if for all prm , any input $x \in \mathcal{X}_{\text{prm}}$ and function $f \in \mathcal{F}_{\text{prm}}$, we have

$$\Pr \left[\begin{array}{l} (\text{mpk}, \text{msk}) \leftarrow \text{Setup}(1^\lambda, \text{prm}), \text{sk}_f \leftarrow \text{KeyGen}(\text{msk}, f), \\ \text{Dec}(\text{mpk}, \text{sk}_f, f, \text{Enc}(\text{mpk}, x)) = f(x) \end{array} \right] = 1.$$

We define our security notion next. At a high level, our notion says that so long as the output of the functionality is pseudorandom, the ciphertext is pseudorandom. For notational brevity, we denote this by prCT security.

Definition 5.13 (prCT Security). For a prFE scheme for function family $\{\mathcal{F}_{\text{prm}} = \{f : \mathcal{X}_{\text{prm}} \rightarrow \mathcal{Y}_{\text{prm}}\}\}_{\text{prm}}$, parameter $\text{prm} = \text{prm}(\lambda)$, let Samp be a PPT algorithm that on input 1^λ , outputs

$$(f_1, \dots, f_{Q_{\text{key}}}, x_1, \dots, x_{Q_{\text{msg}}}, \text{aux} \in \{0, 1\}^*)$$

where Q_{key} is the number of key queries, Q_{msg} is the number of message queries, and $f_i \in \mathcal{F}_{\text{prm}}$, $x_j \in \mathcal{X}_{\text{prm}}$ for all $i \in [Q_{\text{key}}]$, $j \in [Q_{\text{msg}}]$.

We define the following advantage functions:

$$\begin{aligned} \text{Adv}_{\mathcal{A}_0}^{\text{PRE}}(\lambda) &\stackrel{\text{def}}{=} \Pr[\mathcal{A}_0(\text{aux}, \{f_i, f_i(x_j)\}_{i \in [Q_{\text{key}}], j \in [Q_{\text{msg}}]}) = 1] \\ &\quad - \Pr[\mathcal{A}_0(\text{aux}, \{f_i, \Delta_{i,j} \leftarrow \mathcal{Y}_{\text{prm}}\}_{i \in [Q_{\text{key}}], j \in [Q_{\text{msg}}]}) = 1] \end{aligned} \tag{5.9}$$

$$\begin{aligned} \text{Adv}_{\mathcal{A}_1}^{\text{POST}}(\lambda) \stackrel{\text{def}}{=} & \Pr[\mathcal{A}_1(\text{mpk}, \text{aux}, \{f_i, \text{Enc}(\text{mpk}, x_j), \text{sk}_{f_i}\}_{i \in [Q_{\text{key}}], j \in [Q_{\text{msg}}]}) = 1] \\ & - \Pr[\mathcal{A}_1(\text{mpk}, \text{aux}, \{f_i, \delta_j \leftarrow \mathcal{CT}, \text{sk}_{f_i}\}_{i \in [Q_{\text{key}}], j \in [Q_{\text{msg}}]}) = 1] \end{aligned} \quad (5.10)$$

where $(f_1, \dots, f_{Q_{\text{key}}}, x_1, \dots, x_{Q_{\text{msg}}}, \text{aux} \in \{0, 1\}^*) \leftarrow \text{Samp}(1^\lambda)$, $(\text{mpk}, \text{msk}) \leftarrow \text{Setup}(1^\lambda, \text{prm})$ and \mathcal{CT} is the ciphertext space. We say that a prFE scheme for function family \mathcal{F}_{prm} satisfies prCT security with respect to the sampler class \mathcal{SC} if for every PPT sampler $\text{Samp} \in \mathcal{SC}$ there exists a polynomial $Q(\cdot)$ such that for every PPT adversary \mathcal{A}_1 , there exists another PPT \mathcal{A}_0 such that

$$\text{Adv}_{\mathcal{A}_0}^{\text{PRE}}(\lambda) \geq \text{Adv}_{\mathcal{A}_1}^{\text{POST}}(\lambda)/Q(\lambda) - \text{negl}(\lambda) \quad (5.11)$$

and $\text{Time}(\mathcal{A}_0) \leq \text{Time}(\mathcal{A}_1) \cdot Q(\lambda)$.

Remark 19. It is shown in [19] that there is no prFE that satisfies prCT security for all general samplers. Therefore, when we use the security of prFE, we invoke the security *with respect to a specific sampler class* that is induced by the respective applications. For simplicity, we sometimes will treat as if there was prFE that is secure for all the samplers.

Remark 20. We note that the above security definition is in the multi-challenge flavor. One may wonder whether single-challenge version of the definition where $Q_{\text{msg}} = 1$ implies the multi-challenge version or not. However, unlike the standard security notions for public key primitives (e.g., indistinguishability security for functional encryption [91]), this does not seem to be the case. This is because the standard hybrid argument to prove the multi-challenge security from the single challenge security where we replace the ciphertext to be random one-by-one fails. To see why, recall that in these hybrids, we simulate some of the ciphertexts honestly, while we try to change a particular honest ciphertext to be a random one. To generate honest ciphertexts, we may have to know the corresponding plaintexts. However, this may ruin the precondition for invoking the single challenge security for the target ciphertext, since knowing some of the inputs (say, x_1) may make the output (say, $f_1(x_2)$) not pseudorandom any more when the inputs are

correlated with each other.

Definition 5.14 (Compactness). A prFE scheme is said to be compact if for any input message $x \in \mathcal{X}$, the running time of the encryption algorithm is polynomial in the security parameter and the size of x . In particular, it does not depend on the circuit description size or the output length of any function f supported by the scheme.

5.3.2 Construction

In this section, we provide our construction of a functional encryption scheme for pseudorandom functionalities for function family $\mathcal{F}_{L(\lambda), \ell(\lambda), \text{dep}(\lambda)} = \{f : \{0, 1\}^L \rightarrow \{0, 1\}^\ell\}$, where the depth of a function $f \in \mathcal{F}$ is at most $\text{dep}(\lambda) = \text{poly}(\lambda)$. We denote the information of the parameters representing the supported class of the circuits by $\text{prm} = (1^{L(\lambda)}, 1^{\ell(\lambda)}, 1^{\text{dep}(\lambda)})$.

Ingredients. Our construction needs a pseudorandom function $\text{PRF} : \{0, 1\}^\lambda \times \{0, 1\}^\lambda \rightarrow [-q/4 + B, q/4 - B]^{1 \times \ell}$ that can be evaluated by a circuit of depth at most $\text{dep}(\lambda) = \text{poly}(\lambda)$. Here B is chosen to be exponentially smaller than $q/4$. We note that for our choice of B the statistical distance between the uniform distribution over $[-q/4, q/4]$ and $[-q/4 + B, q/4 - B]$ is negligible.

$\text{Setup}(1^\lambda, \text{prm}) \rightarrow (\text{mpk}, \text{msk})$. The setup algorithm parses $\text{prm} = (1^{L(\lambda)}, 1^{\ell(\lambda)}, 1^{\text{dep}(\lambda)})$ and does the following.

- Sample appropriate parameters $q, n, m, \tau, \sigma, \sigma_{\mathbf{B}}, B$ and M such that M divides q , as in Equation (5.12)⁵.
- Samples appropriate parameters as in Equation (5.12).
- Set $L_X = m(\lambda + L)(n + 1)\lceil \log q \rceil$, sample $\mathbf{A}_{\text{att}} \leftarrow \mathbb{Z}_q^{(n+1) \times (L_X+1)m}$ and $(\mathbf{B}, \mathbf{B}_\tau^{-1}) \leftarrow \text{TrapGen}(1^{n+1}, 1^{mw}, q)$, where $w \in O(\log q)$.
- Output $\text{mpk} := (\mathbf{A}_{\text{att}}, \mathbf{B}, M)$ and $\text{msk} := \mathbf{B}_\tau^{-1}$.

⁵We assume these parameters to be part of the mpk.

$\text{KeyGen}(\text{msk}, f) \rightarrow \text{sk}_f$. The key generation algorithm parses $\text{msk} = \mathbf{B}_\tau^{-1}$ and does the following.

- Sample $\mathbf{r} \leftarrow \{0, 1\}^\lambda$ and define function $F = F[f, \mathbf{r}]$ with f, \mathbf{r} hardwired as follows⁶:
On input (\mathbf{x}, sd) , compute and output $f(\mathbf{x}) \lfloor q/2 \rfloor + \text{PRF}(\text{sd}, \mathbf{r}) \in \mathbb{Z}_q^{1 \times \ell}$.
- Parse $F[f, \mathbf{r}](\mathbf{x}, \text{sd}) = M \cdot f_{\text{high}}(\mathbf{x}, \text{sd}) + f_{\text{low}}(\mathbf{x}, \text{sd})$, where $f_{\text{high}}(\mathbf{x}, \text{sd}) \in [0, q/M]^\ell$ and $f_{\text{low}}(\mathbf{x}, \text{sd}) \in [0, M-1]^\ell$. Using the fact that the PRF and $f(\mathbf{x})$ can be computed by a circuit of depth at most $\text{dep}(\lambda) = \text{poly}(\lambda)$, the function $F[f, \mathbf{r}]$ can be computed by a circuit of depth at most $d = \text{poly}(\text{dep})$.
- Define functions $F_{\text{high}} := M \cdot f_{\text{high}}$ and $F_{\text{low}} := M \cdot f_{\text{low}}$, which on input (\mathbf{x}, sd) outputs $M \cdot f_{\text{high}}(\mathbf{x}, \text{sd})$ and $M \cdot f_{\text{low}}(\mathbf{x}, \text{sd})$, respectively. We note that these functions can be computed by a circuit of depth at most $d = \text{poly}(\text{dep})$.
- Define $\text{VEval}_{\text{high}} = \text{MakeVEvalCkt}(n, m, q, F_{\text{high}})$ and $\text{VEval}_{\text{low}} = \text{MakeVEvalCkt}(n, m, q, F_{\text{low}})$. From Theorem 4.5, the depth of $\text{VEval}_{\text{high}}$ and $\text{VEval}_{\text{low}}$ is bounded by $(dO(\log m \log \log q) + O(\log^2 \log q))$.
- Compute $\mathbf{H}_{\mathbf{A}_{\text{att}}}^{\text{F}_{\text{high}}} = \text{MEvalC}(\mathbf{A}_{\text{att}}, \text{VEval}_{\text{high}})$, $\mathbf{H}_{\mathbf{A}_{\text{att}}}^{\text{F}_{\text{low}}} = \text{MEvalC}(\mathbf{A}_{\text{att}}, \text{VEval}_{\text{low}}) \in \mathbb{Z}_q^{(L_X+1)m \times \ell}$.
- Compute $\mathbf{A}_{\text{high}} = \mathbf{A}_{\text{att}} \cdot \mathbf{H}_{\mathbf{A}_{\text{att}}}^{\text{F}_{\text{high}}}$ and $\mathbf{A}_{\text{low}} = \mathbf{A}_{\text{att}} \cdot \mathbf{H}_{\mathbf{A}_{\text{att}}}^{\text{F}_{\text{low}}}$.
- Compute
$$\mathbf{A}_F = M \cdot \left\lfloor \frac{\mathbf{A}_{\text{high}}}{M} \right\rfloor + \left\lfloor \frac{\mathbf{A}_{\text{low}}}{M} \right\rfloor$$
and sample $\mathbf{K} \leftarrow \mathbf{B}_\tau^{-1}(\mathbf{A}_F)$.
- Output $\text{sk}_f = (\mathbf{K}, \mathbf{r})$.

$\text{Enc}(\text{mpk}, \mathbf{x}) \rightarrow \text{ct}$. The encryption parse $\text{mpk} = (\mathbf{A}_{\text{att}}, \mathbf{B}, M)$ algorithm does the following.

- Sample $\bar{\mathbf{s}} \leftarrow \mathcal{D}_{\mathbb{Z}, \sigma_s}^n$ and set $\mathbf{s} = (\bar{\mathbf{s}}^\top, -1)^\top$.
- Sample $\mathbf{e}_\mathbf{B} \leftarrow \mathcal{D}_{\mathbb{Z}, \sigma_\mathbf{B}}^{mw}$ and compute $\mathbf{c}_\mathbf{B}^\top := \mathbf{s}^\top \mathbf{B} + \mathbf{e}_\mathbf{B}^\top$.

⁶The circuit representation of the function F is the universal circuit with F hardwired.

- Sample $\mathbf{sd} \leftarrow \{0, 1\}^\lambda$, $\bar{\mathbf{A}}_{\text{fhe}} \leftarrow \mathbb{Z}_q^{n \times m}$, $\mathbf{e}_{\text{fhe}} \leftarrow \mathcal{D}_{\mathbb{Z}, \sigma}^m$, $\mathbf{R} \leftarrow \{0, 1\}^{m \times m(\lambda+L)}$ and compute a GSW encryption as follows.

$$\mathbf{A}_{\text{fhe}} := \begin{pmatrix} \bar{\mathbf{A}}_{\text{fhe}} \\ \bar{\mathbf{s}}^\top \bar{\mathbf{A}}_{\text{fhe}} + \mathbf{e}_{\text{fhe}}^\top \end{pmatrix}, \quad \mathbf{X} = \mathbf{A}_{\text{fhe}} \mathbf{R} - (\mathbf{x}, \mathbf{sd}) \otimes \mathbf{G} \in \mathbb{Z}_q^{(n+1) \times m(\lambda+L)}.$$

Let $L_X = m(\lambda + L)(n + 1) \lceil \log q \rceil$ be the bit length of \mathbf{X} .

- Compute a BGG^+ encoding as follows.

$$\mathbf{e}_{\text{att}} \leftarrow \mathcal{D}_{\mathbb{Z}, \sigma}^{(L_X+1)m}, \quad \mathbf{c}_{\text{att}}^\top := \mathbf{s}^\top (\mathbf{A}_{\text{att}} - \text{bits}(1, \mathbf{X}) \otimes \mathbf{G}) + \mathbf{e}_{\text{att}}^\top.$$

- Output $\text{ct} = (\mathbf{c}_B, \mathbf{c}_{\text{att}}, \mathbf{X})$.

$\text{Dec}(\text{mpk}, \text{sk}_f, f, \text{ct}) \rightarrow \mathbf{y}$. The decryption algorithm does the following.

- Parse $\text{mpk} = (\mathbf{A}_{\text{att}}, \mathbf{B}, M)$, $\text{sk}_f = (\mathbf{K}, \mathbf{r})$ and $\text{ct} = (\mathbf{c}_B, \mathbf{c}_{\text{att}}, \mathbf{X})$.
- Compute $\mathbf{H}_{\mathbf{A}_{\text{att}}, \mathbf{X}}^{\text{Fhigh}} = \text{MEvalCX}(\mathbf{A}_{\text{att}}, \text{VEval}_{\text{high}}, \mathbf{X})$ and $\mathbf{H}_{\mathbf{A}_{\text{att}}, \mathbf{X}}^{\text{Flow}} = \text{MEvalCX}(\mathbf{A}_{\text{att}}, \text{VEval}_{\text{low}}, \mathbf{X})$ for circuits $\text{VEval}_{\text{high}}$ and $\text{VEval}_{\text{low}}$ as defined in KeyGen algorithm.
- Compute
$$\mathbf{z} := \mathbf{c}_B^\top \cdot \mathbf{K} - \left(M \cdot \left\lfloor \frac{\mathbf{c}_{\text{att}}^\top \cdot \mathbf{H}_{\mathbf{A}_{\text{att}}, \mathbf{X}}^{\text{Fhigh}}}{M} \right\rfloor + \left\lfloor \frac{\mathbf{c}_{\text{att}}^\top \cdot \mathbf{H}_{\mathbf{A}_{\text{att}}, \mathbf{X}}^{\text{Flow}}}{M} \right\rfloor \right).$$
- For $i \in [\ell]$, set $y_i = 0$ if $z_i \in [-q/4, q/4]$ and $y_i = 1$ otherwise, where z_i is the i -th coordinate of \mathbf{z} .
- Output $\mathbf{y} = (y_1, \dots, y_\ell)$.

Parameters. We set our parameters as follows.

$$\begin{aligned} \beta &= 2^{O(\text{dep} \cdot \log \lambda)}, \quad q = 2^{12\lambda} \beta, \quad M = 2^{4\lambda} \beta, \quad n = \text{poly}(\lambda, \text{dep}), \quad m = O(n \log q), \quad B = 2^{10\lambda} \beta, \\ \tau &= O\left(\sqrt{(n+1) \log q}\right) \quad \sigma_s = \sigma = 2^{2\lambda}, \quad \sigma_B = 2^{9\lambda} \beta, \quad \sigma_1 = 2^{8\lambda + O(1)} \beta / \text{poly}(\lambda). \end{aligned} \quad (5.12)$$

Efficiency. Using the above set parameters, we have

$$|\text{mpk}| = L \cdot \text{poly}(\text{dep}, \lambda), \quad |\text{sk}_f| = \ell \cdot \text{poly}(\text{dep}, \lambda), \quad |\text{ct}| = L \cdot \text{poly}(\text{dep}, \lambda).$$

Correctness We analyze the correctness of our scheme below.

– First, we note that

$$\begin{aligned} \mathbf{c}_{\text{att}}^\top \cdot \mathbf{H}_{\mathbf{A}_{\text{att}}, \mathbf{X}}^{\text{Fhigh}} &= (\mathbf{s}^\top (\mathbf{A}_{\text{att}} - \text{bits}(1, \mathbf{X}) \otimes \mathbf{G}) + \mathbf{e}_{\text{att}}^\top) \mathbf{H}_{\mathbf{A}_{\text{att}}, \mathbf{X}}^{\text{Fhigh}} \\ &= \mathbf{s}^\top \mathbf{A}_{\text{att}} \mathbf{H}_{\mathbf{A}_{\text{att}}}^{\text{Fhigh}} - \mathbf{s}^\top \text{VEval}_{\text{high}}(\text{bits}(\mathbf{X})) + \mathbf{e}_{\text{att}}^\top \mathbf{H}_{\mathbf{A}_{\text{att}}, \mathbf{X}}^{\text{Fhigh}} \\ &= \mathbf{s}^\top \mathbf{A}_{\text{high}} - F_{\text{high}}(\mathbf{x}, \text{sd}) + \mathbf{e}_{\text{fhe}}^\top \mathbf{R}_{\text{high}} + \mathbf{e}_{\text{att}}^\top \mathbf{H}_{\mathbf{A}_{\text{att}}, \mathbf{X}}^{\text{Fhigh}} \\ \Rightarrow \left\lfloor \frac{\mathbf{c}_{\text{att}}^\top \cdot \mathbf{H}_{\mathbf{A}_{\text{att}}, \mathbf{X}}^{\text{Fhigh}}}{M} \right\rfloor &= \left\lfloor \frac{\mathbf{s}^\top \mathbf{A}_{\text{high}} - M \cdot f_{\text{high}}(\mathbf{x}, \text{sd}) + \mathbf{e}_{\text{fhe}}^\top \mathbf{R}_{\text{high}} + \mathbf{e}_{\text{att}}^\top \mathbf{H}_{\mathbf{A}_{\text{att}}, \mathbf{X}}^{\text{Fhigh}}}{M} \right\rfloor \\ &= \left\lfloor \frac{\mathbf{s}^\top \mathbf{A}_{\text{high}} + \mathbf{e}_{\text{fhe}}^\top \mathbf{R}_{\text{high}} + \mathbf{e}_{\text{att}}^\top \mathbf{H}_{\mathbf{A}_{\text{att}}, \mathbf{X}}^{\text{Fhigh}}}{M} \right\rfloor - f_{\text{high}}(\mathbf{x}, \text{sd}) \quad (5.13) \end{aligned}$$

where $\text{VEval}_{\text{high}}(\text{bits}(\mathbf{X})) = \mathbf{A}_{\text{fhe}} \mathbf{R}_{\text{high}} - \begin{pmatrix} \mathbf{0}_{n \times \ell} \\ F_{\text{high}}(\mathbf{x}, \text{sd}) \end{pmatrix}$. Using Theorem 4.5, we have

$$\begin{aligned} \|\mathbf{R}_{\text{high}}^\top\| &\leq (m+2)^d \lceil \log q \rceil \cdot m = (m+2)^d \lceil \log q \rceil \cdot 3(n+1) \lceil \log q \rceil \\ &\leq (m+2)^d O(\log q) \leq \beta. \end{aligned}$$

and using the depth bound from Section 4.2.8,

$$\left\| \left(\mathbf{H}_{\mathbf{A}_{\text{att}}, \mathbf{X}}^{\text{Fhigh}} \right)^\top \right\| \leq (m+2)^{d_{\text{VEval}_{\text{high}}}} \lceil \log q \rceil \leq 2^{d \cdot O(\log \lambda)} \leq \beta$$

where $d_{\text{VEval}_{\text{high}}}$ denotes the depth of the circuit $\text{VEval}_{\text{high}}$. So we have

$$\left\| \mathbf{e}_{\text{fhe}}^\top \mathbf{R}_{\text{high}} + \mathbf{e}_{\text{att}}^\top \mathbf{H}_{\mathbf{A}_{\text{att}}, \mathbf{X}}^{\text{Fhigh}} \right\| \leq 2^{2\lambda+1} \sqrt{\lambda} \beta \leq 2^{3\lambda} \beta < M. \text{ Using this in Equation (5.13),}$$

$$\begin{aligned} \left\lfloor \frac{\mathbf{c}_{\text{att}}^\top \cdot \mathbf{H}_{\mathbf{A}_{\text{att}}, \mathbf{X}}^{\text{Fhigh}}}{M} \right\rfloor &= \left\lfloor \frac{\mathbf{s}^\top \mathbf{A}_{\text{high}}}{M} \right\rfloor - f_{\text{high}}(\mathbf{x}, \text{sd}) + \mathbf{err}_{\text{high}} \\ &= \mathbf{s}^\top \left\lfloor \frac{\mathbf{A}_{\text{high}}}{M} \right\rfloor + \mathbf{e}_{\text{s,high}}^\top - f_{\text{high}}(\mathbf{x}, \text{sd}) + \mathbf{err}_{\text{high}}^\top \quad (5.14) \end{aligned}$$

where $\mathbf{err}_{\text{high}}^\top \in \{0, 1\}^\ell$, is the rounding error which is 1 if $\left\| (\mathbf{s}^\top \mathbf{A}_{\text{high}})^\top + \mathbf{e}_{\text{fhe}}^\top \mathbf{R}_{\text{high}} + \mathbf{e}_{\text{att}}^\top \mathbf{H}_{\mathbf{A}_{\text{att}}, \mathbf{X}}^{\text{Fhigh}} \right\| \geq M$ and 0 otherwise, and $\|\mathbf{e}_{\text{s,high}}\| \leq (n+1) \cdot \|\mathbf{s}\|$. To see the latter, we use the fact that

$$\begin{aligned} \lfloor \mathbf{s}^\top X \rfloor - \mathbf{s}^\top \lfloor X \rfloor &= \lfloor \mathbf{s}^\top X - \mathbf{s}^\top \lfloor X \rfloor \rfloor = \lfloor \mathbf{s}^\top (X - \lfloor X \rfloor) \rfloor, \text{ where } X - \lfloor X \rfloor < 1. \text{ So} \\ \mathbf{e}_{\mathbf{s},\text{high}}^\top &= \left\lfloor \mathbf{s}^\top \left(\frac{\mathbf{A}_{\text{high}}}{M} - \left\lfloor \frac{\mathbf{A}_{\text{high}}}{M} \right\rfloor \right) \right\rfloor \quad \text{and} \\ \|\mathbf{e}_{\mathbf{s},\text{high}}\| &\leq \|\mathbf{s}\| \left\| \left(\frac{\mathbf{A}_{\text{high}}}{M} - \left\lfloor \frac{\mathbf{A}_{\text{high}}}{M} \right\rfloor \right)^\top \right\| < (n+1) \|\mathbf{s}\|. \end{aligned}$$

Using a similar analysis as to obtain Equation (5.14), we get

$$\left\lfloor \frac{\mathbf{c}_{\text{att}}^\top \cdot \mathbf{H}_{\mathbf{A}_{\text{att}},\mathbf{X}}^{\text{F}_{\text{low}}}}{M} \right\rfloor = \mathbf{s}^\top \left\lfloor \frac{\mathbf{A}_{\text{low}}}{M} \right\rfloor + \mathbf{e}_{\mathbf{s},\text{low}}^\top - f_{\text{low}}(\mathbf{x}, \text{sd}) + \mathbf{err}_{\text{low}}^\top \quad (5.15)$$

where $\mathbf{err}_{\text{low}}^\top \in \{0, 1\}^\ell$ and $\|\mathbf{e}_{\mathbf{s},\text{low}}\| \leq (n+1) \cdot \|\mathbf{s}\|$. Using Equation (5.14) and Equation (5.15), we get

$$\begin{aligned} &M \cdot \left\lfloor \frac{\mathbf{c}_{\text{att}}^\top \cdot \mathbf{H}_{\mathbf{A}_{\text{att}},\mathbf{X}}^{\text{F}_{\text{high}}}}{M} \right\rfloor + \left\lfloor \frac{\mathbf{c}_{\text{att}}^\top \cdot \mathbf{H}_{\mathbf{A}_{\text{att}},\mathbf{X}}^{\text{F}_{\text{low}}}}{M} \right\rfloor \\ &= \mathbf{s}^\top \left(M \cdot \left\lfloor \frac{\mathbf{A}_{\text{high}}}{M} \right\rfloor + \left\lfloor \frac{\mathbf{A}_{\text{low}}}{M} \right\rfloor \right) - (M \cdot f_{\text{high}}(\mathbf{x}, \text{sd}) + f_{\text{low}}(\mathbf{x}, \text{sd})) + \mathbf{err} \\ &= \mathbf{s}^\top \left(M \cdot \left\lfloor \frac{\mathbf{A}_{\text{high}}}{M} \right\rfloor + \left\lfloor \frac{\mathbf{A}_{\text{low}}}{M} \right\rfloor \right) - F[f, \mathbf{r}](\mathbf{x}, \text{sd}) + \mathbf{err} \end{aligned} \quad (5.16)$$

where

$$\begin{aligned} \mathbf{err} &= M \cdot \mathbf{e}_{\mathbf{s},\text{high}}^\top + \mathbf{e}_{\mathbf{s},\text{low}}^\top + M \cdot \mathbf{err}_{\text{high}} + \mathbf{err}_{\text{low}} \\ &= M \cdot \left(\mathbf{s}^\top \left\lfloor \frac{\mathbf{A}_{\text{high}}}{M} \right\rfloor - \left\lfloor \mathbf{s}^\top \frac{\mathbf{A}_{\text{high}}}{M} \right\rfloor \right) + \left(\mathbf{s}^\top \left\lfloor \frac{\mathbf{A}_{\text{low}}}{M} \right\rfloor - \left\lfloor \mathbf{s}^\top \frac{\mathbf{A}_{\text{low}}}{M} \right\rfloor \right) + M \cdot \mathbf{err}_{\text{high}} + \mathbf{err}_{\text{low}} \end{aligned} \quad (5.17)$$

where $\mathbf{err}_{\text{high}}, \mathbf{err}_{\text{low}} \in \{0, 1\}^\ell$ are rounding errors and matrices $\mathbf{A}_{\text{high}}, \mathbf{A}_{\text{low}}$ are publicly computable matrices and

$$\begin{aligned} \|\mathbf{err}\| &\leq M \cdot ((n+1) \cdot \|\mathbf{s}\| + 1) + (n+1) \cdot \|\mathbf{s}\| + 1 \leq 2M \cdot ((n+1) \cdot \|\mathbf{s}\| + 1) \\ &= 2^{4\lambda+1} \beta \left((n+1) \cdot 2^{2\lambda+1} \sqrt{\lambda} \right) \leq 2^{7\lambda} \beta \end{aligned}$$

– Next, we note that

$$\mathbf{c}_{\mathbf{B}}^\top \cdot \mathbf{K} = \mathbf{s}^\top \left(M \cdot \left\lfloor \frac{\mathbf{A}_{\text{high}}}{M} \right\rfloor + \left\lfloor \frac{\mathbf{A}_{\text{low}}}{M} \right\rfloor \right) + \mathbf{c}_{\mathbf{B}}^\top \cdot \mathbf{K}. \quad (5.18)$$

where $\|(\mathbf{c}_{\mathbf{B}}^\top \cdot \mathbf{K})^\top\| \leq 2^{9\lambda} \beta \sqrt{\lambda} \cdot \tau$ from our parameter setting.

– Using Equations (5.16) and (5.18), we get

$$\begin{aligned}
\mathbf{z} &= \mathbf{c}_B^\top \cdot \mathbf{K} - \left(M \cdot \left[\frac{\mathbf{c}_{\text{att}}^\top \cdot \mathbf{H}_{\mathbf{A}_{\text{att}}, \mathbf{X}}^{\text{F}_{\text{high}}}}{M} \right] + \left[\frac{\mathbf{c}_{\text{att}}^\top \cdot \mathbf{H}_{\mathbf{A}_{\text{att}}, \mathbf{X}}^{\text{F}_{\text{low}}}}{M} \right] \right) \\
&= F[f, \mathbf{r}](\mathbf{x}, \mathbf{sd}) + \mathbf{e}_B^\top \cdot \mathbf{K} - \mathbf{err} \\
&= f(\mathbf{x}) \lfloor q/2 \rfloor + \text{PRF}(\mathbf{sd}, \mathbf{r}) + \mathbf{e}_B^\top \cdot \mathbf{K} - \mathbf{err}
\end{aligned}$$

where

$$\begin{aligned}
\|\text{PRF}(\mathbf{sd}, \mathbf{r}) + \mathbf{e}_B^\top \cdot \mathbf{K} - \mathbf{err}\| &\leq \|\text{PRF}(\mathbf{sd}, \mathbf{r})\| + 2^{9\lambda} \beta \sqrt{\lambda} \cdot \tau + 2^{7\lambda} \beta \\
&\leq \|\text{PRF}(\mathbf{sd}, \mathbf{r})\| + 2^{9\lambda+1} \beta \sqrt{\lambda} \cdot \tau < q/4 - B + B < q/4
\end{aligned}$$

Hence the last step of decryption outputs \mathbf{y} correctly with probability 1.

5.3.3 Security Proof for Pseudorandom Functionalities

Theorem 5.15. *Let $\mathcal{SC}_{\text{prFE}}$ be a sampler class for prFE. Assuming LWE (Assumption 2.5) and private coin Evasive LWE (Assumption 2.6) with respect to the sampler class that contains all $\text{Samp}_{\text{evs}}(1^\lambda)$ induced by $\text{Samp}_{\text{prFE}} \in \mathcal{SC}_{\text{prFE}}$ as defined in Figure 5.1, our prFE scheme satisfies prCT security with respect to $\mathcal{SC}_{\text{prFE}}$ as defined in Definition 5.13.*

Proof. Consider a sampler $\text{Samp}_{\text{prFE}}$ that generates the following:

1. **Key Queries.** It issues Q_{key} number of functions $f_1, \dots, f_{Q_{\text{key}}}$ for key queries.
2. **Ciphertext Queries.** It issues Q_{msg} ciphertext queries $\mathbf{x}_1, \dots, \mathbf{x}_{Q_{\text{msg}}}$.
3. **Auxiliary Information.** It outputs the auxiliary information $\text{aux}_{\mathcal{A}}$.

To prove the prCT security as per Definition 5.13, we show

$$\left(\begin{array}{l} \text{mpk} = (\mathbf{A}_{\text{att}}, \mathbf{B}, M), \text{aux}_{\mathcal{A}}, \mathbf{C}_B = \mathbf{S}\mathbf{B} + \mathbf{E}_B, \\ \{\mathbf{X}_j = \mathbf{A}_{\text{fhe}, j} \mathbf{R}_j - (\mathbf{x}_j, \mathbf{sd}_j) \otimes \mathbf{G}\}_{j \in [Q_{\text{msg}}]}, \\ \{\mathbf{c}_{\text{att}, j}^\top = \mathbf{s}_j^\top (\mathbf{A}_{\text{att}} - \text{bits}(1, \mathbf{X}_j) \otimes \mathbf{G}) + \mathbf{e}_{\text{att}, j}^\top\}_{j \in [Q_{\text{msg}}]}, \\ \{\mathbf{K}_k, \mathbf{r}_k\}_{k \in [Q_{\text{key}}]} \end{array} \right) \approx_c \left(\begin{array}{l} \text{mpk} = (\mathbf{A}_{\text{att}}, \mathbf{B}, M), \text{aux}_{\mathcal{A}}, \mathbf{C}_B \leftarrow \mathbb{Z}_q^{Q_{\text{msg}} \times m w}, \\ \{\mathbf{X}_j \leftarrow \mathbb{Z}_q^{(n+1) \times m(\lambda+L)}\}_{j \in [Q_{\text{msg}}]}, \\ \{\mathbf{c}_{\text{att}, j} \leftarrow \mathbb{Z}_q^{(L_X+1)m}\}_{j \in [Q_{\text{msg}}]}, \\ \{\mathbf{K}_k, \mathbf{r}_k\}_{k \in [Q_{\text{key}}]} \end{array} \right) \quad (5.19)$$

where

$$\mathbf{S} = \begin{pmatrix} \mathbf{s}_1^\top \\ \vdots \\ \mathbf{s}_{Q_{\text{msg}}}^\top \end{pmatrix}, \mathbf{E}_{\mathbf{B}} = \begin{pmatrix} \mathbf{e}_{\mathbf{B},1}^\top \\ \vdots \\ \mathbf{e}_{\mathbf{B},Q_{\text{msg}}}^\top \end{pmatrix},$$

$$(\mathbf{aux}_{\mathcal{A}}, \{f_k\}_{k \in [Q_{\text{key}}]}, \{\mathbf{x}_j\}_{j \in [Q_{\text{msg}}]}) \leftarrow \text{Samp}_{\text{prFE}}(1^\lambda)$$

and for $j \in [Q_{\text{msg}}]$, $\mathbf{s}_j, \mathbf{e}_{\mathbf{B},j}, \mathbf{A}_{\text{fhe},j}, \mathbf{R}_j, \mathbf{sd}_j, \mathbf{e}_{\text{att},j}$ are sampled as in the construction, for $k \in [Q_{\text{key}}]$, we have $\mathbf{r}_k \leftarrow \{0,1\}^\lambda$, $F_k = F[f_k, \mathbf{r}_k]$ and \mathbf{A}_{F_k} is as defined in the construction, and $\mathbf{K}_k = \mathbf{B}_\tau^{-1}(\mathbf{A}_{F_k})$

assuming we have

$$(1^\lambda, \mathbf{aux}_{\mathcal{A}}, \{f_k, f_k(x_j)\}_{j \in [Q_{\text{msg}}], k \in [Q_{\text{key}}]}) \approx_c (1^\lambda, \mathbf{aux}_{\mathcal{A}}, \{f_k, \Delta_{j,k} \leftarrow \{0,1\}^\ell\}_{j \in [Q_{\text{msg}}], k \in [Q_{\text{key}}]}). \quad (5.20)$$

We invoke evasive LWE assumption for a matrix \mathbf{B} with the private coin sampler Samp_{evs} that outputs $(\mathbf{S}, \mathbf{P}, \mathbf{aux} = (\mathbf{aux}_1, \mathbf{aux}_2))$ with private coin $\text{coins}_{\text{priv}}^{\text{Samp}_{\text{evs}}} = \{\mathbf{sd}_j, \mathbf{R}_j, \mathbf{e}_{\text{att},j}, \mathbf{A}_{\text{fhe},j}\}_{j \in [Q_{\text{msg}}]}$, defined as follows.

By Theorem 2.7, to prove Equation (5.19) assuming evasive LWE, it suffices to show

$$\left(\begin{array}{l} \mathbf{aux}_2, \mathbf{B}, \mathbf{C}_{\mathbf{B}} = \mathbf{SB} + \mathbf{E}_{\mathbf{B}}, \\ \{\mathbf{X}_j = \mathbf{A}_{\text{fhe},j} \mathbf{R}_j - (\mathbf{x}_j, \mathbf{sd}_j) \otimes \mathbf{G}\}_{j \in [Q_{\text{msg}}]}, \\ \{\mathbf{c}_{\text{att},j}^\top = \mathbf{s}_j^\top (\mathbf{A}_{\text{att}} - \text{bits}(1, \mathbf{X}_j) \otimes \mathbf{G}) + \mathbf{e}_{\text{att},j}^\top\}_{j \in [Q_{\text{msg}}]}, \\ \mathbf{C}_{\mathbf{P}} = \mathbf{SP} + \mathbf{E}_{\mathbf{P}} \end{array} \right) \approx_c \left(\begin{array}{l} \mathbf{aux}_2, \mathbf{B}, \mathbf{C}_{\mathbf{B}} \leftarrow \mathbb{Z}_q^{Q_{\text{msg}} \times mw}, \\ \{\mathbf{X}_j \leftarrow \mathbb{Z}_q^{(n+1) \times m(\lambda+L)}\}_{j \in [Q_{\text{msg}}]}, \\ \{\mathbf{c}_{\text{att},j} \leftarrow \mathbb{Z}_q^{(L_X+1)m}\}_{j \in [Q_{\text{msg}}]}, \\ \mathbf{C}_{\mathbf{P}} \leftarrow \mathbb{Z}_q^{Q_{\text{msg}} \times \ell \cdot Q_{\text{key}}} \end{array} \right) \quad (5.21)$$

Samp_{evs}(1^λ)

The sampler does the following.

- Runs the prFE sampler Samp_{prFE} to obtain $(\{f_k\}_{k \in [Q_{\text{key}}]}, \{\mathbf{x}_j\}_{j \in [Q_{\text{msg}}]}, \text{aux}_{\mathcal{A}})$ where $f_k : \{0, 1\}^L \rightarrow \{0, 1\}^\ell$, $\mathbf{x}_j \in \{0, 1\}^L$ and $\text{aux}_{\mathcal{A}} \in \{0, 1\}^*$.
- Set appropriate parameters as in Equation (5.12)^a.
- Samples $\text{sd}_j \leftarrow \{0, 1\}^\lambda$, $\bar{\mathbf{A}}_{\text{fhe},j} \leftarrow \mathbb{Z}_q^{n \times m}$, $\mathbf{e}_{\text{fhe},j} \leftarrow \mathcal{D}_{\mathbb{Z}, \sigma}^m$, $\mathbf{R}_j \leftarrow \{0, 1\}^{m \times m(\lambda+L)}$ and computes $\mathbf{X}_j = \mathbf{A}_{\text{fhe},j} \mathbf{R}_j - (\mathbf{x}_j, \text{sd}_j) \otimes \mathbf{G}$ for $j \in [Q_{\text{msg}}]$ where $\mathbf{A}_{\text{fhe},j} = \left(\bar{\mathbf{s}}_j^\top \bar{\mathbf{A}}_{\text{fhe},j} + \mathbf{e}_{\text{fhe},j}^\top \right)$ $\forall j \in [Q_{\text{msg}}]$.
- Samples $\bar{\mathbf{s}}_j \leftarrow \mathcal{D}_{\mathbb{Z}, \sigma_s}^n$, $\mathbf{e}_{\text{att},j} \leftarrow \mathcal{D}_{\mathbb{Z}, \sigma}^{(L_{\text{X}}+1)m}$, sets $\mathbf{s}_j = (\bar{\mathbf{s}}_j^\top, -1)^\top$ and computes $\mathbf{c}_{\text{att},j}^\top = \mathbf{s}_j^\top (\mathbf{A}_{\text{att}} - \text{bits}(1, \mathbf{X}_j) \otimes \mathbf{G}) + \mathbf{e}_{\text{att},j}^\top \forall j \in [Q_{\text{msg}}]$.
- Samples $\mathbf{r}_k \leftarrow \{0, 1\}^\lambda$, defines $F[f_k, \mathbf{r}_k]$ and computes \mathbf{A}_{F_k} , for $k \in [Q_{\text{key}}]$, as in the key generation algorithm.
- It outputs

$$\mathbf{S} = \begin{pmatrix} \mathbf{s}_1^\top \\ \vdots \\ \mathbf{s}_{Q_{\text{msg}}}^\top \end{pmatrix}, \quad \mathbf{P} = [\mathbf{A}_{F_1} \parallel \dots \parallel \mathbf{A}_{F_{Q_{\text{key}}}}]$$

$$\text{aux}_1 = \left(\{\mathbf{X}_j = \mathbf{A}_{\text{fhe},j} \mathbf{R}_j - (\mathbf{x}_j, \text{sd}_j) \otimes \mathbf{G}\}_{j \in [Q_{\text{msg}}]}, \{\mathbf{c}_{\text{att},j}^\top = \mathbf{s}_j^\top (\mathbf{A}_{\text{att}} - \text{bits}(1, \mathbf{X}_j) \otimes \mathbf{G}) + \mathbf{e}_{\text{att},j}^\top\}_{j \in [Q_{\text{msg}}]} \right),$$

$$\text{aux}_2 = (f_1, \dots, f_{Q_{\text{key}}}, \text{aux}_{\mathcal{A}}, \mathbf{r}_1, \dots, \mathbf{r}_{Q_{\text{key}}}, \mathbf{A}_{\text{att}}, M).$$

^aWe assume the parameters to be output as a part of aux_2 , even though we do not explicitly write so.

Figure 5.1: Description of the Sampler for Evasive LWE

where $\mathbf{E}_{\mathbf{P}} \leftarrow \mathcal{D}_{\mathbb{Z}, \sigma_1}^{Q_{\text{msg}} \times \ell \cdot Q_{\text{key}}}$. Using the representation

$$\mathbf{C}_{\mathbf{B}} = \begin{pmatrix} \mathbf{c}_{\mathbf{B},1}^\top = \mathbf{s}_1^\top \mathbf{B} + \mathbf{e}_{\mathbf{B},1}^\top \\ \vdots \\ \mathbf{c}_{\mathbf{B},Q_{\text{msg}}}^\top = \mathbf{s}_{Q_{\text{msg}}}^\top \mathbf{B} + \mathbf{e}_{\mathbf{B},Q_{\text{msg}}}^\top \end{pmatrix} = \{\mathbf{c}_{\mathbf{B},j}^\top\}_{j \in [Q_{\text{msg}}]},$$

$$\mathbf{C}_{\mathbf{P}} = \begin{pmatrix} \mathbf{c}_{\mathbf{P},1}^\top = \mathbf{s}_1^\top \mathbf{A}_{F_1} + \mathbf{e}_{\mathbf{P},1,1}^\top \parallel \dots \parallel \mathbf{s}_1^\top \mathbf{A}_{F_{Q_{\text{key}}}} + \mathbf{e}_{\mathbf{P},1,Q_{\text{key}}}^\top \\ \vdots \\ \mathbf{c}_{\mathbf{P},Q_{\text{msg}}}^\top = \mathbf{s}_{Q_{\text{msg}}}^\top \mathbf{A}_{F_1} + \mathbf{e}_{\mathbf{P},Q_{\text{msg},1}}^\top \parallel \dots \parallel \mathbf{s}_{Q_{\text{msg}}}^\top \mathbf{A}_{F_{Q_{\text{key}}}} + \mathbf{e}_{\mathbf{P},Q_{\text{msg}},Q_{\text{key}}}^\top \end{pmatrix} = \{\mathbf{c}_{\mathbf{P},j,k}^\top\}_{j \in [Q_{\text{msg}}], k \in [Q_{\text{key}}]},$$

we rewrite Equation (5.21) as follows.

$$\left(\begin{array}{l} \text{aux}_2, \mathbf{B}, \{\mathbf{c}_{\mathbf{B},j}^\top = \mathbf{s}_j^\top \mathbf{B} + \mathbf{e}_{\mathbf{B},j}^\top\}_{j \in [Q_{\text{msg}}]}, \\ \{\mathbf{X}_j = \mathbf{A}_{\text{fhe},j} \mathbf{R}_j - (\mathbf{x}_j, \mathbf{sd}_j) \otimes \mathbf{G}\}_{j \in [Q_{\text{msg}}]}, \\ \{\mathbf{c}_{\text{att},j}^\top = \mathbf{s}_j^\top (\mathbf{A}_{\text{att}} - \text{bits}(1, \mathbf{X}_j) \otimes \mathbf{G}) + \mathbf{e}_{\text{att},j}^\top\}_{j \in [Q_{\text{msg}}]}, \\ \{\mathbf{c}_{\mathbf{P},j,k}^\top = \mathbf{s}_j^\top \mathbf{A}_{\mathbf{F}k} + \mathbf{e}_{\mathbf{P},j,k}^\top\}_{j \in [Q_{\text{msg}}], k \in [Q_{\text{key}}]} \end{array} \right) \approx_c \left(\begin{array}{l} \text{aux}_2, \mathbf{B}, \{\mathbf{c}_{\mathbf{B},j} \leftarrow \mathbb{Z}_q^{mw}\}_{j \in [Q_{\text{msg}}]}, \\ \{\mathbf{X}_j \leftarrow \mathbb{Z}_q^{(n+1) \times m(\lambda+L)}\}_{j \in [Q_{\text{msg}}]}, \\ \{\mathbf{c}_{\text{att},j} \leftarrow \mathbb{Z}_q^{(L_X+1)m}\}_{j \in [Q_{\text{msg}}]}, \\ \{\mathbf{c}_{\mathbf{P},j,k} \leftarrow \mathbb{Z}_q^\ell\}_{j \in [Q_{\text{msg}}], k \in [Q_{\text{key}}]} \end{array} \right) \quad (5.22)$$

where $\mathbf{e}_{\mathbf{P},j,k} \leftarrow \mathcal{D}_{\mathbb{Z}, \sigma_1}^\ell$. Now, to prove Equation (5.19) it suffices to show Equation (5.22).

We prove Equation (5.22) via the following sequence of hybrids.

Hyb₀. This is L.H.S distribution of Equation (5.22).

Hyb₁. This hybrid is same as Hyb₀, except we compute $\mathbf{c}_{\mathbf{P},j,k}^\top$ as

$$\mathbf{c}_{\mathbf{P},j,k}^\top = M \cdot \left\lfloor \frac{\mathbf{c}_{\text{att},j}^\top \cdot \mathbf{H}_{\mathbf{A}_{\text{att}}, \mathbf{X}_j}^{\text{Fhigh},k}}{M} \right\rfloor + \left\lfloor \frac{\mathbf{c}_{\text{att},j}^\top \cdot \mathbf{H}_{\mathbf{A}_{\text{att}}, \mathbf{X}_j}^{\text{Flow},k}}{M} \right\rfloor + f_k(\mathbf{x}_j) \lfloor q/2 \rfloor + \text{PRF}(\mathbf{sd}_j, \mathbf{r}_k) + \mathbf{e}_{\mathbf{P},j,k}^\top$$

where $\mathbf{e}_{\mathbf{P},j,k} \leftarrow \mathcal{D}_{\mathbb{Z}, \sigma_1}^\ell$. We claim that Hyb₀ and Hyb₁ are statistically indistinguishable. To see this, we observe the following.

– From Equation (5.16) we note that

$$\begin{aligned} M \cdot \left\lfloor \frac{\mathbf{c}_{\text{att},j}^\top \cdot \mathbf{H}_{\mathbf{A}_{\text{att}}, \mathbf{X}_j}^{\text{Fhigh},k}}{M} \right\rfloor + \left\lfloor \frac{\mathbf{c}_{\text{att},j}^\top \cdot \mathbf{H}_{\mathbf{A}_{\text{att}}, \mathbf{X}_j}^{\text{Flow},k}}{M} \right\rfloor + f_k(\mathbf{x}_j) \lfloor q/2 \rfloor + \text{PRF}(\mathbf{sd}_j, \mathbf{r}_k) + \mathbf{e}_{\mathbf{P},j,k}^\top \\ = \mathbf{s}_j^\top \left(M \cdot \left\lfloor \frac{\mathbf{A}_{\text{high},k}}{M} \right\rfloor + \left\lfloor \frac{\mathbf{A}_{\text{low},k}}{M} \right\rfloor \right) + \mathbf{err}_{j,k} + \mathbf{e}_{\mathbf{P},j,k}^\top \\ = \mathbf{s}_j^\top \mathbf{A}_{\mathbf{F}k} + \mathbf{err}_{j,k} + \mathbf{e}_{\mathbf{P},j,k}^\top \end{aligned}$$

where $\|\mathbf{err}_{j,k}\| \leq 2^{7\lambda} \beta$.

– Next, we note that $\|\mathbf{err}_{j,k}\| \leq 2^{8\lambda+O(1)} \beta / \text{poly}(\lambda) = \chi_1 = \|\mathbf{e}_{\mathbf{P},j,k}\|$. Thus by noise flooding (Theorem 2.3) we have $\mathbf{e}_{\mathbf{P},j,k}^\top \approx_s \mathbf{err}_{j,k} + \mathbf{e}_{\mathbf{P},j,k}^\top$ with a statistical distance of $\text{poly}(\lambda) 2^{-\lambda}$.

From the above, we have

$$\Delta(\text{Hyb}_0, \text{Hyb}_1) = \frac{Q_{\text{key}} \cdot Q_{\text{msg}} \cdot \text{poly}(\lambda)}{2^\lambda}.$$

Thus, it suffices to show pseudorandomness of the following distribution given aux_2

$$\left(\begin{array}{l} \mathbf{B}, \{ \mathbf{c}_{\mathbf{B},j}^\top = \mathbf{s}_j^\top \mathbf{B} + \mathbf{e}_{\mathbf{B},j}^\top \}_{j \in [Q_{\text{msg}}]} \\ \{ \mathbf{X}_j = \mathbf{A}_{\text{fhe},j} \mathbf{R}_j - (\mathbf{x}_j, \text{sd}_j) \otimes \mathbf{G} \}_{j \in [Q_{\text{msg}}]}, \\ \{ \mathbf{c}_{\text{att},j}^\top = \mathbf{s}_j^\top (\mathbf{A}_{\text{att}} - \text{bits}(1, \mathbf{X}_j) \otimes \mathbf{G}) + \mathbf{e}_{\text{att},j}^\top \}_{j \in [Q_{\text{msg}}]}, \\ \{ \tilde{\mathbf{F}}_{j,k} = f_k(\mathbf{x}_j) \lfloor q/2 \rfloor + \text{PRF}(\text{sd}_j, \mathbf{r}_k) + \mathbf{e}_{\mathbf{P},j,k}^\top \}_{j \in [Q_{\text{msg}}], k \in [Q_{\text{key}}]} \end{array} \right)$$

Hyb₂. This hybrid is same as **Hyb₁** except that for all $j \in [Q_{\text{msg}}]$ we sample $\mathbf{c}_{\mathbf{B},j} \leftarrow \mathbb{Z}_q^{mw}$, $\mathbf{c}_{\text{att},j} \leftarrow \mathbb{Z}_q^{(L_X+1)m}$ and $\mathbf{A}_{\text{fhe},j} \leftarrow \mathbb{Z}_q^{(n+1) \times m}$, where $\mathbf{A}_{\text{fhe},j}$ is the fhe public key used to compute \mathbf{X}_j . We have $\text{Hyb}_1 \approx_c \text{Hyb}_2$ using LWE.

To prove this we consider sub-hybrids **Hyb_{1,i}** for $i \in [Q_{\text{msg}}]$, where in **Hyb_{1,i}** we sample $\mathbf{c}_{\mathbf{B},j} \leftarrow \mathbb{Z}_q^{mw}$, $\mathbf{c}_{\text{att},j} \leftarrow \mathbb{Z}_q^{(L_X+1)m}$ and $\mathbf{A}_{\text{fhe},j} \leftarrow \mathbb{Z}_q^{(n+1) \times m}$ for $1 \leq j \leq i$. We set $\text{Hyb}_1 = \text{Hyb}_{1,0}$ and $\text{Hyb}_2 = \text{Hyb}_{1,Q_{\text{msg}}}$. Next, we prove that for all $i \in [Q_{\text{msg}}]$, $\text{Hyb}_{1,i-1} \approx_c \text{Hyb}_{1,i}$ via the following claim.

Claim 5.16. $\text{Hyb}_{1,i-1} \approx_c \text{Hyb}_{1,i}$, for $i \in [Q_{\text{msg}}]$, assuming the security of LWE.

Proof. We show that if there exists an adversary \mathcal{A} who can distinguish between the two hybrids with non-negligible advantage, then there is a reduction \mathcal{B} that breaks LWE security with non-negligible advantage. The reduction is as follows.

1. The adversary \mathcal{A} sends the function queries $f_1, \dots, f_{Q_{\text{key}}}$, message queries $\mathbf{x}_1, \dots, \mathbf{x}_{Q_{\text{msg}}}$ and auxiliary input $\text{aux}_{\mathcal{A}}$ to the reduction.
2. \mathcal{B} initiates the LWE security game with the LWE challenger. The challenger sends $\mathbf{A}_{\text{lwe}} \in \mathbb{Z}_q^{n \times mw + m + (L_X+1)m}$ and $\mathbf{b} \in \mathbb{Z}_q^{mw + m + (L_X+1)m}$ to \mathcal{B} .
3. \mathcal{B} parses $\mathbf{A}_{\text{lwe}} = (\mathbf{B}', \hat{\mathbf{A}}_{\text{fhe}}, \mathbf{A}'_{\text{att}})$, where $\mathbf{B}' \in \mathbb{Z}_q^{n \times mw}$, $\hat{\mathbf{A}}_{\text{fhe}} \in \mathbb{Z}_q^{n \times m}$, $\mathbf{A}'_{\text{att}} \in \mathbb{Z}_q^{n \times (L_X+1)m}$ and $\mathbf{b}^\top = (\mathbf{b}_{\mathbf{B}}^\top, \mathbf{b}_{\text{fhe}}^\top, \mathbf{b}_{\text{att}}^\top)$. For $j \in [Q_{\text{msg}}]$, it computes

$\mathbf{c}_{\mathbf{B},j}$, $\mathbf{c}_{\text{att},j}$ and $\mathbf{A}_{\text{fhe},j}$ as follows.

- For $1 \leq j < i$: \mathcal{B} samples $\mathbf{c}_{\mathbf{B},j} \leftarrow \mathbb{Z}_q^{mw}$, $\mathbf{c}_{\text{att},j} \leftarrow \mathbb{Z}_q^{(L_X+1)m}$ and $\mathbf{A}_{\text{fhe},j} \leftarrow \mathbb{Z}_q^{(n+1) \times m}$.

- For $j = i$: \mathcal{B} does the following.

- Samples $\underline{\mathbf{b}} \leftarrow \mathbb{Z}_q^{mw}$ and sets $\mathbf{B} = \begin{pmatrix} \mathbf{B}' \\ \underline{\mathbf{b}}^\top \end{pmatrix}$ and $\mathbf{c}_{\mathbf{B},i}^\top := \mathbf{b}_{\mathbf{B}}^\top - \underline{\mathbf{b}}^\top$.
- Sets $\mathbf{A}_{\text{fhe},i} := \begin{pmatrix} \hat{\mathbf{A}}_{\text{fhe}} \\ \mathbf{b}_{\text{fhe}}^\top \end{pmatrix}$ and computes $\mathbf{X}_i = \mathbf{A}_{\text{fhe},i} \mathbf{R}_i - (\mathbf{x}_i, \text{sd}_i) \otimes \mathbf{G}$ as in the construction.
- Sets $\bar{\mathbf{A}}_{\text{att}} = \mathbf{A}'_{\text{att}} + \text{bits}(1, \mathbf{X}_i) \otimes \bar{\mathbf{G}}$, $\mathbf{A}_{\text{att}} = \begin{pmatrix} \bar{\mathbf{A}}_{\text{att}} \\ \underline{\mathbf{a}}_{\text{att}}^\top \end{pmatrix}$, where $\underline{\mathbf{a}}_{\text{att}} \leftarrow \mathbb{Z}_q^{(L_X+1)m}$, and $\mathbf{c}_{\text{att},i}^\top = \mathbf{b}_{\text{att}}^\top - (\underline{\mathbf{a}}_{\text{att}}^\top - \text{bits}(1, \mathbf{X}_i) \otimes \underline{\mathbf{G}})$, where $\bar{\mathbf{G}}$ and $\underline{\mathbf{G}}$ denotes the first n rows and $n+1$ -th row of the gadget matrix $\mathbf{G} \in \mathbb{Z}_q^{(n+1) \times m}$, respectively.

- For $j > i$: \mathcal{B} computes $\mathbf{c}_{\mathbf{B},j}^\top$, \mathbf{X}_j and $\mathbf{c}_{\text{att},j}^\top$ as in the construction, where

$$\mathbf{c}_{\text{att},j}^\top \text{ is computed using } \mathbf{A}_{\text{att}} = \begin{pmatrix} \bar{\mathbf{A}}_{\text{att}} \\ \underline{\mathbf{a}}_{\text{att}}^\top \end{pmatrix}.$$

4. \mathcal{B} sets $\text{aux}_2 = (f_1, \dots, f_{Q_{\text{key}}}, \text{aux}_{\mathcal{A}}, \mathbf{r}_1, \dots, \mathbf{r}_{Q_{\text{key}}}, \mathbf{A}_{\text{att}}, M)$ where $\mathbf{r}_k \leftarrow \{0, 1\}^\lambda$ and computes $\tilde{\mathbf{F}}_{j,k}$ as in Hyb_1 . It sends $(\text{aux}_2, \{\mathbf{c}_{\mathbf{B},j}^\top, \mathbf{X}_j, \mathbf{c}_{\text{att},j}^\top, \tilde{\mathbf{F}}_{j,k}\})$ to the adversary.

5. \mathcal{A} outputs a bit β' . \mathcal{B} forwards the bit β' to the LWE challenger.

We note that if the LWE challenger sent $\mathbf{b} = \bar{\mathbf{s}}\mathbf{A}_{\text{lwe}} + \mathbf{e}_{\text{lwe}}$, then \mathcal{B} simulated $\text{Hyb}_{1,i-1}$ with \mathcal{A} else if LWE challenger sent random $\mathbf{b} \leftarrow \mathbb{Z}_q^{mw+m+(L_X+1)m}$ then \mathcal{B} simulated $\text{Hyb}_{1,i}$ with \mathcal{A} .

To see the former case, we note that if $\mathbf{b} = \bar{\mathbf{s}}\mathbf{A}_{\text{lwe}} + \mathbf{e}_{\text{lwe}}^\top = \bar{\mathbf{s}}(\mathbf{B}', \hat{\mathbf{A}}_{\text{fhe}}, \mathbf{A}'_{\text{att}}) + (\mathbf{e}_{\mathbf{B}}^\top, \mathbf{e}_{\text{fhe}}^\top, \mathbf{e}_{\text{att}}^\top)$, then $\mathbf{b}_{\mathbf{B}} = \bar{\mathbf{s}}\mathbf{B}' + \mathbf{e}_{\mathbf{B}}^\top$, $\mathbf{b}_{\text{fhe}} := \bar{\mathbf{s}}\hat{\mathbf{A}}_{\text{fhe}} + \mathbf{e}_{\text{fhe}}^\top$, and $\mathbf{b}_{\text{att}} := \bar{\mathbf{s}}\mathbf{A}'_{\text{att}} + \mathbf{e}_{\text{att}}^\top$.

Thus we have

$$\mathbf{c}_{\mathbf{B},i}^\top = (\bar{\mathbf{s}}, -1) \begin{pmatrix} \mathbf{B}' \\ \underline{\mathbf{b}}^\top \end{pmatrix} + \mathbf{e}_{\mathbf{B}}^\top, \quad \mathbf{A}_{\text{fhe},i} = \begin{pmatrix} \hat{\mathbf{A}}_{\text{fhe}} \\ \bar{\mathbf{s}}\hat{\mathbf{A}}_{\text{fhe}} + \mathbf{e}_{\text{fhe}}^\top \end{pmatrix}, \quad \mathbf{c}_{\text{att},i}^\top = (\bar{\mathbf{s}}, -1) \left(\begin{pmatrix} \bar{\mathbf{A}}_{\text{att}} \\ \underline{\mathbf{a}}_{\text{att}}^\top \end{pmatrix} - \text{bits}(1, \mathbf{X}_i) \otimes \mathbf{G} \right) + \mathbf{e}_{\text{att}}^\top$$

To see the latter case, we note that if $\mathbf{b} \leftarrow \mathbb{Z}_q^{mw+m+(L_X+1)m}$ then it implies $\mathbf{b}_B \leftarrow \mathbb{Z}_q^{mw}$, $\mathbf{b}_{fhe} \leftarrow \mathbb{Z}_q^m$, $\mathbf{b}_{att} \leftarrow \mathbb{Z}_q^{(L_X+1)m}$. This implies the following.

- Randomness of \mathbf{b}_B implies the randomness of $\mathbf{c}_{B,i}^\top := \mathbf{b}_B^\top - \underline{\mathbf{b}}^\top$.
- Randomness of \mathbf{b}_{fhe} implies $\mathbf{A}_{fhe,i} \leftarrow \mathbb{Z}_q^{(n+1) \times m}$.
- Randomness of \mathbf{b}_{att} implies randomness of $\mathbf{c}_{att,i}^\top = \mathbf{b}_{att}^\top - (\mathbf{a}_{att}^\top - \text{bits}(1, \mathbf{X}_i) \otimes \underline{\mathbf{G}})$.

■

Thus, it suffices to show pseudorandomness of the following distribution given aux_2

$$\left(\begin{array}{l} \mathbf{B}, \{\mathbf{c}_{B,j} \leftarrow \mathbb{Z}_q^{mw}\}_{j \in [Q_{\text{msg}}]}, \{\mathbf{X}_j = \mathbf{A}_{fhe,j} \mathbf{R}_j - (\mathbf{x}_j, \text{sd}_j) \otimes \mathbf{G}\}_{j \in [Q_{\text{msg}}]}, \\ \{\mathbf{c}_{att,j} \leftarrow \mathbb{Z}_q^{(L_X+1)m}\}_{j \in [Q_{\text{msg}}]}, \{\tilde{\mathbf{F}}_{j,k} = f_k(\mathbf{x}_j) \lfloor q/2 \rfloor + \text{PRF}(\text{sd}_j, \mathbf{r}_k) + \mathbf{e}_{\mathbf{P},j,k}^\top\}_{j \in [Q_{\text{msg}}], k \in [Q_{\text{key}}]} \end{array} \right)$$

where $\mathbf{A}_{fhe,j} \leftarrow \mathbb{Z}_q^{(n+1) \times m}$.

Hyb₃. This hybrid is same as **Hyb₂** except that for $j \in [Q_{\text{msg}}]$ we sample $\mathbf{X}_j \leftarrow \mathbb{Z}_q^{(n+1) \times m(\lambda+L)}$. We have **Hyb₂** \approx_s **Hyb₃** using leftover hash lemma. By leftover hash lemma (Theorem 2.4) we have that the statistical distance between $\mathbf{A}_{fhe,j} \mathbf{R}_j$ and a uniform matrix $U \leftarrow \mathbb{Z}_q^{(n+1) \times m(\lambda+L)}$ is $m(\lambda+L)/2^n$. This implies that the statistical distance between $\mathbf{X}_j = \mathbf{A}_{fhe,j} \mathbf{R}_j - (\mathbf{x}_j, \text{sd}_j) \otimes \mathbf{G}$ and $\mathbf{X}_j \leftarrow \mathbb{Z}_q^{(n+1) \times m(\lambda+L)}$ is $m(\lambda+L)/2^n$ and we have

$$\Delta(\text{Hyb}_2, \text{Hyb}_3) \leq \frac{Q_{\text{msg}} \cdot m(\lambda+L)}{2^n} \leq \frac{Q_{\text{msg}} \cdot \text{poly}(\lambda)}{2^\lambda}.$$

Thus, it suffices to show pseudorandomness of the following distribution given aux_2

$$\left(\begin{array}{l} \mathbf{B}, \{\mathbf{c}_{B,j} \leftarrow \mathbb{Z}_q^{mw}, \mathbf{X}_j \leftarrow \mathbb{Z}_q^{(n+1) \times m(\lambda+L)}, \mathbf{c}_{att,j} \leftarrow \mathbb{Z}_q^{(L_X+1)m}\}_{j \in [Q_{\text{msg}}]}, \\ \{\tilde{\mathbf{F}}_{j,k} = f_k(\mathbf{x}_j) \lfloor q/2 \rfloor + \text{PRF}(\text{sd}_j, \mathbf{r}_k) + \mathbf{e}_{\mathbf{P},j,k}^\top\}_{j \in [Q_{\text{msg}}], k \in [Q_{\text{key}}]} \end{array} \right)$$

Hyb₄. This hybrid is the same as the previous one except that we replace $\text{PRF}(\text{sd}_j, \cdot)$ with the real random function $\mathbf{R}^j(\cdot)$ for each $j \in [q_{\text{msg}}]$. Since sd_j is not used anywhere else, we can use the security of PRF to conclude that this hybrid is computationally indistinguishable from the previous one.

Hyb₅. This hybrid is same as the previous one except that we output a failure symbol if the set $\{\mathbf{r}_k\}_{k \in [Q_{\text{key}}]}$, in aux_2 , contains a collision. We prove that the probability with which there occurs a collision is negligible in λ . To prove this it suffices to show that there is no $k, k' \in [Q_{\text{key}}]$ such that $k \neq k'$ and $\mathbf{r}_k = \mathbf{r}_{k'}$. The probability of this happening can be bounded by $Q_{\text{key}}^2/2^\lambda$ by taking the union bound with respect to all the combinations of k, k' . Thus the probability of outputting the failure symbol is $Q_{\text{key}}^2/2^\lambda$ which is $\text{negl}(\lambda)$.

Hyb₆. In this hybrid we compute $\tilde{\mathbf{F}}_{j,k}$ as

$$\tilde{\mathbf{F}}_{j,k} = f_k(\mathbf{x}_j) \lfloor q/2 \rfloor + R_{j,k} + \mathbf{e}_{\mathbf{P},j,k}^\top$$

for all $j \in [Q_{\text{msg}}], k \in [Q_{\text{key}}]$. Namely, we use fresh randomness $R_{j,k} \leftarrow [-q/4 + B, q/4 - B]^{1 \times \ell}$ instead of deriving the randomness by $\mathbf{R}^j(\mathbf{r}_k)$. We claim that this change is only conceptual. To see this, we observe that unless the failure condition introduced in **Hyb₅** is satisfied, every invocation of the function \mathbf{R}^j is with respect to a fresh input and thus the output can be replaced with a fresh randomness.

Thus, it suffices to show pseudorandomness of the following distribution given aux_2

$$\left(\mathbf{B}, \{ \mathbf{c}_{\mathbf{B},j} \leftarrow \mathbb{Z}_q^{mw}, \mathbf{X}_j \leftarrow \mathbb{Z}_q^{(n+1) \times m(\lambda+L)}, \mathbf{c}_{\text{att},j} \leftarrow \mathbb{Z}_q^{(L_X+1)m} \}_{j \in [Q_{\text{msg}}]}, \{ \tilde{\mathbf{F}}_{j,k} = f_k(\mathbf{x}_j) \lfloor q/2 \rfloor + \mathbf{R}_{j,k} + \mathbf{e}_{\mathbf{P},j,k}^\top \}_{j \in [Q_{\text{msg}}], k \in [Q_{\text{key}}]} \right)$$

Hyb₇. This hybrid is same as the previous one except we sample $R_{j,k} \leftarrow [-q/4, q/4]^{1 \times \ell}$.

We note that $\text{Hyb}_6 \approx_s \text{Hyb}_7$. To see this note that the statistical distance between the uniform distributions $U_1 = [-q/4 + B, q/4 - B]$ and $U_2 = [-q/4, q/4]$ is

$$\Delta(U_1, U_2) = \frac{1}{2} \left| \frac{2}{q - 4B} - \frac{2}{q} \right| \leq \frac{4B}{q} \leq \frac{\text{poly}(\lambda)}{2^\lambda}$$

by our parameter setting. Therefore,

$$\Delta(\text{Hyb}_2, \text{Hyb}_3) \leq \frac{Q_{\text{key}} \cdot Q_{\text{msg}} \cdot \text{poly}(\lambda)}{2^\lambda}.$$

Hyb₈. This hybrid is same as the previous one except we sample $\tilde{F}_{j,k} \leftarrow \mathbb{Z}_q^\ell$. This follows from the pseudorandomness of $\{f_k(x_j)\}_{j,k}$. To see this note that we have

$$(1^\lambda, \text{aux}_{\mathcal{A}}, \{f_k, f_k(x_j)\}_{j \in [Q_{\text{msg}}], k \in [Q_{\text{key}}]}) \approx_c (1^\lambda, \text{aux}_{\mathcal{A}}, \{f_k, \Delta_{j,k} \leftarrow \{0, 1\}^\ell\}_{j \in [Q_{\text{msg}}], k \in [Q_{\text{key}}]})$$

which implies

$$(1^\lambda, \text{aux}_{\mathcal{A}}, \{f_k, \tilde{F}_{j,k} = f_k(x_j) \lfloor q/2 \rfloor + R_{j,k} + \mathbf{e}_{\mathbf{P},j,k}^\top\}_{j \in [Q_{\text{msg}}], k \in [Q_{\text{key}}]}) \quad (5.23)$$

$$\approx_c (1^\lambda, \text{aux}_{\mathcal{A}}, \{f_k, \tilde{F}_{j,k} \leftarrow \mathbb{Z}_q^\ell\}_{j \in [Q_{\text{msg}}], k \in [Q_{\text{key}}]}) \quad (5.24)$$

where $R_{j,k} \leftarrow [-q/4, q/4]^{1 \times \ell}$ and $\mathbf{e}_{\mathbf{P},j,k} \leftarrow \mathcal{D}_{\mathbb{Z}, \sigma_1}^\ell$.

Thus, using Equation (5.23) and noting that adding random strings does not make the task of distinguishing the two distributions any easier, we achieve the following distribution

$$\left(\text{aux}_{\mathcal{A}}, \mathbf{B}, \{\mathbf{c}_{\mathbf{B},j} \leftarrow \mathbb{Z}_q^{mw}, \mathbf{X}_j \leftarrow \mathbb{Z}_q^{(n+1) \times m(\lambda+L)}, \mathbf{c}_{\text{att},j} \leftarrow \mathbb{Z}_q^{(L_X+1)m}\}_{j \in [Q_{\text{msg}}]}, \{\tilde{F}_{j,k} \leftarrow \mathbb{Z}_q^\ell\}_{j \in [Q_{\text{msg}}], k \in [Q_{\text{key}}]} \right)$$

which is the R.H.S distribution of Equation (5.22), hence the proof. ■

5.3.4 Basing Security on Variant of Circular Evasive LWE ([122])

In this section, we provide our construction of a functional encryption scheme for pseudorandom functionalities for function family $\mathcal{F}_{L(\lambda), \ell(\lambda), \text{dep}(\lambda)} = \{f : \{0, 1\}^L \rightarrow$

$\{0, 1\}^\ell$ basing the security on a variant of the circular evasive assumption introduced by [122], which is considered public-coin.

Assumptions

Here, we state the assumptions used in this section.

Assumption 5.17 (Circular Small Secret LWE). [122] Let n, m, m', q, χ, χ' be functions of λ and

$$\begin{aligned} \bar{\mathbf{A}}_{\text{fhe}} &\leftarrow \mathbb{Z}_q^{n \times m}, \bar{\mathbf{A}}' \leftarrow \mathbb{Z}_q^{n \times m'}, \mathbf{r} \leftarrow \mathcal{D}_{\mathbb{Z}, \chi}^n, \mathbf{s} \leftarrow (\mathbf{r}^\top, -1)^\top, \mathbf{e}_{\text{fhe}} \leftarrow \mathcal{D}_{\mathbb{Z}, \chi}^m, \mathbf{e}' \leftarrow \mathcal{D}_{\mathbb{Z}, \chi'}^{m'}, \\ \mathbf{R} &\leftarrow \{0, 1\}^{m \times (n+1) \lceil \log_2 q \rceil m}, \delta_{\text{fhe}} \leftarrow \mathbb{Z}_q^m, \delta' \leftarrow \mathbb{Z}_q^{m'}, \Delta \leftarrow \mathbb{Z}_q^{(n+1) \times (n+1) \lceil \log_2 q \rceil m} \end{aligned}$$

The circular small-secret LWE assumption $\text{csLWE}_{n, m, m', q, \chi, \chi'}$ states that

$$\begin{aligned} &\left\{ \left(1^\lambda, \begin{pmatrix} \bar{\mathbf{A}}_{\text{fhe}} \\ \mathbf{r}^\top \bar{\mathbf{A}}_{\text{fhe}} + \mathbf{e}_{\text{fhe}}^\top \end{pmatrix}, \begin{pmatrix} \bar{\mathbf{A}}_{\text{fhe}} \\ \mathbf{r}^\top \bar{\mathbf{A}}_{\text{fhe}} + \mathbf{e}_{\text{fhe}}^\top \end{pmatrix} \mathbf{R} - \text{bits}(\mathbf{s}) \otimes \mathbf{G}, \bar{\mathbf{A}}', \mathbf{r}^\top \bar{\mathbf{A}}' + (\mathbf{e}')^\top \right) \right\}_{\lambda \in \mathbb{N}} \\ &\approx \left\{ \left(1^\lambda, \begin{pmatrix} \bar{\mathbf{A}}_{\text{fhe}} \\ \delta_{\text{fhe}}^\top \end{pmatrix}, \Delta, \bar{\mathbf{A}}', (\delta_i)^\top \right) \right\}_{\lambda \in \mathbb{N}} \end{aligned}$$

Next, we define a variant of evcsLWE assumption introduced by [122] further refined to avoid the attacks as discusses in [19].

Assumption 5.18 (evcsLWE). Let $\mathcal{S}(1^\lambda; \text{aux})$ be an algorithm that, given randomness aux , outputs

$$\mathbf{A}_{\text{circ}} \leftarrow \mathbb{Z}_q^{(n+1) \times (m(n+1)^2 \lceil \log_2 q \rceil^2 + 1)m}, \bar{\mathbf{A}}' \in \mathbb{Z}_q^{n \times m'}, \mathbf{P} \in \mathbb{Z}_q^{n \times J}, \sigma, \sigma', \sigma_{-1}, \sigma_{\text{post}}, \sigma_{\text{pre}}$$

where $m \geq m_0(n, q)$ and $\sigma_{-1} \geq \sigma_0(n, m)$ and $\sigma_{\text{post}} \geq \sigma_{\text{pre}}$. Suppose

$$\begin{aligned} \bar{\mathbf{A}}_{\text{fhe}} &\leftarrow \mathbb{Z}_q^{n \times m}, (\mathbf{B}, \tau) \leftarrow \text{TrapGen}(1^n, 1^m, q), \mathbf{K} \leftarrow \mathbf{B}^{-1}(\mathbf{P}), \\ \mathbf{e}_{\text{fhe}} &\leftarrow \mathcal{D}_{\mathbb{Z}, \sigma}^m, \mathbf{e}_{\text{circ}} \leftarrow \mathcal{D}_{\mathbb{Z}, \sigma'}^{(m(n+1)^2 \lceil \log_2 q \rceil^2 + 1)m}, \mathbf{e}' \leftarrow \mathcal{D}_{\mathbb{Z}, \sigma'}^{m'}, \mathbf{e}_{\mathbf{B}} \in \mathbb{Z}^m, \mathbf{e}_{\mathbf{P}} \in \mathbb{Z}^J, \\ \delta_{\text{fhe}} &\leftarrow \mathbb{Z}_q^m, \delta_{\text{circ}} \leftarrow \mathbb{Z}_q^{(m(n+1)^2 \lceil \log_2 q \rceil^2 + 1)m}, \delta' \leftarrow \mathbb{Z}_q^{m'}, \delta_{\mathbf{B}} \leftarrow \mathbb{Z}_q^m, \delta_{\mathbf{P}} \leftarrow \mathbb{Z}_q^J, \\ \mathbf{r} &\leftarrow \mathcal{D}_{\mathbb{Z}, \sigma}^n, \mathbf{s} \leftarrow (\mathbf{r}^\top, -1)^\top, \mathbf{R} \leftarrow \{0, 1\}^{m \times (n+1) \lceil \log_2 q \rceil m}, \Delta \leftarrow \mathbb{Z}_q^{(n+1) \times (n+1) \lceil \log_2 q \rceil m}, \end{aligned}$$

$$\mathbf{S} = \begin{pmatrix} \bar{\mathbf{A}}_{\text{fhe}} \\ \mathbf{r}^\top \bar{\mathbf{A}}_{\text{fhe}} + \mathbf{e}_{\text{fhe}}^\top \end{pmatrix} \mathbf{R} - (\mathbf{x}, \text{bits}(\mathbf{s})) \otimes \mathbf{G} \text{ for } \mathbf{x} \in \{0, 1\}^L$$

In the precondition, the entries of $\mathbf{e}_{\mathbf{B}}, \mathbf{e}_{\mathbf{P}}$ are independent and follow $\mathcal{D}_{\mathbb{Z}, \sigma_{\text{pre}}}$, and $\text{evcsLWE}_{\text{pre}}^{\mathcal{S}}$ states that

$$\left\{ \begin{pmatrix} 1^\lambda, \text{aux}, \bar{\mathbf{A}}_{\text{fhe}}, \mathbf{B}, \mathbf{r}^\top \bar{\mathbf{A}}_{\text{fhe}} + \mathbf{e}_{\text{fhe}}^\top, \mathbf{S}, \\ \mathbf{s}^\top (\mathbf{A}_{\text{circ}} - (1, \text{bits}(\mathbf{S})) \otimes \mathbf{G}) + \mathbf{e}_{\text{circ}}^\top \\ \mathbf{r}^\top \bar{\mathbf{A}}' + (\mathbf{e}')^\top, \mathbf{r}^\top \mathbf{B} + \mathbf{e}_{\mathbf{B}}^\top, \mathbf{r}^\top \mathbf{P} + \mathbf{e}_{\mathbf{P}}^\top \end{pmatrix} \right\}_{\lambda \in \mathbb{N}} \approx \left\{ \begin{pmatrix} 1^\lambda, \text{aux}, \bar{\mathbf{A}}_{\text{fhe}}, \mathbf{B}, \delta_{\text{fhe}}^\top, \Delta, \\ \delta_{\text{circ}}^\top, \\ (\delta')^\top, \delta_{\mathbf{B}}^\top, \delta_{\mathbf{P}}^\top \end{pmatrix} \right\}_{\lambda \in \mathbb{N}}.$$

In the postcondition, the entries of $\mathbf{e}_{\mathbf{B}}$ are independent and follow $\mathcal{D}_{\mathbb{Z}, \sigma_{\text{post}}}$, and $\text{evcsLWE}_{\text{post}}^{\mathcal{S}}$ states that

$$\left\{ \begin{pmatrix} 1^\lambda, \text{aux}, \bar{\mathbf{A}}_{\text{fhe}}, \mathbf{B}, \mathbf{r}^\top \bar{\mathbf{A}}_{\text{fhe}} + \mathbf{e}_{\text{fhe}}^\top, \mathbf{S}, \\ \mathbf{s}^\top (\mathbf{A}_{\text{circ}} - (1, \text{bits}(\mathbf{S})) \otimes \mathbf{G}) + \mathbf{e}_{\text{circ}}^\top \\ \mathbf{r}^\top \bar{\mathbf{A}}' + (\mathbf{e}')^\top, \mathbf{r}^\top \mathbf{B} + \mathbf{e}_{\mathbf{B}}^\top, \mathbf{K} \end{pmatrix} \right\}_{\lambda \in \mathbb{N}} \approx \left\{ \begin{pmatrix} 1^\lambda, \text{aux}, \bar{\mathbf{A}}_{\text{fhe}}, \mathbf{B}, \delta_{\text{fhe}}^\top, \Delta, \\ \delta_{\text{circ}}^\top, \\ (\delta')^\top, \delta_{\mathbf{B}}^\top, \mathbf{K} \end{pmatrix} \right\}_{\lambda \in \mathbb{N}}.$$

The *evasive circular small-secret LWE assumption* with respect to the sampler class \mathcal{SC} states that $\text{evcsLWE}_{\text{pre}}^{\mathcal{S}}$ implies $\text{evcsLWE}_{\text{post}}^{\mathcal{S}}$ for all efficient samplers \mathcal{S} .

We conjecture that for reasonable class of samplers, the evasive LWE assumption holds. In particular, we conjecture that our sampler $\text{Samp}_{\text{prFE}}(1^\lambda)$ used for the security proof of our prFE for natural class of functions should be in the secure class of samplers \mathcal{SC} for which the evasive LWE holds.

Remark 21. In our assumption, the FHE encoding encode the attribute \mathbf{x} , whereas the one in [122] it only encode the FHE secret key. This is created by the difference between ABE and FE, since the attribute \mathbf{x} can be public in the former and must be hidden in the latter.

Construction

The construction is same as in Section 5.3.2 except the following changes.

1. We use a concrete pseudorandom function PRF, $\text{PRF} : \mathbb{Z}^{(n+1) \times m'} \times \mathbb{Z}^{n+1} \rightarrow [-q/4 + B, q/4 - B]^{1 \times \ell}$ defined as $\text{PRF}(\mathbf{A}, \mathbf{s}) = [\mathbf{G}_q \mathbf{G}_p^{-1} (\lfloor (\mathbf{s}^\top \mathbf{A})^\top \rfloor_p)]_B$. Here

$m' = \lceil \ell(\lceil \log q \rceil / \lceil \log p \rceil) \rceil$, $\mathbf{G}_q = \mathbf{I}_\ell \otimes (1, 2, 2^2, \dots, 2^{\lceil \log_2 q \rceil - 1})$, $\mathbf{G}_p^{-1}(\mathbf{x})$ for a vector $\mathbf{x} \in \mathbb{Z}^{m'}$ is $(\text{bits}(\mathbf{x}[1]), \dots, \text{bits}(\mathbf{x}[m']))^\top$, where $\text{bits}(\mathbf{x}[i]) \in \{0, 1\}^{\log p}$ and $[\mathbf{x}]_B$ for any $\mathbf{x} \in \mathbb{Z}_q^\ell$ represents truncating the range to $[-q/2 + B, q/2 - B]^\ell$ by mapping the out-of-range values to 0. Here B is chosen to be exponentially smaller than $q/4$. We note that for our choice of B the statistical distance between the uniform distribution over $[-q/4, q/4]$ and $[-q/4 + B, q/4 - B]$ is negligible. We show how to set p later.

2. In the setup algorithm, we set $L_X = m(L + (n+1)\lceil \log q \rceil)(n+1)\lceil \log q \rceil$ and replace the notation \mathbf{A}_{att} with \mathbf{A}_{circ} to match the notation of the underlying assumption.

3. The $\text{KeyGen}(\text{msk}, f)$ algorithm has the following changes.

- It samples $\mathbf{A}' \leftarrow \mathbb{Z}_q^{(n+1) \times m'}$ where $m' = \lceil \ell(\lceil \log q \rceil / \lceil \log p \rceil) \rceil$ instead of $\mathbf{r} \leftarrow \{0, 1\}^\lambda$.
- It defines function $F[f, \mathbf{A}']$ (instead of $F = F[f, \mathbf{r}]$), with f, \mathbf{A}' hardwired, as follows:
On input $(\mathbf{x} \in \{0, 1\}^L, \text{bits}(\mathbf{s}) \in \{0, 1\}^{(n+1)\lceil \log q \rceil})$, first recover $\mathbf{s} \in \mathbb{Z}^{n+1}$ from $\text{bits}(\mathbf{s})$ and then compute and output $f(\mathbf{x}) \lfloor q/2 \rfloor + \text{PRF}(\mathbf{A}', \mathbf{s}) \in \mathbb{Z}_q^{1 \times \ell}$.

The rest of the algorithm remains as it is. It outputs $\text{sk}_f = (\mathbf{K}, \mathbf{A}')$.

4. The encryption algorithm has the following changes

- It samples \mathbf{R} differently as $\mathbf{R} \leftarrow \{0, 1\}^{m \times m(L + (n+1)\lceil \log q \rceil)}$ and computes $\mathbf{X} = \mathbf{A}_{\text{the}} \mathbf{R} - (\mathbf{x}, \text{bits}(\mathbf{s})) \otimes \mathbf{G} \in \mathbb{Z}_q^{(n+1) \times m(L + (n+1)\lceil \log q \rceil)}$.
- We also change the notation of \mathbf{c}_{att} and denote this as $\mathbf{c}_{\text{circ}}^\top = \mathbf{s}^\top (\mathbf{A}_{\text{circ}} - \text{bits}(1, \mathbf{X}) \otimes \mathbf{G}) + \mathbf{e}_{\text{circ}}^\top$ for $\mathbf{e}_{\text{circ}} \leftarrow \mathcal{D}_{\mathbb{Z}, \sigma}^{(L_X + 1)m}$.

The rest of the algorithm remains as it is.

The above modified prFE scheme satisfies correctness. This can be argued using same steps as for the construction in Section 5.3.2.

Security

Theorem 5.19. *Let SC_{prFE} be a sampler class for prFE. Assuming circular small-secret LWE (Assumption 5.17) and evasive circular small-secret LWE (Assumption 5.18) with respect to the sampler class that contains all $\text{Samp}_{\text{evcs}}(1^\lambda)$ induced by $\text{Samp}_{\text{prFE}} \in SC_{\text{prFE}}$ as defined in Figure 5.2, our prFE scheme satisfies prCT security with respect to SC_{prFE} as defined in Definition 5.13.*

Proof. For a prFE sampler $\text{Samp}_{\text{prFE}}$ as defined in the proof for Theorem 5.15, we show

$$\begin{pmatrix} \text{mpk} = (\mathbf{A}_{\text{circ}}, \mathbf{B}, M), \text{aux}_{\mathcal{A}}, \mathbf{C}_{\mathbf{B}} = \mathbf{S}\mathbf{B} + \mathbf{E}_{\mathbf{B}}, \\ \{\mathbf{X}_j = \mathbf{A}_{\text{the},j}\mathbf{R}_j - (\mathbf{x}_j, \text{bits}(\mathbf{s}_j)) \otimes \mathbf{G}\}_{j \in [Q_{\text{msg}}]}, \\ \{\mathbf{c}_{\text{circ},j}^{\top} = \mathbf{s}_j^{\top} (\mathbf{A}_{\text{circ}} - \text{bits}(1, \mathbf{X}_j) \otimes \mathbf{G}) + \mathbf{e}_{\text{circ},j}^{\top}\}_{j \in [Q_{\text{msg}}]}, \\ \{\mathbf{K}_k, \mathbf{A}'_k\}_{k \in [Q_{\text{key}}]} \end{pmatrix} \approx_c \begin{pmatrix} \text{mpk} = (\mathbf{A}_{\text{circ}}, \mathbf{B}, M), \text{aux}_{\mathcal{A}}, \mathbf{C}_{\mathbf{B}} \leftarrow \mathbb{Z}_q^{Q_{\text{msg}} \times mw}, \\ \{\mathbf{X}_j \leftarrow \mathbb{Z}_q^{(n+1) \times m(L+(n+1)\lceil \log q \rceil)}\}_{j \in [Q_{\text{msg}}]}, \\ \{\mathbf{c}_{\text{circ},j} \leftarrow \mathbb{Z}_q^{(L_X+1)m}\}_{j \in [Q_{\text{msg}}]}, \\ \{\mathbf{K}_k, \mathbf{A}'_k\}_{k \in [Q_{\text{key}}]} \end{pmatrix} \quad (5.25)$$

where

$$\mathbf{S} = \begin{pmatrix} \mathbf{s}_1^{\top} \\ \vdots \\ \mathbf{s}_{Q_{\text{msg}}}^{\top} \end{pmatrix}, \mathbf{E}_{\mathbf{B}} = \begin{pmatrix} \mathbf{e}_{\mathbf{B},1}^{\top} \\ \vdots \\ \mathbf{e}_{\mathbf{B},Q_{\text{msg}}}^{\top} \end{pmatrix},$$

$$(\text{aux}_{\mathcal{A}}, \{f_k\}_{k \in [Q_{\text{key}}]}, \{\mathbf{x}_j\}_{j \in [Q_{\text{msg}}]}) \leftarrow \text{Samp}_{\text{prFE}}(1^\lambda)$$

and for $j \in [Q_{\text{msg}}]$, $\mathbf{s}_j, \mathbf{e}_{\mathbf{B},j}, \mathbf{A}_{\text{the},j}, \mathbf{R}_j, \mathbf{s}_d, \mathbf{e}_{\text{circ},j}$ are sampled as in the construction, for $k \in [Q_{\text{key}}]$, we have $\mathbf{A}' \leftarrow \mathbb{Z}_q^{(n+1) \times m'}$, $F_k = F[f_k, \mathbf{A}'_k]$ and \mathbf{A}_{F_k} is as defined in the construction, and $\mathbf{K}_k = \mathbf{B}_\tau^{-1}(\mathbf{A}_{F_k})$

assuming we have

$$(1^\lambda, \text{aux}_{\mathcal{A}}, \{f_k, f_k(x_j)\}_{j \in [Q_{\text{msg}}], k \in [Q_{\text{key}}]}) \approx_c (1^\lambda, \text{aux}_{\mathcal{A}}, \{f_k, \Delta_{j,k} \leftarrow \{0,1\}^\ell\}_{j \in [Q_{\text{msg}}], k \in [Q_{\text{key}}]}). \quad (5.26)$$

Consider the evcsLWE sampler $\text{Samp}_{\text{evcs}}(1^\lambda, \text{aux})$ as defined in Figure 5.2. To prove Equation (5.25), assuming evcsLWE assumption (Assumption 5.18) with sampler

$$\text{Samp}_{\text{evcs}}(1^\lambda, \text{aux})$$

The sampler does the following.

- Parse $\text{aux} = (\text{aux}_{\mathcal{A}}, \{f_k\}_{k \in [Q_{\text{key}}]}, \{\mathbf{A}'_k \leftarrow \mathbb{Z}_q^{(n+1) \times m'}\}_{k \in [Q_{\text{key}}]}, \text{aux}_{\text{Samp}})$.
- Sample parameters as in Equation (5.12). Additionally set $p = q/(2^\lambda \sigma \sqrt{\lambda})^a$.
- Samples $\mathbf{A}_{\text{circ}} \leftarrow \mathbb{Z}_q^{(n+1) \times (L_X+1)m}$ using the randomness aux_{Samp} .
- Defines $F_k[f_k, \mathbf{A}'_k]$ and computes \mathbf{A}_{F_k} as in the construction. \mathbf{A}_{F_k} is publicly computable given f_k, \mathbf{A}'_k, M and \mathbf{A}_{circ} .
- It outputs

$$\mathbf{A}_{\text{circ}}, \quad \mathbf{P} = [\mathbf{A}_{F_1} \parallel \dots \parallel \mathbf{A}_{F_{Q_{\text{key}}}}], \quad \text{aux} = (\text{aux}_{\text{Samp}}, \text{aux}_{\text{prFE}}, \{\mathbf{A}'_k\}_{k \in [Q_{\text{key}}]})$$

^aWe assume the parameters to be part of the output of $\text{Samp}_{\text{evcs}}$, even though we do not explicitly write so.

Figure 5.2: Description of the Sampler for circular small-secret evasive LWE

$\text{Samp}_{\text{evcs}}$ it suffices to show

$$\left(\begin{array}{l} \text{aux}_1, \mathbf{B}, \mathbf{C}_B = \mathbf{S}\mathbf{B} + \mathbf{E}_B, \\ \{\mathbf{X}_j = \mathbf{A}_{\text{fhe},j} \mathbf{R}_j - (\mathbf{x}_j, \text{bits}(\mathbf{s}_j)) \otimes \mathbf{G}\}_{j \in [Q_{\text{msg}}]}, \\ \{\mathbf{c}_{\text{circ},j}^\top = \mathbf{s}_j^\top (\mathbf{A}_{\text{circ}} - \text{bits}(1, \mathbf{X}_j) \otimes \mathbf{G}) + \mathbf{e}_{\text{circ},j}^\top\}_{j \in [Q_{\text{msg}}]}, \\ \mathbf{C}_P = \mathbf{S}\mathbf{P} + \mathbf{E}_P \end{array} \right) \approx_c \left(\begin{array}{l} \text{aux}_1, \mathbf{B}, \mathbf{C}_B \leftarrow \mathbb{Z}_q^{Q_{\text{msg}} \times mw}, \\ \{\mathbf{X}_j \leftarrow \mathbb{Z}_q^{(n+1) \times m(L+(n+1)\lceil \log q \rceil)}\}_{j \in [Q_{\text{msg}}]}, \\ \{\mathbf{c}_{\text{circ},j} \leftarrow \mathbb{Z}_q^{(L_X+1)m}\}_{j \in [Q_{\text{msg}}]}, \\ \mathbf{C}_P \leftarrow \mathbb{Z}_q^{Q_{\text{msg}} \times \ell \cdot Q_{\text{key}}} \end{array} \right) \quad (5.27)$$

where $\text{aux}_1 = (\text{aux}, M)$ and $\mathbf{E}_P \leftarrow \mathcal{D}_{\mathbb{Z}, \sigma_1}^{Q_{\text{msg}} \times \ell \cdot Q_{\text{key}}}$. Using the representation in

Section 5.3.3, we rewrite the Equation (5.27) as

$$\left(\begin{array}{l} \text{aux}_1, \mathbf{B}, \{\mathbf{c}_{B,j}^\top = \mathbf{s}_j^\top \mathbf{B} + \mathbf{e}_{B,j}^\top\}_{j \in [Q_{\text{msg}}]}, \\ \{\mathbf{X}_j = \mathbf{A}_{\text{fhe},j} \mathbf{R}_j - (\mathbf{x}_j, \text{bits}(\mathbf{s}_j)) \otimes \mathbf{G}\}_{j \in [Q_{\text{msg}}]}, \\ \{\mathbf{c}_{\text{circ},j}^\top = \mathbf{s}_j^\top (\mathbf{A}_{\text{circ}} - \text{bits}(1, \mathbf{X}_j) \otimes \mathbf{G}) + \mathbf{e}_{\text{circ},j}^\top\}_{j \in [Q_{\text{msg}}]}, \\ \{\mathbf{c}_{P,j,k}^\top = \mathbf{s}_j^\top \mathbf{A}_{F_k} + \mathbf{e}_{P,j,k}^\top\}_{j \in [Q_{\text{msg}}], k \in [Q_{\text{key}}]} \end{array} \right) \approx_c \left(\begin{array}{l} \text{aux}_1, \mathbf{B}, \{\mathbf{c}_{B,j} \leftarrow \mathbb{Z}_q^{mw}\}_{j \in [Q_{\text{msg}}]}, \\ \{\mathbf{X}_j \leftarrow \mathbb{Z}_q^{(n+1) \times m(L+(n+1)\lceil \log q \rceil)}\}_{j \in [Q_{\text{msg}}]}, \\ \{\mathbf{c}_{\text{circ},j} \leftarrow \mathbb{Z}_q^{(L_X+1)m}\}_{j \in [Q_{\text{msg}}]}, \\ \{\mathbf{c}_{P,j,k} \leftarrow \mathbb{Z}_q^\ell\}_{j \in [Q_{\text{msg}}], k \in [Q_{\text{key}}]} \end{array} \right) \quad (5.28)$$

where $\mathbf{e}_{\mathbf{P},j,k} \leftarrow \mathcal{D}_{\mathbb{Z},\sigma_1}^\ell$.

We prove Equation (5.28) using the following sequence of hybrids.

Hyb₀. This is L.H.S distribution of Equation (5.28).

Hyb₁. This hybrid is same as **Hyb₀**, except we compute $\mathbf{c}_{\mathbf{P},j,k}^\top$ as

$$\mathbf{c}_{\mathbf{P},j,k}^\top = M \cdot \left\lfloor \frac{\mathbf{c}_{\text{circ},j}^\top \cdot \mathbf{H}_{\mathbf{A}_{\text{circ}},\mathbf{X}_j}^{\text{High},k}}{M} \right\rfloor + \left\lfloor \frac{\mathbf{c}_{\text{circ},j}^\top \cdot \mathbf{H}_{\mathbf{A}_{\text{circ}},\mathbf{X}_j}^{\text{Flow},k}}{M} \right\rfloor + f_k(\mathbf{x}_j) \lfloor q/2 \rfloor + \text{PRF}(\mathbf{A}'_k, \mathbf{s}_j) + \mathbf{e}_{\mathbf{P},j,k}^\top$$

where $\mathbf{e}_{\mathbf{P},j,k} \leftarrow \mathcal{D}_{\mathbb{Z},\sigma_1}^\ell$. The proof of $\text{Hyb}_0 \approx_s \text{Hyb}_1$ is exactly the same as the proof of $\text{Hyb}_0 \approx_s \text{Hyb}_1$ in Theorem 5.15, hence we omit it. So it suffices to show pseudorandomness of the following distribution given aux_1

$$\left(\begin{array}{l} \mathbf{B}, \{ \mathbf{c}_{\mathbf{B},j}^\top = \mathbf{s}_j^\top \mathbf{B} + \mathbf{e}_{\mathbf{B},j}^\top \}_{j \in [Q_{\text{msg}}]} \\ \{ \mathbf{X}_j = \mathbf{A}_{\text{fhe},j} \mathbf{R}_j - (\mathbf{x}_j, \text{bits}(\mathbf{s}_j)) \otimes \mathbf{G} \}_{j \in [Q_{\text{msg}}]}, \\ \{ \mathbf{c}_{\text{circ},j}^\top = \mathbf{s}_j^\top (\mathbf{A}_{\text{circ}} - \text{bits}(1, \mathbf{X}_j) \otimes \mathbf{G}) + \mathbf{e}_{\text{circ},j}^\top \}_{j \in [Q_{\text{msg}}]}, \\ \{ \tilde{\mathbf{F}}_{j,k} = f_k(\mathbf{x}_j) \lfloor q/2 \rfloor + [\mathbf{G}_q \mathbf{G}_p^{-1} (\lfloor (\mathbf{s}_j^\top \mathbf{A}'_k)^\top \rfloor_p)]_B + \mathbf{e}_{\mathbf{P},j,k}^\top \}_{j \in [Q_{\text{msg}}], k \in [Q_{\text{key}}]} \end{array} \right)$$

$$\text{where we write } \text{PRF}(\mathbf{A}'_k, \mathbf{s}_j) = \left[\mathbf{G}_q \mathbf{G}_p^{-1} \left(\left\lfloor (\mathbf{s}_j^\top \mathbf{A}'_k)^\top \right\rfloor_p \right) \right]_B.$$

Hyb₂: This hybrid is same as **Hyb₁** except that for all $j \in [Q_{\text{msg}}], k \in [Q_{\text{key}}]$, we

compute the PRF differently. We compute

$$\text{PRF}(\mathbf{A}'_k, \mathbf{s}_j) = \left[\mathbf{G}_q \mathbf{G}_p^{-1} \left(\left\lfloor (\bar{\mathbf{s}}_j^\top \bar{\mathbf{A}}'_k - (\underline{\mathbf{a}}'_k)^\top + \mathbf{e}_k^\top)^\top \right\rfloor_p \right) \right]_B \text{ for } \mathbf{e}_k \leftarrow \mathcal{D}_{\mathbb{Z},\sigma}^\ell, \text{ where}$$

$\bar{\mathbf{A}}'_k$ denotes the first n rows and $\underline{\mathbf{a}}'_k$ denotes the last row of the matrix \mathbf{A}'_k . We

claim that $\text{Hyb}_1 \approx_s \text{Hyb}_2$ with all but negligible probability. To see this, we note the following.

- The statistical distance between $(\bar{\mathbf{s}}_j^\top \bar{\mathbf{A}}'_k - (\underline{\mathbf{a}}'_k)^\top + \mathbf{e}_k^\top)^\top$ and $(\bar{\mathbf{s}}_j^\top \bar{\mathbf{A}}'_k - (\underline{\mathbf{a}}'_k)^\top)^\top$ is bounded by $\|\mathbf{e}_k\| \leq \sigma\sqrt{\lambda}$.

- So,

$$\Pr \left[\left[(\bar{\mathbf{s}}_j^\top \bar{\mathbf{A}}'_k - \underline{\mathbf{a}}'_k)^\top \right]_p \neq \left[(\bar{\mathbf{s}}_j^\top \bar{\mathbf{A}}'_k - \underline{\mathbf{a}}'_k)^\top + \mathbf{e}_k^\top \right]_p \right] \leq \sigma \sqrt{\lambda} p / q = 1/2^\lambda$$
 due to our parameter setting.
- Taking the probability over all $k \in [Q_{\text{key}}]$, we get $\text{Hyb}_1 \approx_s \text{Hyb}_2$ with all but $Q_{\text{key}}/2^\lambda$ probability.

So it suffices to show pseudorandomness of the following distribution given aux_1

$$\left(\begin{array}{l} \mathbf{B}, \{ \mathbf{c}_{\mathbf{B},j}^\top = \mathbf{s}_j^\top \mathbf{B} + \mathbf{e}_{\mathbf{B},j}^\top \}_{j \in [Q_{\text{msg}}]} \\ \{ \mathbf{X}_j = \mathbf{A}_{\text{fhe},j} \mathbf{R}_j - (\mathbf{x}_j, \text{bits}(\mathbf{s}_j)) \otimes \mathbf{G} \}_{j \in [Q_{\text{msg}}]}, \\ \{ \mathbf{c}_{\text{circ},j}^\top = \mathbf{s}_j^\top (\mathbf{A}_{\text{circ}} - \text{bits}(1, \mathbf{X}_j) \otimes \mathbf{G}) + \mathbf{e}_{\text{circ},j}^\top \}_{j \in [Q_{\text{msg}}]}, \\ \{ \tilde{\mathbf{F}}_{j,k} = f_k(\mathbf{x}_j) \lfloor q/2 \rfloor + [\mathbf{G}_q \mathbf{G}_p^{-1} (\lfloor (\bar{\mathbf{s}}_j^\top \bar{\mathbf{A}}'_k - \underline{\mathbf{a}}'_k)^\top + \mathbf{e}_k^\top \rfloor_p)]_B + \mathbf{e}_{\mathbf{P},j,k}^\top \}_{j \in [Q_{\text{msg}}], k \in [Q_{\text{key}}]} \end{array} \right)$$

Hyb₃: This hybrid is same as Hyb_2 except that for all $j \in [Q_{\text{msg}}]$ we sample $\mathbf{c}_{\mathbf{B},j} \leftarrow \mathbb{Z}_q^{mw}$, $\mathbf{c}_{\text{circ},j} \leftarrow \mathbb{Z}_q^{(L_X+1)m}$, $\mathbf{X}_j \leftarrow \mathbb{Z}_q^{(n+1) \times m(L+(n+1)\lceil \log q \rceil)}$ and $R_{j,k} \leftarrow \mathbb{Z}_q^\ell$, where $R_{j,k} = \mathbf{G}_q \mathbf{G}_p^{-1} (\lfloor (\bar{\mathbf{s}}_j^\top \bar{\mathbf{A}}'_k - \underline{\mathbf{a}}'_k)^\top + \mathbf{e}_k^\top \rfloor_p)$ in Hyb_2 . We have $\text{Hyb}_1 \approx_c \text{Hyb}_2$ assuming the hardness of circular small secret LWE.

To prove this we consider sub-hybrids $\text{Hyb}_{1,i}$ for $i \in [Q_{\text{msg}}]$, where in $\text{Hyb}_{1,i}$ we sample $\mathbf{c}_{\mathbf{B},j} \leftarrow \mathbb{Z}_q^{mw}$, $\mathbf{c}_{\text{att},j} \leftarrow \mathbb{Z}_q^{(L_X+1)m}$, $\mathbf{X}_j \leftarrow \mathbb{Z}_q^{(n+1) \times m(L+(n+1)\lceil \log q \rceil)}$ and $R_{j,k} \leftarrow \mathbb{Z}_q^\ell$ for $1 \leq j \leq i$. We set $\text{Hyb}_1 = \text{Hyb}_{1,0}$ and $\text{Hyb}_2 = \text{Hyb}_{1,Q_{\text{msg}}}$. We prove that for all $i \in [Q_{\text{msg}}]$, $\text{Hyb}_{1,i-1} \approx_c \text{Hyb}_{1,i}$ in Claim 5.20. Thus, it suffices to show pseudorandomness of the following distribution given aux_2

$$\left(\begin{array}{l} \mathbf{B}, \{ \mathbf{c}_{\mathbf{B},j} \leftarrow \mathbb{Z}_q^{mw} \}_{j \in [Q_{\text{msg}}]}, \{ \mathbf{X}_j \leftarrow \mathbb{Z}_q^{(n+1) \times m(L+(n+1)\lceil \log q \rceil)} \}_{j \in [Q_{\text{msg}}]}, \\ \{ \mathbf{c}_{\text{circ},j} \leftarrow \mathbb{Z}_q^{(L_X+1)m} \}_{j \in [Q_{\text{msg}}]}, \{ \tilde{\mathbf{F}}_{j,k} = f_k(\mathbf{x}_j) \lfloor q/2 \rfloor + [R_{j,k}]_B + \mathbf{e}_{\mathbf{P},j,k}^\top \}_{j \in [Q_{\text{msg}}], k \in [Q_{\text{key}}]} \end{array} \right).$$

Hyb₄. This hybrid is same as the previous one except we sample $\tilde{\mathbf{F}}_{j,k} \leftarrow \mathbb{Z}_q^\ell$. This follows from the pseudorandomness of $\{f_k(x_j)\}_{j,k}$. Thus we achieve the following

distribution

$$\left(\text{aux}_{\mathcal{A}}, \mathbf{B}, \{ \mathbf{c}_{\mathbf{B},j} \leftarrow \mathbb{Z}_q^{mw}, \mathbf{X}_j \leftarrow \mathbb{Z}_q^{(n+1) \times m(L+(n+1)\lceil \log q \rceil)}, \mathbf{c}_{\text{circ},j} \leftarrow \mathbb{Z}_q^{(L_X+1)m} \}_{j \in [Q_{\text{msg}}]}, \{ \tilde{\mathbf{F}}_{j,k} \leftarrow \mathbb{Z}_q^\ell \}_{j \in [Q_{\text{msg}}], k \in [Q_{\text{key}}]} \right)$$

which is the R.H.S distribution of Equation (5.22), hence the proof. \blacksquare

Claim 5.20. $\text{Hyb}_{1,i-1} \approx_c \text{Hyb}_{1,i}$.

Proof. We show that if there exists an adversary \mathcal{A} who can distinguish between the two hybrids with non-negligible advantage, then there is a reduction \mathcal{B} that breaks csLWE security with non-negligible advantage. The reduction is as follows.

1. The adversary \mathcal{A} sends the function queries $f_1, \dots, f_{Q_{\text{key}}}$, message queries $\mathbf{x}_1, \dots, \mathbf{x}_{Q_{\text{msg}}}$ and auxiliary input $\text{aux}_{\mathcal{A}}$ to the reduction.
2. \mathcal{B} initiates the csLWE security game with the csLWE challenger. The csLWE challenger samples a bit $\beta \leftarrow \{0, 1\}$ and does the following

- It samples $\bar{\mathbf{A}}_{\text{fhe},j} \leftarrow \mathbb{Z}_q^{n \times m}$, $\bar{\mathbf{A}}'_j \leftarrow \mathbb{Z}_q^{n \times (mw + (L_X+1)m + m' \cdot Q_{\text{key}})}$, $\bar{\mathbf{s}}_j \leftarrow \mathcal{D}_{\mathbb{Z}, \sigma_s}^n$, $\mathbf{e}_{\text{fhe},j} \leftarrow \mathcal{D}_{\mathbb{Z}, \sigma}^m$, $\mathbf{e}'_j \leftarrow \mathcal{D}_{\mathbb{Z}, \sigma}^{mw + (L_X+1)m + m' \cdot Q_{\text{key}}}$, $\mathbf{R}_j \leftarrow \{0, 1\}^{m \times (n+1)\lceil \log_2 q \rceil m}$, $\delta_{\text{fhe},j} \leftarrow \mathbb{Z}_q^m$, $\delta'_j \leftarrow \mathbb{Z}_q^{mw + (L_X+1)m + m' \cdot Q_{\text{key}}}$, $\Delta_j \leftarrow \mathbb{Z}_q^{(n+1) \times (n+1)\lceil \log_2 q \rceil m}$. It sets $\mathbf{s}_j = (\bar{\mathbf{s}}_j^\top, -1)^\top$.

- If $\beta = 0$, it sets $\mathbf{b}_{\text{fhe},j}^\top := \bar{\mathbf{s}}_j^\top \bar{\mathbf{A}}_{\text{fhe},j} + \mathbf{e}_{\text{fhe},j}^\top$, $\mathbf{S}_j := \left(\bar{\mathbf{s}}_j^\top \bar{\mathbf{A}}_{\text{fhe},j} + \mathbf{e}_{\text{fhe},j}^\top \right) \mathbf{R}_j - \text{bits}(\mathbf{s}_j) \otimes \mathbf{G}$, and $(\mathbf{b}'_j)^\top = \bar{\mathbf{s}}_j^\top \bar{\mathbf{A}}'_j + (\mathbf{e}'_j)^\top$.

- If $\beta = 1$, it sets $\mathbf{b}_{\text{fhe},j} := \delta_{\text{fhe},j}$, $\mathbf{S}_j := \Delta_j$, and $\mathbf{b}'_j := \delta'_j$.

- It returns $(\mathbf{A}_{\text{fhe},j} = (\bar{\mathbf{A}}_{\text{fhe},j} \mathbf{b}_{\text{fhe},j}^\top)^\top, \mathbf{S}_j, \bar{\mathbf{A}}'_j, (\mathbf{b}'_j)^\top)$ to the reduction.

3. \mathcal{B} does the following.

- For $1 \leq j < i$: \mathcal{B} samples $\mathbf{c}_{\mathbf{B},j} \leftarrow \mathbb{Z}_q^{mw}$, $\mathbf{c}_{\text{circ},j} \leftarrow \mathbb{Z}_q^{(L_X+1)m}$, $\mathbf{X}_j \leftarrow \mathbb{Z}_q^{(n+1) \times m(L+(n+1)\lceil \log q \rceil)}$ and $R_{j,k} \leftarrow \mathbb{Z}_q^\ell$.

- For $j = i$, it does as follows.

- a) It samples $\mathbf{R}'_j \leftarrow \{0, 1\}^{m \times mL}$, computes $\mathbf{C}_j := \mathbf{A}_{\text{fhe},j} \mathbf{R}'_j - \mathbf{x}_j \otimes \mathbf{G}$, sets $\mathbf{X}_j = \mathbf{S}_j + \mathbf{C}_j$.
- b) It parses $\bar{\mathbf{A}}'_j = [\bar{\mathbf{B}}, \bar{\mathbf{A}}'_{\text{circ},j}, \bar{\mathbf{A}}'_{j,1}, \dots, \bar{\mathbf{A}}'_{j,Q_{\text{key}}}]$, where $\bar{\mathbf{B}} \in \mathbb{Z}_q^{n \times mw}$, $\bar{\mathbf{A}}'_{\text{circ},j} \in \mathbb{Z}_q^{n \times (L_X+1)m}$, $\bar{\mathbf{A}}'_{j,k} \in \mathbb{Z}_q^{n \times m'}$, and similarly $(\mathbf{b}'_j)^\top = ((\mathbf{b}_j)^\top, (\mathbf{b}'_{\text{circ},j})^\top, (\mathbf{b}'_{j,1})^\top, \dots, (\mathbf{b}'_{j,Q_{\text{key}}})^\top)$, where $\mathbf{b}_j \in \mathbb{Z}_q^{mw}$, $\mathbf{b}'_{\text{circ},j} \in \mathbb{Z}_q^{(L_X+1)m}$ and $\mathbf{b}'_{j,k} \in \mathbb{Z}_q^{m'}$ for $k \in [Q_{\text{key}}]$.
- c) It samples $\underline{\mathbf{b}} \leftarrow \mathbb{Z}_q^{mw}$, $\underline{\mathbf{a}}_{\text{circ},j} \leftarrow \mathbb{Z}_q^{(L_X+1)m}$ and sets $\mathbf{c}_{\mathbf{B},j} := (\mathbf{b}_j)^\top - \underline{\mathbf{b}}^\top$ and $\mathbf{c}_{\text{circ},j}^\top := (\mathbf{b}'_{\text{circ},j})^\top - (\underline{\mathbf{a}}_{\text{circ},j}^\top - \text{bits}(1, \mathbf{X}_j) \otimes \mathbf{G})$.
- d) For $k \in [Q_{\text{key}}]$, it samples $\underline{\mathbf{a}}'_k \leftarrow \mathbb{Z}_q^{m'}$ and sets $R_{j,k} = (\mathbf{b}'_{j,k})^\top - (\underline{\mathbf{a}}'_k)^\top$.

– For $j > i$, it computes $\mathbf{c}_{\mathbf{B},j}$, $\mathbf{c}_{\text{circ},j}$, \mathbf{X}_j and $R_{j,k}$ as in Hyb_2 .

4. It sets $\text{aux}_1 = (\text{aux}, M)$, $\mathbf{B} = (\bar{\mathbf{B}} \quad \underline{\mathbf{b}}^\top)^\top$ and $\tilde{\mathbf{F}}_{j,k} = f_k(\mathbf{x}_j) \lfloor q/2 \rfloor + [\mathbf{G}_q \mathbf{G}_p^{-1} (\lfloor (R_{j,k})^\top \rfloor_p)]_B + \mathbf{e}_{\mathbf{P},j,k}^\top$ and sends $(\text{aux}_1, \mathbf{B}, \mathbf{c}_{\mathbf{B},j}, \mathbf{X}_j, \mathbf{c}_{\text{circ},j}, \tilde{\mathbf{F}}_{j,k})$ for all $j \in [Q_{\text{msg}}]$ and $k \in [Q_{\text{key}}]$ to \mathcal{A} .

5. \mathcal{A} outputs a bit β' . \mathcal{B} forwards the bit β' to the csLWE challenger.

We note that if the csLWE challenger choose $\beta = 0$ then \mathcal{B} simulated $\text{Hyb}_{1,i-1}$ with \mathcal{A} else it simulated $\text{Hyb}_{1,i}$ with \mathcal{A} . ■

5.4 LACONIC PSEUDORANDOM POLY-DOMAIN OBFUSCATION

In this section, we introduce the notion of laconic pPRIO and construct it from several ingredients, which are all implied by the evasive LWE and LWE. The construction of laconic pPRIO will be used in Section 5.5. Since the intuition for this notion was discussed in Section 5.1, we proceed directly to the formal definition.

5.4.1 Definition

Syntax. A laconic pPRIO scheme supporting any circuit consists of the following algorithms.

$\text{LDigest}(1^\lambda, X = \{X_i\}_{i \in [N]}) \rightarrow \text{dig}$. The digest algorithm takes as input the security parameter λ and an input space X of the form $X = \{X_i \in \{0, 1\}^\ell\}_{i \in [N]}$ for some

$\ell = \ell(\lambda)$ and $N \in \mathbb{N}$. We assume that X encodes the information of ℓ and N and one can retrieve them efficiently. It outputs a string dig .

$\text{LObfuscate}(1^\lambda, \text{dig}, E) \rightarrow \text{Lobf}$. The obfuscate algorithm takes as input the security parameter λ , string dig and a circuit $E : \{0, 1\}^\ell \rightarrow \{0, 1\}^L$ whose size is S . It outputs a ciphertext Lobf .

We decompose this algorithm into two phases.

$\text{LObfOff}(1^\lambda, 1^S) \rightarrow (\text{Lobf}_{\text{off}}, \text{st})$. The offline obfuscate algorithm takes as input the security parameter λ and the circuit size S . It outputs Lobf_{off} and a state st .

$\text{LObfOn}(\text{st}, \text{dig}, E) \rightarrow \text{Lobf}_{\text{on}}$. The online obfuscate algorithm takes as input the state st , string dig and circuit E . It outputs Lobf_{on} .

With the above decomposition, the obfuscate algorithm outputs $\text{Lobf} = (\text{Lobf}_{\text{off}}, \text{Lobf}_{\text{on}})$.

$\text{LEval}(X, \text{Lobf}) \rightarrow Y$. The evaluation algorithm takes as input $X = \{X_i \in \{0, 1\}^\ell\}_{i \in [N]}$ and Lobf . It outputs $Y = \{Y_i \in \{0, 1\}^L\}_{i \in [N]}$.

Definition 5.15 (Compactness). For all $\ell, N \in \mathbb{N}$ and $X = \{X_i \in \{0, 1\}^\ell\}_{i \in [N]}$, we need that the size of $\text{dig} = \text{LDigest}(1^\lambda, X)$ should be bounded by $\text{poly}(\lambda)$. In particular, the size of dig should be independent of N .

Definition 5.16 (Correctness). For all $\ell, N \in \mathbb{N}$, $X = \{X_i \in \{0, 1\}^\ell\}_{i \in [N]}$, and circuit $E : \{0, 1\}^\ell \rightarrow \{0, 1\}^L$ such that $|E| \leq S$ for an arbitrary polynomial $S = S(\lambda)$, we have

$$\Pr \left[\begin{array}{c} \text{LDigest}(1^\lambda, X = \{X_i\}_{i \in [N]}) \rightarrow \text{dig} \\ \text{LEval}(X, (\text{Lobf}_{\text{off}}, \text{Lobf}_{\text{on}})) = \{E(X_i)\}_{i \in [N]} : \begin{array}{l} \text{LObfOff}(1^\lambda, 1^S) \rightarrow (\text{Lobf}_{\text{off}}, \text{st}) \\ \text{LObfOn}(\text{st}, \text{dig}, E) \rightarrow \text{Lobf}_{\text{on}} \end{array} \end{array} \right] = 1.$$

Remark 22. We note that the above correctness requirement implies that for the correctness to hold, LObfOff only has to know the upper bound S on the size of the

circuit E that is going to be input to LObfOn and does not have to know anything else.

Definition 5.17 (Security). Let Samp be a PPT algorithm that on input 1^λ , outputs

$$\left(\text{aux}, 1^S, X^1 = \{X_i^1\}_{i \in [N^1]}, \dots, X^Q = \{X_i^Q\}_{i \in [N^Q]}, E^1, \dots, E^Q \right)$$

where for all for $k \in [Q], i \in [N^k], X_i^k \in \{0, 1\}^{\ell^k}, E^k : \{0, 1\}^{\ell^k} \rightarrow \{0, 1\}^{L^k}$ and $|E^k| \leq S$. We say that a laconic pPRIO scheme is secure with respect to the sampler class \mathcal{SC} if for every PPT sampler $\text{Samp} \in \mathcal{SC}$ the following holds.

$$\begin{aligned} & \text{If } \left(\text{aux}, 1^S, X^1, \dots, X^Q, \{E^1(X_i^1)\}_{i \in [N^1]}, \dots, \{E^Q(X_i^Q)\}_{i \in [N^Q]}, \right) \\ & \quad \approx_c \left(\text{aux}, 1^S, X^1, \dots, X^Q, \{\Delta_i^1\}_{i \in [N^1]}, \dots, \{\Delta_i^Q\}_{i \in [N^Q]} \right) \\ & \text{then } \left(\text{aux}, X^1, \dots, X^Q, \text{Lobf}_{\text{off}}, \text{Lobf}_{\text{on}}^1, \dots, \text{Lobf}_{\text{on}}^Q \right) \\ & \quad \approx_c \left(\text{aux}, X^1, \dots, X^Q, \text{Lobf}_{\text{off}}, \delta^1, \dots, \delta^Q \right), \end{aligned}$$

where

$$\begin{aligned} \Delta_i^k & \leftarrow \{0, 1\}^{L^k} \text{ for } k \in [Q], i \in [N^k], \\ (\text{Lobf}_{\text{off}}, \text{st}) & \leftarrow \text{LObfOff}(1^\lambda, 1^S), \text{dig}^k \leftarrow \text{LDigest}(1^\lambda, X^k), \text{Lobf}_{\text{on}}^k \leftarrow \text{LObfOn}(\text{dig}^k, E^k), \\ \delta^k & \leftarrow O_{\text{on}} \text{ for } k \in [Q], \text{ where } O_{\text{on}} \text{ is the co-domain of } \text{LObfOn} \text{ algorithm.} \end{aligned}$$

Remark 23. It is shown in [19, 65] that there is no laconic pPRIO scheme satisfying the above security for all general samplers. Therefore, when we use the security of laconic pPRIO, we invoke the security *with respect to a specific sampler class* that is induced by the respective applications. For simplicity, we sometimes will treat as if there was laconic pPRIO that is secure for all the samplers.

5.4.2 Construction

In this section we construct a laconic pPRIO scheme for any circuit.

Building Blocks. Below, we list the ingredients for our construction.

1. A pseudorandom function $\text{PRF} : \{0, 1\}^\lambda \times \{0, 1\}^\lambda \rightarrow \{0, 1\}^\lambda$ with key space, input space and output space as $\{0, 1\}^\lambda$. The input to the PRF will be in the form of $(i||j||b)$ where $i \in [N]$, $j \in [\ell]$, and $b \in \{0, 1\}$. Here, N and ℓ are some polynomial functions in λ . Since we have $2^\lambda > \text{poly}(\lambda)$, the input space of the PRF is large enough to accommodate such inputs with an appropriate encoding. It is known that PRF can be constructed from one-way functions.
2. A blind garbling scheme $\text{bGC} = (\text{Garble}, \text{bGC.Eval}, \text{bGC.Sim})$ (defined in Section 5.2.1) for any circuit. Without loss of generality, we assume the labels are in $\{0, 1\}^\lambda$. We also assume the randomness space of Garble algorithm to be $\{0, 1\}^\lambda$. If longer randomness is required by the algorithm, we can expand it by using a PRF. We require the scheme to be secure as per Definition 5.3 and Definition 5.4. We can construct bGC with the required properties assuming one-way functions (See Theorem 5.13).
3. A pPRIO scheme $\text{pPRIO}(\text{ObfOff}, \text{ObfOn}, \text{Eval})$. We require the scheme to be secure as per Definition 5.11. We can construct pPRIO with the required properties assuming evasive LWE and LWE (See Theorem 5.14).
4. A blind batch encryption scheme $\text{BBE} = \text{BBE}(\text{Setup}, \text{Gen}, \text{SingleEnc}, \text{SingleDec})$ with message space $\{0, 1\}^\lambda$. We assume the randomness space of SingleEnc algorithm to be $\{0, 1\}^\lambda$ and denote the ciphertext space by $\mathcal{CT}_{\text{BBE}} = \{0, 1\}^{\ell_{\text{BBE}}^{\text{ct}}}$. We require the scheme to be secure as per Definition 5.8 and Definition 5.9. We also require that the CRS is a uniformly random string. As for the efficiency requirement, we need the size of the CRS to be $|\text{crs}| = \text{poly}(\lambda, \log N, \log \ell)$ and the size of the single ciphertext to be $|\text{BBE.ct}| = \text{poly}(\lambda, \log N, \log \ell)$. We can construct BBE with the required properties assuming the LWE assumption (See Theorem 5.39).

Now, we describe our construction.

$\text{LDigest}(1^\lambda, X = \{X_i\}_{i \in [N]})$. The digest algorithm does the following.

- Retrieve N and ℓ from the input and run $\text{crs} \leftarrow \text{BBE.Setup}(1^\lambda, 1^{\ell N})$.
- Compute $h := \text{BBE.Gen}(\text{crs}, X_1 || \dots || X_N)$, where $X_1 || \dots || X_N \in \{0, 1\}^{\ell N}$ is the concatenation of the bit strings $X_1, \dots, X_N \in \{0, 1\}^\ell$.
- Output $\text{dig} := (\text{crs}, h, N, \ell)$.

$\text{LObfOff}(1^\lambda, 1^S)$. The offline obfuscate algorithm does the following.

- Generate $(\text{pPRIO.obf}_{\text{off}}, \text{pPRIO.st}) \leftarrow \text{pPRIO.ObfOff}(1^\lambda, 1^{\text{size}})$ where $\text{size} = \text{poly}(S, \lambda)$ is the maximum size of the circuit $C[\text{dig}, E, \text{sd}]$ defined in Figure 5.3 when the size of E is bounded by S .

- Output $\text{Lobf}_{\text{off}} := \text{pPRIO.obf}_{\text{off}}$ and $\text{st} := \text{pPRIO.st}$.

$\text{LobfOn}(\text{st}, \text{dig}, E)$. The online obfuscate algorithm does the following.

- Parse $\text{dig} = (\text{crs}, h, N, \ell)$ and $\text{st} = \text{pPRIO.st}$.
- Sample a PRF key $\text{sd} \leftarrow \{0, 1\}^\lambda$.
- Construct a circuit $C[\text{dig}, E, \text{sd}]$ as in Figure 5.3 and compute $\text{pPRIO.obf}_{\text{on}} \leftarrow \text{pPRIO.ObfOn}(\text{pPRIO.st}, C[\text{dig}, E, \text{sd}])$.
- Output $\text{Lobf}_{\text{on}} := (\text{crs}, \text{pPRIO.obf}_{\text{on}})$.

$\text{LEval}(X, \text{Lobf})$. The evaluation algorithm does the following.

- Parse $X = \{X_i\}_{i \in [N]}$ and $\text{Lobf} = (\text{Lobf}_{\text{off}} = \text{pPRIO.obf}_{\text{off}}, \text{Lobf}_{\text{on}} = (\text{crs}, \text{pPRIO.obf}_{\text{on}}))$.
- Set $\text{pPRIO.Obf} := (\text{pPRIO.obf}_{\text{off}}, \text{pPRIO.obf}_{\text{on}})$ and run $\text{pPRIO.Eval}(\text{pPRIO.Obf}, i) \rightarrow y_i$ for $i \in [N]$.
- Parse $y_i = (\{\text{BBE.ct}_{i,j}\}_{j \in [\ell]}, \tilde{E}_i)$ for each $i \in [N]$.
- Compute $\text{lab}_{i,j} \leftarrow \text{BBE.SingleDec}(\text{crs}, X, \ell(i-1) + j, \text{BBE.ct}_{i,j})$ for each $i \in [N]$ and $j \in [\ell]$ and set $\text{lab}_i := \{\text{lab}_{i,j}\}$.
- Compute $z_i = \text{bGC.Eval}(\text{lab}_i, \tilde{E}_i)$ for $i \in [N]$.
- Output $\{z_i\}_{i \in [N]}$.

Correctness. For $X = \{X_i\}_{i \in [N]}$ and $\text{Lobf} = (\text{Lobf}_{\text{off}} = \text{pPRIO.obf}_{\text{off}}, \text{Lobf}_{\text{on}} = (\text{crs}, \text{pPRIO.obf}_{\text{on}}))$, we have $\text{pPRIO.obf}_{\text{on}} \leftarrow \text{pPRIO.ObfOn}(\text{pPRIO.st}, C[\text{dig}, E, \text{sd}])$. From the correctness of pPRIO and the definition of $C[\text{dig}, E, \text{sd}]$, for each $i \in [N]$ we get

$$\text{pPRIO.Eval}(\text{pPRIO.Obf}, i) = y_i = (\{\text{BBE.ct}_{i,j}\}_{j \in [\ell]}, \tilde{E}_i)$$

where $\text{BBE.ct}_{i,j} \leftarrow \text{BBE.SingleEnc}(\text{crs}, h, (i-1)\ell + j, (\text{lab}_{i,j,0}, \text{lab}_{i,j,1}); S_{i,j})$ for $(\{\text{lab}_{i,j,b}\}_{j \in [\ell], b \in \{0,1\}}, \tilde{E}_i) \leftarrow \text{Garble}(1^\lambda, E_i; R_i)$.

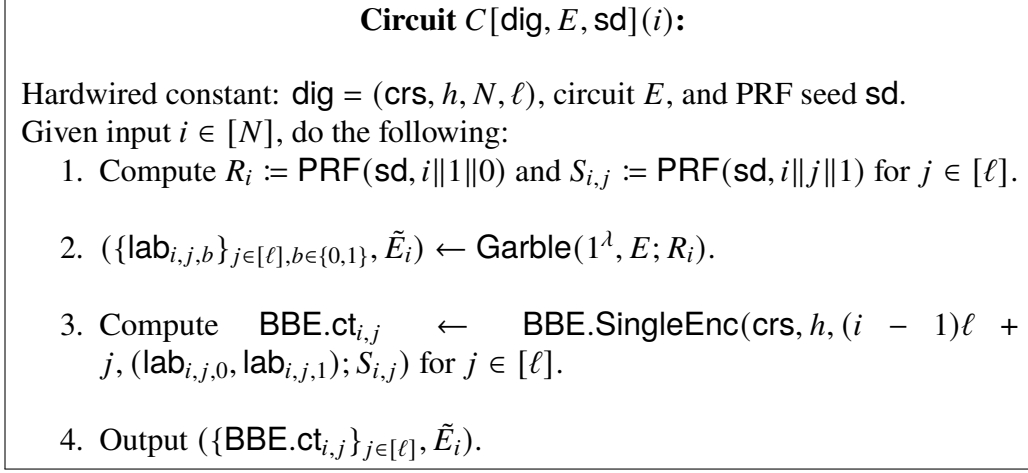


Figure 5.3: The Circuit $C[\text{dig}, E, \text{sd}](i)$.

Next, from the perfect correctness of the BBE scheme we get

$$\text{BBE.SingleDec}(\text{crs}, X, \ell(i - 1) + j, \text{BBE.ct}_{i,j}) = \text{lab}_{i,j} \text{ for } i \in [N], j \in [\ell]$$

Next, we set $\text{lab}_i := \{\text{lab}_{i,j}\}$, which are the labels corresponding to input X_i for $i \in [N]$.

Now, from the correctness of the bGC scheme and definition of E_i , for each $i \in [N]$ we get

$$\text{bGC.Eval}(\text{lab}_i, \tilde{E}_i) = E(X_i)$$

and hence the correctness.

Efficiency. First, we note the following.

1. Instantiating the BBE scheme as in Theorem 5.39, we have $|\text{crs}| = \text{poly}(\lambda, \log N, \log \ell)$ and $|\text{BBE.ct}| = \text{poly}(\lambda, \log N, \log \ell)$, which are bounded by a fixed polynomial $\text{poly}(\lambda)$ for any polynomially bounded N and ℓ .
2. Instantiating the pPRIO scheme as in Theorem 5.14, we have $|\text{pPRIO.Obf}_{\text{off}}| = \text{poly}(S, \lambda)$, $|\text{Obf}_{\text{on}}| = \text{poly}(S, \lambda)$.

From the above observations, our laconic pPRIO scheme satisfies

$$|\text{dig}| = O(\lambda), \quad |\text{Lobf}_{\text{off}}| = \text{poly}(S, \lambda), \quad |\text{Lobf}_{\text{on}}| = \text{poly}(S, \lambda).$$

We formalise the instantiation using the following theorem. The security of our laconic pPRIO will be proven in Theorem 5.22.

Theorem 5.21. *Assuming LWE and evasive LWE assumptions, there exists a secure (Definition 5.17) laconic pPRIO scheme satisfying*

$$|\text{dig}| = O(\lambda), \quad |\text{Lobf}_{\text{off}}| = \text{poly}(S, \lambda), \quad |\text{Lobf}_{\text{on}}| = \text{poly}(S, \lambda).$$

where $\text{dig} \leftarrow \text{LDigest}(1^\lambda, X = \{X_i\}_{i \in [N]})$,
 $(\text{Lobf}_{\text{off}}, \text{Lobf}_{\text{on}}) \leftarrow \text{LObfuscate}(1^\lambda, \text{dig}, E)$ for circuit $E : \{0, 1\}^\ell \rightarrow \{0, 1\}^L$ whose size is bounded by $S = S(\lambda)$.

5.4.3 Security Proof

Theorem 5.22. *Let $SC_{l\text{-pPRIO}}$ be a sampler class for laconic pPRIO. Assume pPRIO is secure (Definition 5.11) with respect to the sampler class that contains all $\text{Samp}_{\text{pPRIO}}(1^\lambda)$, induced by $\text{Samp}_{l\text{-pPRIO}} \in SC_{l\text{-pPRIO}}$, as defined in Section 5.4.3, BBE satisfies security (Definition 5.8) and strong blindness (Definition 5.9), bGC satisfies security (Definition 5.3) and blindness (Definition 5.4) and PRF is secure. Then the above construction of laconic pPRIO satisfies security as defined in Definition 5.17.*

Proof. Let us consider a sampler $\text{Samp}_{l\text{-pPRIO}}$ that outputs

$$\left(\text{aux}, \{X^k = \{X_i^k \in \{0, 1\}^{\ell^k}\}_{i \in [N^k]}\}_{k \in [Q]}, \{E^k\}_{k \in [Q]} \right).$$

To prove the theorem, we show that

$$\begin{aligned} & \left(\text{aux}, \text{Lobf}_{\text{off}} = \text{pPRIO.obf}_{\text{off}}, \{X^k, \text{Lobf}_{\text{on}}^k = (\text{crs}^k, \text{pPRIO.obf}_{\text{on}}^k)\}_{k \in [Q]} \right) \\ & \approx_c \left(\text{aux}, \text{Lobf}_{\text{off}} = \text{pPRIO.obf}_{\text{off}}, \{X^k, \delta^k \leftarrow \mathcal{O}_{\text{on}}\}_{k \in [Q]} \right) \end{aligned} \quad (5.29)$$

holds assuming

$$\left(\text{aux}, \{X^k, \{E^k(X_i^k)\}_{i \in [N^k]}\}_{k \in [Q]} \right) \approx_c \left(\text{aux}, \{X^k, \{\Delta_i^k \leftarrow \{0, 1\}^{L^k}\}_{i \in [N^k]}\}_{k \in [Q]} \right) \quad (5.30)$$

where O_{on} is the co-domain of LObfOn algorithm. Recalling that each crs^k is a random string, it suffices to prove that $\{\text{pPRIO.obf}_{\text{on}}^k\}_k$ is pseudorandom.

We invoke the security of pPRIO scheme with respect to a sampler $\text{Samp}_{\text{pPRIO}}(1^\lambda)$ that outputs

$$\left(1^{N^1 + \dots + N^Q}, \text{aux}_{\text{pPRIO}} = (\text{aux}, \{\text{crs}^k\}_{k \in [Q]}, \{X^k\}_{k \in [Q]}), \left\{ C^k[\text{dig}^k, E^k, \text{sd}^k] \right\}_{k \in [Q]} \right)$$

where $\text{crs}^k \leftarrow \text{BBE.Setup}(1^\lambda, 1^{\ell^k N^k})$, $h^k = \text{BBE.Gen}(\text{crs}^k, X_1^k \parallel \dots \parallel X_{N^k}^k)$, $\text{dig}^k = (\text{crs}^k, h^k, N^k, \ell^k)$, and $\text{sd}^k \leftarrow \{0, 1\}^\lambda$. From the security of pPRIO , we can see that it suffices to prove

$$\begin{aligned} & \left(\text{aux}, \{\text{crs}^k\}_{k \in [Q]}, \{X^k\}_{k \in [Q]}, \left\{ C^k[\text{dig}^k, E^k, \text{sd}^k](i) \right\}_{k \in [Q], i \in [N^k]} \right) \\ & \approx_c \left(\text{aux}, \{\text{crs}^k\}_{k \in [Q]}, \{X^k\}_{k \in [Q]}, \{\gamma_i^k\}_{k \in [Q], i \in [N^k]} \right) \end{aligned} \quad (5.31)$$

where $C^k[\text{dig}^k, E^k, \text{sd}^k](i) = (\{\text{BBE.ct}_{i,j}^k\}_{j \in [\ell^k]}, \tilde{E}_i^k)$ for $i \in [N^k], k \in [Q]$ and $\gamma_i^k \leftarrow \mathcal{CT}_{\text{BBE}}^{\ell^k} \times \{0, 1\}^{\ell_{\text{bGC}}^k}$. Here, ℓ_{bGC}^k is the length of the binary string \tilde{E}_i^k . To prove Equation (5.31), we consider the following sequence of games.

Hyb₀. This is the LHS distribution of Equation (5.31). Rearranging the terms and recalling the definition of the circuit, we can rewrite the distribution as

$$\left(\text{aux}, \left\{ \text{crs}^k, X^k, \left\{ \{\text{BBE.ct}_{i,j}^k\}_{j \in [\ell^k]}, \tilde{E}_i^k \right\}_{i \in [N^k]} \right\}_{k \in [Q]} \right)$$

where $(\{\text{lab}_{i,j,b}^k\}_{j \in [\ell^k], b \in \{0,1\}}, \tilde{E}_i^k) \leftarrow \text{Garble}(1^\lambda, E^k; R_i^k)$ and $\text{BBE.ct}_{i,j}^k \leftarrow \text{SingleEnc}(\text{crs}, h^k, (i-1)\ell^k + j, (\text{lab}_{i,j,0}^k, \text{lab}_{i,j,1}^k); S_{i,j}^k)$ for $R_i^k := \text{PRF}(\text{sd}^k, i \parallel 1 \parallel 0)$ and $S_{i,j}^k := \text{PRF}(\text{sd}^k, i \parallel j \parallel 1)$.

Hyb₁. This hybrid is same as the previous one except that we compute each R_i^k and $S_{i,j}^k$ differently. Concretely, we sample $R_i^k, S_{i,j}^k \leftarrow \{0, 1\}^\lambda$ for all $j \in [\ell^k], i \in [N^k]$ and $k \in [Q]$. This hybrid is computationally indistinguishable from the previous one due to the security of PRF.

Hyb₂. In this hybrid, we change how we compute $\text{BBE.ct}_{i,j}^k$. Namely, we set

$$\text{BBE.ct}_{i,j}^k \leftarrow \text{SingleEnc}(\text{crs}, h^k, (i-1)\ell^k + j, (\overline{\text{lab}}_{i,j,0}^k, \overline{\text{lab}}_{i,j,1}^k))$$

$$\text{where } \overline{\text{lab}}_{i,j,b}^k \begin{cases} = \text{lab}_{i,j,b}^k & \text{if } b = X_{i,j}^k \\ \leftarrow \{0, 1\}^\lambda & \text{otherwise} \end{cases}.$$

In the above, $X_{i,j}^k$ is the j -th bit of X_i^k . By the security of BBE, this hybrid is computationally indistinguishable from the previous one. To see this, observe that the decryption outputs the labels corresponding to the j -th bit of X_i^k and we only substitute $\overline{\text{lab}}_{i,j,b}^k \leftarrow \{0, 1\}^\lambda$ when $b \neq X_{i,j}^k$.

Hyb₃. In this hybrid, we change how we compute \tilde{E}_i^k and $\overline{\text{lab}}_{i,j,b}^k$. Namely, we set

$$(\{\text{lab}_{i,j}^k\}_{j \in [\ell^k]}, \tilde{E}_i^k) \leftarrow \text{bGC.Sim}(1^\lambda, 1^{|E^k|}, 1^{\ell^k}, E^k(X_i^k)) \quad \text{and} \quad \overline{\text{lab}}_{i,j,b}^k \begin{cases} = \text{lab}_{i,j}^k & \text{if } b = X_{i,j}^k \\ \leftarrow \{0, 1\}^\lambda & \text{otherwise} \end{cases}. \quad (5.32)$$

By the simulation security of bGC (Definition 5.3), this hybrid is computationally indistinguishable from the previous one. To see this, it suffices to observe that only the information of $\{\text{lab}_{i,j,X_{i,j}^k}^k\}_{i,j,k}$ is necessary for simulating Hyb₂ and the labels $\{\text{lab}_{i,j,1-X_{i,j}^k}^k\}_{i,j,k}$ are not necessary.

Hyb₄. In this hybrid, we change how we compute \tilde{E}_i^k and $\text{lab}_{i,j}^k$. Namely, we set

$$(\{\text{lab}_{i,j}^k\}_{j \in [\ell]}, \tilde{E}_i^k) \leftarrow \text{bGC.Sim}(1^\lambda, 1^{|E^k|}, 1^\ell, \Delta_i^k),$$

where $\Delta_i^k \leftarrow \{0, 1\}^{L^k}$ is chosen uniformly at random. This hybrid is computationally indistinguishable from the previous one by Equation (5.30). To see this note that Equation (5.30) implies $\{E^k(X_i^k)\}_{i \in [N^k]} \approx_c \{\Delta_i^k \leftarrow \{0, 1\}^{L^k}\}_{i \in [N^k]}$, for $k \in [Q]$, given $\text{aux}, \{X^k\}_{k \in [Q]}$.

Hyb₅. In this hybrid, we sample $\text{lab}_{i,j}^k$ and \tilde{E}_i^k as random strings. In particular, we sample $\text{lab}_{i,j}^k \leftarrow \{0, 1\}^\lambda$ and $\tilde{E}_i^k \leftarrow \{0, 1\}^{\ell_{\text{bGC}}^k}$. By the blindness of bGC scheme (Definition 5.4), this hybrid is computationally indistinguishable from the previous one. To see this, note that in the previous hybrid the simulator **bGC.Sim** takes as input Δ_i^k which is a uniformly random string and thus by the blindness property of bGC we can replace the output of **bGC.Sim** by a completely random string.

The view of the adversary in this hybrid is as follows.

$$\left(\text{aux}, \left\{ \text{crs}^k, X^k, \left\{ \{\text{BBE.ct}_{i,j}^k\}_{j \in [\ell]}, \tilde{E}_i^k \right\}_{i \in [N^k]} \right\}_{k \in [Q]} \right)$$

where $\text{BBE.ct}_{i,j}^k \leftarrow \text{SingleEnc}(\text{crs}, h^k, (i-1)\ell^k + j, (\overline{\text{lab}}_{i,j,0}^k, \overline{\text{lab}}_{i,j,1}^k))$ for $\overline{\text{lab}}_{i,j,0}^k \leftarrow \{0, 1\}^\lambda, \overline{\text{lab}}_{i,j,1}^k \leftarrow \{0, 1\}^\lambda$ and $\tilde{E}_i^k \leftarrow \text{SIM}_{\text{circ}}$.

Hyb₆. In this hybrid, we replace $\text{BBE.ct}_{i,j}^k$ with a random string. Namely, we sample $\text{BBE.ct}_{i,j}^k \leftarrow \mathcal{CT}_{\text{BBE}}$ for all i, j, k . This hybrid is indistinguishable from previous one by strong blindness (Definition 5.9) property of the BBE scheme. To see this, note that $\text{BBE.ct}_{i,j}^k$ encrypts random strings $\overline{\text{lab}}_{i,j,0}^k, \overline{\text{lab}}_{i,j,1}^k \leftarrow \{0, 1\}^\lambda$ in the previous hybrid and thus the blindness property allows us to replace each $\text{BBE.ct}_{i,j}^k$ with a random string.

The view of the adversary in this hybrid is as follows.

$$\left(\text{aux}, \left\{ \text{crs}^k, X^k, \left\{ \{\text{BBE.ct}_{i,j}^k \leftarrow \mathcal{CT}_{\text{BBE}}\}_{j \in [\ell]}, \tilde{E}_i^k \leftarrow \{0, 1\}^{\ell_{\text{bGC}}^k} \right\}_{i \in [N^k]} \right\}_{k \in [Q]} \right).$$

Rearranging the terms, we can observe that the distribution in **Hyb₆** corresponds to the RHS distribution of Equation (5.31). This concludes the proof of Theorem 5.22. \blacksquare

5.5 PARTIAL-HIDING prFE FOR UNBOUNDED DEPTH WITH OPTIMAL PARAMETERS

In this section, we extend the notion of prFE to introduce partially-hiding prFE, where the part of the input to the circuit can be public. We then construct partially hiding prFE with short parameter size in this section from several ingredients, which are all implied by evasive LWE and LWE. Our construction of partially hiding prFE will be used in Section 5.6 and 5.7.

5.5.1 Definition

In this section we give the definitions for partial-hiding functional encryption for pseudorandom functionalities. Consider a circuit class $\{C_{\text{prm}} = \{C : \mathcal{X}_{\text{pub}} \times \mathcal{X}_{\text{priv}} \rightarrow \mathcal{Y}\}\}_{\text{prm}}$ where $C \in C_{\text{prm}}$ takes as input a string $x = (x_{\text{pub}}, x_{\text{priv}}) \in \mathcal{X}_{\text{pub}} \times \mathcal{X}_{\text{priv}}$ and outputs $C(x) \in \mathcal{Y}$.

Syntax. A partial-hiding functional encryption for parameterized circuits C_{prm} by prm consists of four polynomial time algorithms ($\text{Setup}, \text{KeyGen}, \text{Enc}, \text{Dec}$) defined as follows.

$\text{Setup}(1^\lambda, \text{prm}) \rightarrow (\text{mpk}, \text{msk})$. The setup algorithm takes as input the security parameter λ and a parameter prm and outputs a master public key mpk and a master secret key msk . We assume w.l.o.g that msk includes mpk . We also assume that prm is implicitly input to all the algorithms below.

$\text{KeyGen}(\text{msk}, C) \rightarrow \text{sk}_C$. The key generation algorithm takes as input the master secret key msk and a circuit $C \in C_{\text{prm}}$ and it outputs a functional secret key sk_C .

$\text{Enc}(\text{mpk}, x = (x_{\text{pub}}, x_{\text{priv}})) \rightarrow \text{ct}$. The encryption algorithm takes as input the master public key mpk and an input $x \in \mathcal{X}_{\text{prm}}$ and outputs a ciphertext $\text{ct} \in \mathcal{CT}$, where \mathcal{CT} is the ciphertext space.

For the purpose of some applications, we consider a variant where the Enc algorithm can be decomposed into the following two phases.

$\text{EncOff}(\text{mpk}) \rightarrow (\text{ct}_{\text{off}}, \text{st})$. The offline encryption algorithm takes as input the security parameter λ and outputs offline part of the ciphertext ct_{off} and the state st .

$\text{EncOn}(\text{st}, x) \rightarrow \text{ct}_{\text{on}}$. The online encryption algorithm takes as input the state st and the input x and outputs the online part of the ciphertext ct_{on} .

The final output of ct is $\text{ct} = (\text{ct}_{\text{off}}, \text{ct}_{\text{on}})$.

$\text{Dec}(\text{mpk}, \mathbf{x}_{\text{pub}}, \text{sk}_C, C, \text{ct}) \rightarrow y$. The decryption algorithm takes as input the master public key mpk , the public input \mathbf{x}_{pub} , secret key sk_C , circuit C and a ciphertext ct and outputs $y \in \mathcal{Y}_{\text{prm}}$.

Definition 5.18 (Correctness). A PHprFE scheme is said to satisfy *perfect* correctness if for all prm , any input $x = (x_{\text{pub}}, x_{\text{priv}}) \in \mathcal{X}_{\text{pub}} \times \mathcal{X}_{\text{priv}}$ and circuit $C \in \mathcal{C}_{\text{prm}}$, we have

$$\Pr \left[\begin{array}{l} (\text{mpk}, \text{msk}) \leftarrow \text{Setup}(1^\lambda, \text{prm}), \text{sk}_C \leftarrow \text{KeyGen}(\text{msk}, C), \\ \text{Dec}(\text{mpk}, \mathbf{x}_{\text{pub}}, \text{sk}_C, C, \text{Enc}(\text{mpk}, x)) = C(x) \end{array} \right] = 1.$$

Definition 5.19 (Security). For a PHprFE scheme for circuit class $\{\mathcal{C}_{\text{prm}} = \{C : \mathcal{X}_{\text{pub}} \times \mathcal{X}_{\text{priv}} \rightarrow \mathcal{Y}\}\}_{\text{prm}}$ parameterized by $\text{prm} = \text{prm}(\lambda)$, let Samp be a PPT algorithm that on input 1^λ , outputs

$$(C^1, \dots, C^Q, x = (x_{\text{pub}}, x_{\text{priv}}), \text{aux} \in \{0, 1\}^*)$$

where Q is the number of key queries, $C^k \in \mathcal{C}_{\text{prm}}$ for $k \in [Q]$, $x = (x_{\text{pub}}, x_{\text{priv}}) \in \mathcal{X}_{\text{pub}} \times \mathcal{X}_{\text{priv}}$.

We define the following advantage functions:

$$\begin{aligned} \text{Adv}_{\mathcal{A}_0}^{\text{PRE}}(\lambda) &\stackrel{\text{def}}{=} \Pr[\mathcal{A}_0(\text{aux}, \{C^k, x_{\text{pub}}, C^k(x)\}_{k \in [Q]}) = 1] \\ &\quad - \Pr[\mathcal{A}_0(\text{aux}, \{C^k, x_{\text{pub}}, \Delta^k \leftarrow \mathcal{Y}\}_{k \in [Q]}) = 1] \end{aligned} \tag{5.33}$$

$$\begin{aligned} \text{Adv}_{\mathcal{A}_1}^{\text{POST}}(\lambda) &\stackrel{\text{def}}{=} \Pr[\mathcal{A}_1(\text{mpk}, \text{aux}, \{C^k, x_{\text{pub}}, \text{Enc}(\text{mpk}, x), \text{sk}_{C^k}\}_{k \in [Q]}) = 1] \\ &\quad - \Pr[\mathcal{A}_1(\text{mpk}, \text{aux}, \{C^k, x_{\text{pub}}, \delta \leftarrow \text{Sim}(1^\lambda), \text{sk}_{C^k}\}_{k \in [Q]}) = 1] \end{aligned} \quad (5.34)$$

where $(C^1, \dots, C^Q, x = (x_{\text{pub}}, x_{\text{priv}}), \text{aux} \in \{0, 1\}^*) \leftarrow \text{Samp}(1^\lambda)$, $(\text{mpk}, \text{msk}) \leftarrow \text{Setup}(1^\lambda, \text{prm})$ and $C\mathcal{T}$ is the ciphertext space. We say that a PHprFE scheme for circuit class C_{prm} is secure with respect to the sampler class \mathcal{SC} if for every PPT sampler $\text{Samp} \in \mathcal{SC}$, \mathcal{A}_1 and Sim , there exists another PPT \mathcal{A}_0 such that

$$\text{Adv}_{\mathcal{A}_0}^{\text{PRE}}(\lambda) \geq \text{Adv}_{\mathcal{A}_1}^{\text{POST}}(\lambda)/Q(\lambda) - \text{negl}(\lambda) \quad (5.35)$$

and $\text{Time}(\mathcal{A}_0) \leq \text{Time}(\mathcal{A}_1) \cdot Q(\lambda)$.

Remark 24 (prFE as a special PHprFE). We remark that prFE is a special case of PHprFE with $\mathbf{x}_{\text{pub}} = \perp$.

Next, we define the security notion that we require for the PHprFE variant where we decompose the Enc algorithm as $\text{Enc} = (\text{EncOff}, \text{EncOn})$ and reuse the same state output by EncOff multiple times for generating the online part of the ciphertexts. We require the online part of the ciphertexts to be pseudorandom, whereas the offline part may not be.

Definition 5.20 (Reusable Security). For a PHprFE scheme for circuit class $\{C_{\text{prm}} = \{C : \mathcal{X}_{\text{pub}} \times \mathcal{X}_{\text{priv}} \rightarrow \mathcal{Y}\}\}_{\text{prm}}$ parameterized by $\text{prm} = \text{prm}(\lambda)$, let Samp be a PPT algorithm that on input 1^λ , outputs

$$(C^1, \dots, C^{Q_{\text{key}}}, x^1 = (x_{\text{pub}}^1, x_{\text{priv}}^1), \dots, x^{Q_{\text{msg}}} = (x_{\text{pub}}^{Q_{\text{msg}}}, x_{\text{priv}}^{Q_{\text{msg}}}), \text{aux} \in \{0, 1\}^*)$$

where Q_{key} and Q_{msg} are the number of key queries and messages respectively, $x^j = (x_{\text{pub}}^j, x_{\text{priv}}^j) \in \mathcal{X}_{\text{pub}} \times \mathcal{X}_{\text{priv}}$ for $j \in [Q_{\text{msg}}]$, $C^k \in C_{\text{prm}}$ for $k \in [Q]$.

We define the following advantage functions:

$$\begin{aligned} \text{Adv}_{\mathcal{A}_0}^{\text{PRE}}(\lambda) &\stackrel{\text{def}}{=} \Pr[\mathcal{A}_0(\text{aux}, \{C^k, x_{\text{pub}}^j, C^k(x^j)\}_{j \in [Q_{\text{msg}}], k \in [Q_{\text{key}}]}) = 1] \\ &\quad - \Pr[\mathcal{A}_0(\text{aux}, \{C^k, x_{\text{pub}}^j, \Delta^{j,k} \leftarrow \mathcal{Y}\}_{j \in [Q_{\text{msg}}], k \in [Q_{\text{key}}]}) = 1] \end{aligned} \quad (5.36)$$

$$\begin{aligned} \text{Adv}_{\mathcal{A}_1}^{\text{POST}}(\lambda) &\stackrel{\text{def}}{=} \Pr[\mathcal{A}_1(\text{mpk}, \text{aux}, \text{ct}_{\text{off}}, \{x_{\text{pub}}^j, \text{ct}_{\text{on}}^j\}_{j \in [Q_{\text{msg}}]}, \{C^k, \text{sk}_{C^k}\}_{k \in [Q_{\text{key}}]}) = 1] \\ &\quad - \Pr[\mathcal{A}_1(\text{mpk}, \text{aux}, \text{ct}_{\text{off}}, \{x_{\text{pub}}^j, \delta^j\}_{j \in [Q_{\text{msg}}]}, \{C^k, \text{sk}_{C^k}\}_{k \in [Q_{\text{key}}]}) = 1] \end{aligned} \quad (5.37)$$

where $(C^1, \dots, C^{Q_{\text{key}}}, x^1 = (x_{\text{pub}}^1, x_{\text{priv}}^1), \dots, x^{Q_{\text{msg}}} = (x_{\text{pub}}^{Q_{\text{msg}}}, x_{\text{priv}}^{Q_{\text{msg}}}), \text{aux} \in \{0, 1\}^*) \leftarrow \text{Samp}(1^\lambda)$, $(\text{mpk}, \text{msk}) \leftarrow \text{Setup}(1^\lambda, \text{prm})$, $(\text{ct}_{\text{off}}, \text{st}) \leftarrow \text{EncOff}(1^\lambda)$, $\text{ct}_{\text{on}}^j \leftarrow \text{EncOn}(\text{st}, x^j)$, $\delta^j \leftarrow C\mathcal{T}_{\text{on}}$ for $j \in [Q_{\text{msg}}]$ and $C\mathcal{T}_{\text{on}}$ is the online part of the ciphertext space. We say that a PHprFE scheme for circuit class \mathcal{C}_{prm} is secure with respect to the sampler class \mathcal{SC} if for every PPT sampler $\text{Samp} \in \mathcal{SC}$ and \mathcal{A}_1 , there exists another PPT \mathcal{A}_0 such that

$$\text{Adv}_{\mathcal{A}_0}^{\text{PRE}}(\lambda) \geq \text{Adv}_{\mathcal{A}_1}^{\text{POST}}(\lambda)/Q(\lambda) - \text{negl}(\lambda) \quad (5.38)$$

and $\text{Time}(\mathcal{A}_0) \leq \text{Time}(\mathcal{A}_1) \cdot Q(\lambda)$.

Remark 25. Similar to Remark 19, when we use the security of PHprFE, we invoke the security *with respect to a specific sampler class* that is induced by the respective applications. For simplicity, we sometimes will treat as if there was PHprFE that is secure for all the samplers.

Remark 26. We remark that Definition 5.20 is stronger security notion than Definition 5.19, since the former collapses to the latter if we restrict the adversary so that $Q_{\text{msg}} = 1$ and set the simulator so that it runs $\text{EncOff}(\text{mpk}) \rightarrow (\text{ct}_{\text{off}}, \text{st})$, samples $\delta \leftarrow C\mathcal{T}_{\text{on}}$, and outputs $(\text{ct}_{\text{off}}, \delta)$.

Remark 27. We remark that Definition 5.20 is in the single-challenge flavor in that only single offline part of the challenge ciphertext is given to the adversary (though multiple online part of the ciphertext is given to it). While it does not seem to imply multi-challenge flavor of the security definition (See Remark 20), we only define this simpler version of the definition since it suffices for our applications and leave the extension to the multi-challenge flavor for the future work.

5.5.2 Construction

In this section we provide our construction of a PHprFE scheme for circuit family $C_{L_{\text{pub}}, L_{\text{priv}}} = \{C : \{0, 1\}^{L_{\text{pub}}(\lambda)} \times \{0, 1\}^{L_{\text{priv}}(\lambda)} \rightarrow \{0, 1\}\}$. Namely, the construction supports a class of circuits whose public and private input lengths are fixed and output is binary, but its size and depth are unbounded.

Building Blocks. Below, we list the ingredients for our construction.

1. A blind garbling scheme $\text{bGC} = (\text{Garble}, \text{Eval}, \text{bGC.Sim})$ (defined in Section 5.2.1) with decomposability (defined in Definition 5.5). Without loss of generality, we assume the labels are in $\{0, 1\}^\lambda$ and the random coins used by the algorithm $\{\text{Garble}_i\}_i$ is in $\{0, 1\}^\lambda$. The latter is for the sake of notational convenience and can be achieved by using a PRF to derive longer (pseudo-)random coins if needed. We also use $\mathcal{CT}_{\text{bGC}}^i$ to denote the co-domain of $\{\text{Garble}_i\}_i$ algorithm. We can construct bGC with the required properties assuming one-way functions (See Theorem 5.13).
2. A pseudorandom function $\text{PRF} : \{0, 1\}^\lambda \times \{0, 1\}^\lambda \rightarrow \{0, 1\}^\lambda$ with key space, input space and output space as $\{0, 1\}^\lambda$. It is known that PRF can be constructed from one-way functions.
3. A laconic pPRIO scheme $\text{LprIO} = (\text{LDigest}, \text{LObfOff}, \text{LObfOn}, \text{LEval})$. Without loss of generality, the random coins used by LObfOn is in $\{0, 1\}^\lambda$ and the state output by LObfOff is in $\{0, 1\}^\lambda$. The former can be satisfied by using a PRF. The latter can be achieved in two steps. We first let the state Lst to be the randomness used by LObfOff and then replace it with a string in $\{0, 1\}^\lambda$, which can be done by using a PRF. The length of the digest is of fixed polynomial in the security parameter and we denote it by $L_{\text{LDigest}}(\lambda)$. As we show in Theorem 5.21, such a laconic pPRIO can be constructed by assuming evasive LWE and LWE.
4. A prFE scheme $\text{prFE} = \text{prFE}(\text{Setup}, \text{KeyGen}, \text{Enc}, \text{Dec})$ for circuit class $C_{\text{inp}(\lambda), \text{dep}(\lambda), \text{out}(\lambda)}$ consisting of circuits with input length $\text{inp}(\lambda) = L_{\text{priv}} + L_{\text{LDigest}}(\lambda) + 4\lambda$, maximum depth $\text{dep}(\lambda)$ and output length $\text{out}(\lambda)$. We set the maximum depth and output length so that the circuit class supports the circuits $F_0[\mathbf{r}]$, $F_1[\mathbf{r}]$, and $F_2[\mathbf{r}, \text{dig}_C]$ defined as Figure 5.4, 5.5, and 5.6, respectively. We denote the information $(1^{\text{inp}(\lambda)}, 1^{\text{dep}(\lambda)}, 1^{\text{out}(\lambda)})$ specifying the circuit class by prm and the ciphertext space of the prFE scheme by $\mathcal{CT}_{\text{prFE}} = \{0, 1\}^{\ell_{\text{prFE}}^{\text{ct}}}$. Assuming LWE and evasive LWE, we can construct prFE with the required parameters (See Theorem 5.15).

For our construction, we set the following parameters

$\text{Setup}(1^\lambda, 1^{L_{\text{pub}}}, 1^{L_{\text{priv}}})$. The setup algorithm does the following.

- Run $(\text{prFE.mpk}, \text{prFE.msk}) \leftarrow \text{prFE.Setup}(1^\lambda, \text{prm})$.
- Output $\text{mpk} := \text{prFE.mpk}$ and $\text{msk} := \text{prFE.msk}$.

$\text{KeyGen}(\text{msk}, C)$. The key generation algorithm does the following.

- Parse $\text{msk} = \text{prFE.msk}$.
- Sample a string $\mathbf{r} \leftarrow \{0, 1\}^\lambda$.
- Run $\text{dig}_C \leftarrow \text{LDigest}(\{(i, C_i)\}_{i \in [|C|]})$, where C_i is the description of the i -th gate of C and $|C|$ is the number of gates in C . The description of C_i can be encoded into a string of length at most 4λ since it suffices to encode its index, two indices of the incoming wires, and the type of the gate.
- Construct circuits $F_0[\mathbf{r}]$, $F_1[\mathbf{r}]$, and $F_2[\mathbf{r}, \text{dig}_C]$ as in Figure 5.4, 5.5, and 5.6 respectively.
- Run $\begin{array}{ll} \text{prFE.sk}_0 & \leftarrow \text{prFE.KeyGen}(\text{prFE.msk}, F_0[\mathbf{r}]), \\ \text{prFE.sk}_1 & \leftarrow \text{prFE.KeyGen}(\text{prFE.msk}, F_1[\mathbf{r}]), \quad \text{and} \\ \text{prFE.sk}_2 & \leftarrow \text{prFE.KeyGen}(\text{prFE.msk}, F_2[\mathbf{r}, \text{dig}_C]). \end{array}$
- Output $\text{sk}_C := (\mathbf{r}, \text{prFE.sk}_0, \text{prFE.sk}_1, \text{prFE.sk}_2)$.

$\text{Enc}(\text{mpk}, \mathbf{x} = (\mathbf{x}_{\text{pub}}, \mathbf{x}_{\text{priv}}))$. The encryption algorithm does the following. We divide the algorithm into the following two steps.

$\text{EncOff}(\text{mpk})$. It takes as input $\text{mpk} = \text{prFE.mpk}$ and does the following.

- Run $(\text{Lobf}_{\text{off}}, \text{Lst}) \leftarrow \text{LObfOff}(1^\lambda, 1^S)$, where S is the maximum size of the circuits $E_{\text{pub}}[R_0]$ and $E_{\text{cir}}[R_0]$ defined in Figure 5.5 and Figure 5.6, respectively.
- Output $\text{st} := (\text{Lst}, \text{prFE.mpk})$ and $\text{ct}_{\text{off}} := \text{Lobf}_{\text{off}}$.

$\text{EncOn}(\text{st}, \mathbf{x} = (\mathbf{x}_{\text{pub}}, \mathbf{x}_{\text{priv}}))$. It does the following.

- Parse the input as $\text{st} \rightarrow (\text{Lst}, \text{prFE.mpk})$.
- Sample $\text{sd}_0, \text{sd}_1, \text{sd}_2 \leftarrow \{0, 1\}^\lambda$.

- Run $\text{dig}_{\mathbf{x}_{\text{pub}}} \leftarrow \text{LDigest}(1^\lambda, \{(i, \mathbf{x}_{\text{pub},i})\}_{i \in [L_{\text{pub}}]}),$ where $\mathbf{x}_{\text{pub},i} \in \{0, 1\}$ is the i -th bit of $\mathbf{x}_{\text{pub}}.$
- Run $\text{prFE.ct} \leftarrow \text{prFE.Enc}(\text{prFE.mpk}, (\text{dig}_{\mathbf{x}_{\text{pub}}}, \mathbf{x}_{\text{priv}}, \text{Lst}, \text{sd}_0, \text{sd}_1, \text{sd}_2)).$
- Output $\text{ct}_{\text{on}} := \text{prFE.ct}.$

The final output of $\text{Enc}(\text{mpk}, \mathbf{x})$ is $\text{ct} := (\text{Lobf}_{\text{off}}, \text{prFE.ct}).$

$\text{Dec}(\text{mpk}, \mathbf{x}_{\text{pub}}, \text{sk}_C, C, \text{ct}).$ It parses the input as $\text{mpk} = \text{prFE.mpk}, \text{ct} := (\text{Lobf}_{\text{off}}, \text{prFE.ct}), \text{sk}_C = (\mathbf{r}, \text{prFE.sk}_0, \text{prFE.sk}_1, \text{prFE.sk}_2)$ and does the following.

- Compute $\text{prFE.Dec}(\text{prFE.mpk}, \text{prFE.sk}_0, F_0[\mathbf{r}], \text{prFE.ct})$ and parse the output as $\{\widetilde{\text{lab}}'_i\}_{i \in [L_{\text{pub}}+1, L_{\text{pub}}+L_{\text{priv}}]}.$
- Compute $\text{Lobf}_{\text{on},1} = \text{prFE.Dec}(\text{prFE.mpk}, \text{prFE.sk}_1, F_1[\mathbf{r}], \text{prFE.ct}).$
- Compute $\text{Lobf}_{\text{on},2} = \text{prFE.Dec}(\text{prFE.mpk}, \text{prFE.sk}_1, F_2[\mathbf{r}, \text{dig}_C], \text{prFE.ct}).$
- Set $\text{Lobf}_1 := (\text{Lobf}_{\text{off}}, \text{Lobf}_{\text{on},1})$ and compute $\{\widetilde{\text{lab}}_i\}_{i \in [L_{\text{pub}}]} = \text{LEval}(\{(i, \mathbf{x}_{\text{pub},i})\}_{i \in [L_{\text{pub}}]}, \text{Lobf}_1).$
- Set $\text{Lobf}_2 := (\text{Lobf}_{\text{off}}, \text{Lobf}_{\text{on},2})$ and compute $\tilde{C} = \{\tilde{C}_i\}_{i \in [|C|]} = \text{LEval}(\{(i, C_i)\}_{i \in [|C|]}, \text{Lobf}_2).$
- Set $\text{lab}_i := \begin{cases} \widetilde{\text{lab}}_i & \text{if } i \in [L_{\text{pub}}] \\ \text{lab}'_i & \text{if } i \in [L_{\text{pub}}+1, L_{\text{pub}}+L_{\text{priv}}] \end{cases}$ for $i \in [L_{\text{pub}}+L_{\text{priv}}].$
- Compute $z = \text{Eval}(\tilde{C}, \{\text{lab}_i\}_{i \in [L_{\text{pub}}+L_{\text{priv}}]}).$
- Output $z.$

Correctness. We make the following observations.

- From the perfect correctness of prFE scheme and definition of $F_0[\mathbf{r}], F_1[\mathbf{r}], F_2[\mathbf{r}, \text{dig}_C],$ we get

$$\begin{aligned}
\{\text{lab}_{i, \mathbf{x}_{\text{priv}, i-L_{\text{pub}}}}\}_{i \in [L_{\text{pub}}+1, L_{\text{pub}}+L_{\text{priv}}]} &= \text{prFE.Dec}(\text{prFE.mpk}, \text{prFE.sk}_0, F_0[\mathbf{r}], \text{prFE.ct}) \\
&= F_0[\mathbf{r}](\text{dig}_{\mathbf{x}_{\text{pub}}}, \mathbf{x}_{\text{priv}}, \text{Lst}, \text{sd}_0, \text{sd}_1, \text{sd}_2) \\
&= \{\text{Garble}_{\text{inp}, i}(1^\lambda; R_0)\}_{i \in [L_{\text{pub}}+1, L_{\text{pub}}+L_{\text{priv}}]}
\end{aligned}$$

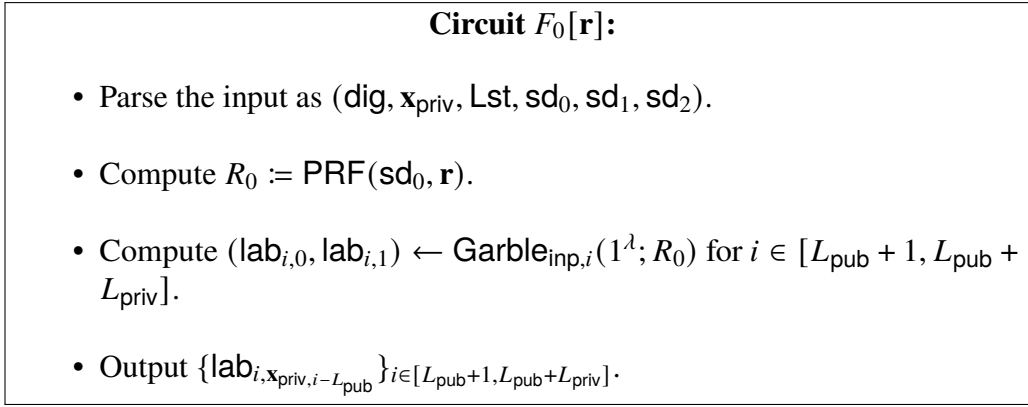


Figure 5.4: Circuit to compute garbled labels corresponding to \mathbf{x}_{priv} .

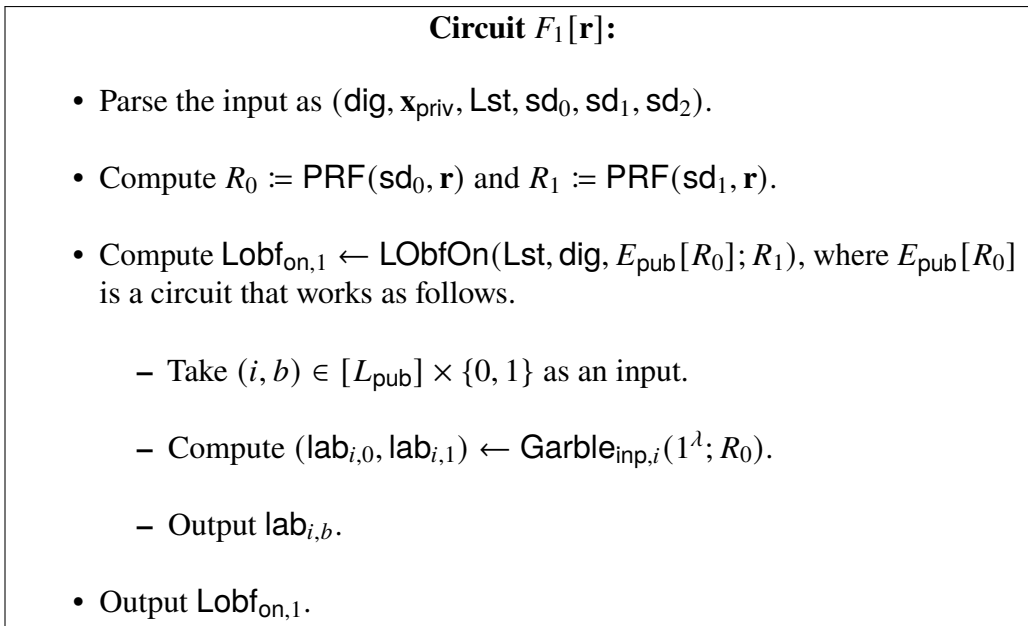


Figure 5.5: The Circuit $F_1[\mathbf{r}]$.

$$\begin{aligned}
 \text{Lobf}_{\text{on},1} &= \text{prFE.Dec}(\text{prFE.mpk}, \text{prFE.sk}_1, F_1[\mathbf{r}], \text{prFE.ct}) \\
 &= F_1[\mathbf{r}](\text{dig}_{\mathbf{x}_{\text{pub}}}, \mathbf{x}_{\text{priv}}, \text{Lst}, \text{sd}_0, \text{sd}_1, \text{sd}_2) \\
 &= \text{LObfOn}(\text{Lst}, \text{dig}_{\mathbf{x}_{\text{pub}}}, E_{\text{pub}}[R_0]; R_1),
 \end{aligned}$$

$$\begin{aligned}
 \text{Lobf}_{\text{on},2} &= \text{prFE.Dec}(\text{prFE.mpk}, \text{prFE.sk}_2, F_2[\mathbf{r}, \text{dig}_C], \text{prFE.ct}) \\
 &= F_2[\mathbf{r}, \text{dig}_C](\text{dig}_{\mathbf{x}_{\text{pub}}}, \mathbf{x}_{\text{priv}}, \text{Lst}, \text{sd}_0, \text{sd}_1, \text{sd}_2) \\
 &= \text{LObfOn}(\text{Lst}, \text{dig}_C, E_{\text{cir}}[R_0]; R_2),
 \end{aligned}$$

where $R_b := \text{PRF}(\text{sd}_b, \mathbf{r})$ for $b \in \{0, 1, 2\}$ and $E_{\text{pub}}[R_0]$ and $E_{\text{cir}}[R_0]$ are the circuits as defined in Figure 5.5 and Figure 5.6, respectively.

Circuit $F_2[\mathbf{r}, \text{dig}_C]$:

- Parse the input as $(\text{dig}, \mathbf{x}_{\text{priv}}, \text{Lst}, \text{sd}_0, \text{sd}_1, \text{sd}_2)$.
- Compute $R_0 := \text{PRF}(\text{sd}_0, \mathbf{r})$ and $R_2 := \text{PRF}(\text{sd}_2, \mathbf{r})$.
- Compute $\text{Lobf}_{\text{on},2} \leftarrow \text{LObfOn}(\text{Lst}, \text{dig}_C, E_{\text{cir}}[R_0]; R_2)$, where $E_{\text{cir}}[R_0]$ is a circuit that works as follows.
 - Take $(i, G) \in [|C|] \times \{0, 1\}^{4\lambda}$ as an input, where G encodes the information of a gate.
 - Compute $\tilde{G} \leftarrow \text{Garble}_i(1^\lambda, G; R_0)$.
 - Output \tilde{G} .
- Output $\text{Lobf}_{\text{on},2}$.

Figure 5.6: The Circuit $F_2[\mathbf{r}, \text{dig}_C]$.

- Next, parsing $\text{Lobf}_1 := (\text{Lobf}_{\text{off}}, \text{Lobf}_{\text{on},1})$, we get

$$\begin{aligned}
 \{\widetilde{\text{lab}}_i\}_{i \in [L_{\text{pub}}]} &= \text{LEval}(\{(i, \mathbf{x}_{\text{pub},i})\}_{i \in [L_{\text{pub}}]}, \text{Lobf}_1) \\
 &= \{E_{\text{pub}}[R_0](i, \mathbf{x}_{\text{pub},i})\}_{i \in [L_{\text{pub}}]} \\
 &= \text{lab}_{i, \mathbf{x}_{\text{pub},i}}
 \end{aligned}$$

where $(\text{lab}_{i,0}, \text{lab}_{i,1}) \leftarrow \text{Garble}_{\text{inp},i}(1^\lambda; R_0)$ and

$$\begin{aligned}
 \tilde{C} &= \{\tilde{C}_i\}_{i \in [|C|]} \\
 &= \text{LEval}(\{(i, C_i)\}_{i \in [|C|]}, \text{Lobf}_2) \\
 &= \{E_{\text{cir}}[R_0](i, C_i)\}_{i \in [|C|]} \\
 &= \{\text{Garble}_i(1^\lambda, C_i; R_0)\}_{i \in [|C|]}
 \end{aligned}$$

from the perfect correctness of the laconic pPRIO scheme and the definition of $E_{\text{pub}}[R_0]$ and $E_{\text{cir}}[R_0]$.

- Next, from the perfect correctness of the garbling scheme bGC, we have, for $\text{lab}_i = (\text{lab}_{\mathbf{x}_{\text{pub}}}, \text{lab}_{\mathbf{x}_{\text{priv}}})$,

$$\text{Eval}(\tilde{C}, \{\text{lab}_i\}_{i \in [L_{\text{pub}}+L_{\text{priv}}]}) = C(\mathbf{x}_{\text{pub}}, \mathbf{x}_{\text{priv}})$$

and hence the perfect correctness of the PHprFE scheme.

Efficiency Here, we show

$$|\text{mpk}| = \text{poly}(\lambda, L_{\text{priv}}), \quad |\text{sk}_C| = \text{poly}(\lambda, L_{\text{priv}}), \quad |\text{ct}| = \text{poly}(\lambda, L_{\text{priv}}).$$

To do so, we first show that the sizes of the circuits $F_0[\mathbf{r}]$, $F_1[\mathbf{r}]$, and $F_2[\mathbf{r}, \text{dig}_C]$ are all bounded by $\text{poly}(\lambda, L_{\text{priv}})$.

- In $F_0[\mathbf{r}]$, the computation of the PRF and $\text{Garble}_{\text{inp},i}$ are performed, where both of them can be implemented by circuits of size $\text{poly}(\lambda)$. Since the latter is repeated for L_{priv} times, the overall size of $F_0[\mathbf{r}]$ is $\text{poly}(\lambda, L_{\text{priv}})$.
- To bound the size of $F_1[\mathbf{r}]$, we first bound the size of $E_{\text{pub}}[R_0]$. $E_{\text{pub}}[R_0]$ performs the computation of $\text{Garble}_{\text{inp},i}$ on input $i \in [L_{\text{pub}}]$ (in binary), which can be implemented by a circuit of size $\text{poly}(\log L_{\text{pub}}, \lambda)$. This can be further bounded by $\text{poly}(\lambda)$ since $L_{\text{pub}} \leq 2^\lambda$. We then observe that $F_1[\mathbf{r}]$ consists of the evaluation of two PRF values and LObfOn on input Lst , dig , and $E_{\text{pub}}[R_0]$. We can bound the input length to LObfOn by a fixed polynomial, since we have $|\text{Lst}| = \lambda$ and $|\text{dig}| = \text{poly}(\lambda)$ and LObfOn runs in polynomial time in the input length, its size is bounded by $\text{poly}(\lambda)$. Therefore, the overall size of $F_1[\mathbf{r}]$ is $\text{poly}(\lambda, L_{\text{priv}})$, where we take into account the input \mathbf{x}_{priv} , which is ignored in the computation.
- To bound the size of $F_2[\mathbf{r}, \text{dig}_C]$, we first bound the size of $E_{\text{cir}}[R_0]$. $E_{\text{cir}}[R_0]$ performs the computation of Garble_i on input $i \in [|C|]$ (in binary) and $G \in \{0, 1\}^{4\lambda}$, which can be implemented by a circuit of size $\text{poly}(\log |C|, \lambda)$. This can be further bounded by $\text{poly}(\lambda)$, since $|C| \leq 2^\lambda$. Similarly to the case of $F_1[\mathbf{r}]$, the size of $F_2[\mathbf{r}]$ can be bounded by $\text{poly}(\lambda, L_{\text{priv}})$.

We then move to discuss the size of the parameters.

- We can bound $|\text{mpk}| = |\text{prFE.mpk}|$ by $\text{poly}(\lambda, \text{inp}, \text{dep}, \text{out})$, since it is output by $\text{prFE.Setup}(1^\lambda, \text{prm} = (1^{\text{inp}}, 1^{\text{dep}}, 1^{\text{out}}))$. We have $\text{inp}(\lambda) \leq L_{\text{priv}} + L_{\text{LDigest}}(\lambda) + 4\lambda \leq \text{poly}(\lambda, L_{\text{priv}})$ and $\text{dep}(\lambda)$ and $\text{out}(\lambda)$ are bounded by the maximum size of the circuits E_{cir} and E_{pub} , which in turn is bounded by $\text{poly}(\lambda, L_{\text{priv}})$. We therefore have $|\text{mpk}| = \text{poly}(\lambda, L_{\text{priv}})$.
- We have $|\text{sk}_C| = \lambda + |\text{prFE.sk}_0| + |\text{prFE.sk}_1| + |\text{prFE.sk}_2|$. We can bound the size of $|\text{prFE.sk}_0|$ by $\text{poly}(\lambda, L_{\text{priv}})$, since it is generated by $\text{prFE.KeyGen}(\text{prFE.msk}, F_0[\mathbf{r}])$, where the input length to prFE.KeyGen is bounded by $\text{poly}(\lambda, L_{\text{priv}})$. Similarly, we can bound $|\text{prFE.sk}_1|$ and $|\text{prFE.sk}_2|$ by $\text{poly}(\lambda, L_{\text{priv}})$. Therefore, we have $|\text{sk}_C| = \text{poly}(\lambda, L_{\text{priv}})$.
- We have $|\text{ct}| = |\text{LObf}_{\text{off}}| + |\text{prFE.ct}|$. We first bound $|\text{LObf}_{\text{off}}|$. From the above discussion, we have $S = \max\{|E_{\text{pub}}[R_0]|, |E_{\text{cir}}[R_0]|\} = \text{poly}(\lambda, L_{\text{priv}})$. This implies $|\text{LObf}_{\text{off}}| = \text{poly}(\lambda, L_{\text{priv}})$, since LObf_{off} is output by $\text{LObfOff}(1^\lambda, 1^S)$. We then bound $|\text{prFE.ct}|$, which is output by

$\text{prFE.ct} \leftarrow \text{prFE.Enc}(\text{prFE.mpk}, (\text{dig}_{\mathbf{x}_{\text{pub}}}, \mathbf{x}_{\text{priv}}, \text{Lst}, \text{sd}_0, \text{sd}_1, \text{sd}_2))$. Since the input length of prFE.Enc is $\text{inp}(\lambda) = L_{\text{priv}} + L_{\text{LDigest}}(\lambda) + 4\lambda = \text{poly}(\lambda, L_{\text{priv}})$. We therefore have $|\text{ct}| = |\text{Lobf}_{\text{off}}| + |\text{prFE.ct}| = \text{poly}(\lambda, L_{\text{priv}})$.

5.5.3 Security

Before proving the security of our scheme, we prove the following useful lemma. The lemma essentially says that if a part of the auxiliary information is pseudorandom in the pre-condition distribution, then it is pseudorandom in the corresponding post-condition distribution where we apply laconic pPRIO security. A conceptually similar lemma is proven in Lemma 3.4 of [22] in the context of evasive LWE.

Lemma 5.23. *Let $\text{LprIO} = (\text{LDigest}, \text{LObfuscate}, \text{LEval})$ be a laconic pPRIO scheme and Samp be a PPT algorithm that takes as input 1^λ and outputs*

$$\left(\text{aux} = (\text{aux}_1, \text{aux}_2) \in \{0, 1\}^* \times \mathcal{X}, 1^S, X^1 = \{X_i^1\}_{i \in [N^1]}, \dots, X^Q = \{X_i^Q\}_{i \in [N^Q]}, E^1, \dots, E^Q \right)$$

for some set \mathcal{X} . Here $X_i^k \in \{0, 1\}^{\ell^k}$, $E^k : \{0, 1\}^{\ell^k} \rightarrow \{0, 1\}^{L^k}$ and $|E^k| \leq S$ for $k \in [Q], i \in [N^k]$.

Let us assume that

$$\begin{aligned} & \left((\text{aux}_1, \text{aux}_2), 1^S, X^1, \dots, X^Q, \{E^1(X_i^1)\}_{i \in [N^1]}, \dots, \{E^Q(X_i^Q)\}_{i \in [N^Q]}, \right) \\ & \approx_c \left((\text{aux}_1, \mathbf{x}), 1^S, X^1, \dots, X^Q, \{\Delta_i^1\}_{i \in [N^1]}, \dots, \{\Delta_i^Q\}_{i \in [N^Q]} \right) \end{aligned}$$

holds for $\mathbf{x} \leftarrow \mathcal{X}$, $\Delta_i^k \leftarrow \{0, 1\}^{L^k}$ for $k \in [Q], i \in [N^k]$ and also assume the security of LprIO with respect to Samp . We then have

$$\begin{aligned} & \left((\text{aux}_1, \text{aux}_2), X^1, \dots, X^Q, \text{Lobf}_{\text{off}}, \text{Lobf}_{\text{on}}^1, \dots, \text{Lobf}_{\text{on}}^Q \right) \\ & \approx_c \left((\text{aux}_1, \mathbf{x}), X^1, \dots, X^Q, \text{Lobf}_{\text{off}}, \delta^1, \dots, \delta^Q \right), \end{aligned} \quad (5.39)$$

where $(\text{Lobf}_{\text{off}}, \text{st}) \leftarrow \text{LObfOff}(1^\lambda, 1^S)$, $\text{dig}^k \leftarrow \text{LDigest}(1^\lambda, X^k)$, $\text{Lobf}_{\text{on}}^k \leftarrow \text{LObfOn}(\text{dig}^k, E^k)$, $\delta^k \leftarrow O_{\text{on}}$ for $k \in [Q]$.

Proof. From the assumption, we have

$$\left((\mathbf{aux}_1, \mathbf{aux}_2), 1^S, X^1, \dots, X^Q, \{E^j(X_i^j)\}_{i \in [N^j], j \in [Q]} \right) \approx_c \left((\mathbf{aux}_1, \mathbf{x}), 1^S, X^1, \dots, X^Q, \{\Delta_i^j\}_{i \in [N^j], j \in [Q]} \right) \quad (5.40)$$

which implies $(\mathbf{aux}_1, \mathbf{aux}_2, 1^S, X^1, \dots, X^Q) \approx_c (\mathbf{aux}_1, \mathbf{x}, 1^S, X^1, \dots, X^Q)$. This further implies

$$(\mathbf{aux}_1, \mathbf{aux}_2, 1^S, X^1, \dots, X^Q, \{\Delta_i^j\}_{i \in [N^j], j \in [Q]}) \approx_c (\mathbf{aux}_1, \mathbf{x}, 1^S, X^1, \dots, X^Q, \{\Delta_i^j\}_{i \in [N^j], j \in [Q]}) \quad (5.41)$$

since adding independently sampled random terms Δ_i^j does not make the task of distinguishing the distributions easier. Equation (5.40) and Equation (5.41) implies

$$(\mathbf{aux}_1, \mathbf{aux}_2, 1^S, X^1, \dots, X^Q, \{E^j(X_i^j)\}_{i \in [N^j], j \in [Q]}) \approx_c (\mathbf{aux}_1, \mathbf{aux}_2, 1^S, X^1, \dots, X^Q, \{\Delta_i^j\}_{i \in [N^j], j \in [Q]}).$$

Applying LprIO security definition with respect to Samp, we get

$$\left(\mathbf{aux}_1, \mathbf{aux}_2, X^1, \dots, X^Q, \text{Lobf}_{\text{off}}, \text{Lobf}_{\text{on}}^1, \dots, \text{Lobf}_{\text{on}}^Q \right) \approx_c \left(\mathbf{aux}_1, \mathbf{aux}_2, X^1, \dots, X^Q, \text{Lobf}_{\text{off}}, \delta^1, \dots, \delta^Q \right). \quad (5.42)$$

Next, from $(\mathbf{aux}_1, \mathbf{aux}_2, 1^S, X^1, \dots, X^Q) \approx_c (\mathbf{aux}_1, \mathbf{x}, 1^S, X^1, \dots, X^Q)$ we have

$$\left(\mathbf{aux}_1, \mathbf{aux}_2, X^1, \dots, X^Q, \text{Lobf}_{\text{off}}, \delta^1, \dots, \delta^Q \right) \approx_c \left(\mathbf{aux}_1, \mathbf{x}, X^1, \dots, X^Q, \text{Lobf}_{\text{off}}, \delta^1, \dots, \delta^Q \right), \quad (5.43)$$

since Lobf_{off} can be sampled using 1^S and $\{\delta^j\}_{j \in [Q]}$ can be sampled independently.

From Equation (5.42) and Equation (5.43) we deduce

$$\left(\mathbf{aux}_1, \mathbf{aux}_2, X^1, \dots, X^Q, \text{Lobf}_{\text{off}}, \text{Lobf}_{\text{on}}^1, \dots, \text{Lobf}_{\text{on}}^Q \right) \approx_c \left(\mathbf{aux}_1, \mathbf{x}, X^1, \dots, X^Q, \text{Lobf}_{\text{off}}, \delta^1, \dots, \delta^Q \right).$$

hence the lemma. ■

The following theorem asserts the security of our construction of PHprFE scheme. This in particular implies that the construction satisfies the security notion as per Definition 5.19.

Theorem 5.24. *The above construction satisfies reusable security as per Definition 5.20.*

Proof. Consider a sampler Samp that generates the following:

1. **Key Queries.** It issues Q_{key} key queries $C^1, \dots, C^{Q_{\text{key}}}$.
2. **Ciphertext Queries.** It issues messages $\mathbf{x}^1 = (\mathbf{x}_{\text{pub}}^1, \mathbf{x}_{\text{priv}}^1), \dots, \mathbf{x}^{Q_{\text{msg}}} = (\mathbf{x}_{\text{pub}}^{Q_{\text{msg}}}, \mathbf{x}_{\text{priv}}^{Q_{\text{msg}}})$.
3. **Auxiliary Information.** It outputs the auxiliary information aux.

To prove the security as per Definition 5.20, we prove

$$\left(\text{mpk} = \text{prFE.mpk}, \text{aux}, \{C^k\}_{k \in [Q_{\text{key}}]}, \left\{ \text{sk}^k := (\mathbf{r}^k, \text{prFE.sk}_0^k, \text{prFE.sk}_1^k, \text{prFE.sk}_2^k) \right\}_{k \in [Q_{\text{key}}]}, \text{Lobf}_{\text{off}}, \left\{ \mathbf{x}_{\text{pub}}^j, \text{prFE.ct}^j \right\}_{j \in [Q_{\text{msg}}]} \right) \approx_c \left(\text{mpk} = \text{prFE.mpk}, \text{aux}, \{C^k\}_{k \in [Q_{\text{key}}]}, \left\{ \text{sk}^k := (\mathbf{r}^k, \text{prFE.sk}_0^k, \text{prFE.sk}_1^k, \text{prFE.sk}_2^k) \right\}_{k \in [Q_{\text{key}}]}, \text{Lobf}_{\text{off}}, \left\{ \mathbf{x}_{\text{pub}}^j, \delta^j \leftarrow C\mathcal{T}_{\text{prFE}} \right\}_{j \in [Q_{\text{msg}}]} \right) \quad (5.44)$$

assuming we have

$$\begin{aligned} & (1^\lambda, \text{aux}, \{\mathbf{x}_{\text{pub}}^j, C^k, C^k(\mathbf{x}^j)\}_{j \in [Q_{\text{msg}}], k \in [Q_{\text{key}}]}) \\ & \approx_c (1^\lambda, \text{aux}, \{\mathbf{x}_{\text{pub}}^j, C^k, \Delta^{j,k} \leftarrow \{0, 1\}\}_{j \in [Q_{\text{msg}}], k \in [Q_{\text{key}}]}) \end{aligned} \quad (5.45)$$

where

$$\begin{aligned} & \left(\{C^k\}_{k \in [Q_{\text{key}}]}, \{\mathbf{x}^j = (\mathbf{x}_{\text{pub}}^j, \mathbf{x}_{\text{priv}}^j)\}_{j \in [Q_{\text{msg}}]}, \text{aux} \in \{0, 1\}^* \right) \leftarrow \text{Samp}(1^\lambda), \\ & (\text{prFE.mpk}, \text{prFE.msk}) \leftarrow \text{prFE.Setup}(1^\lambda, \text{prm}), \\ & \mathbf{r}^k \leftarrow \{0, 1\}^\lambda, \text{prFE.sk}_0^k \leftarrow \text{prFE.KeyGen}(\text{prFE.msk}, F_0[\mathbf{r}^k]), \text{ for } F_0[\mathbf{r}^k] \text{ as defined in Figure 5.4,} \\ & \text{prFE.sk}_1^k \leftarrow \text{prFE.KeyGen}(\text{prFE.msk}, F_1[\mathbf{r}^k]), \text{ for } F_1[\mathbf{r}^k, \text{dig}_{C^k}] \text{ as defined in Figure 5.5} \\ & \text{dig}_{C^k} \leftarrow \text{LDigest}(\{(i, C_i^k)\}_{i \in |C^k|}) \\ & \text{prFE.sk}_2^k \leftarrow \text{prFE.KeyGen}(\text{prFE.msk}, F_2[\mathbf{r}^k, \text{dig}_{C^k}]), \text{ for } F_2[\mathbf{r}^k, \text{dig}_{C^k}] \text{ as defined in Figure 5.6,} \end{aligned}$$

$$(\text{Lobf}_{\text{off}}, \text{Lst}) \leftarrow \text{LObfOff}(1^\lambda, 1^S),$$

$$\text{dig}_{\mathbf{x}_{\text{pub}}^j} \leftarrow \text{LDigest}(1^\lambda, \{(i, \mathbf{x}_{\text{pub},i}^j)\}_{i \in [L_{\text{pub}}]}), \text{sd}_0^j, \text{sd}_1^j, \text{sd}_2^j \leftarrow \{0, 1\}^\lambda,$$

$$\text{prFE.ct}^j \leftarrow \text{prFE.Enc}(\text{prFE.mpk}, (\text{dig}_{\mathbf{x}_{\text{pub}}^j}, \mathbf{x}_{\text{priv}}^j, \text{Lst}, \text{sd}_0^j, \text{sd}_1^j, \text{sd}_2^j)) \quad \text{for } j \in [Q_{\text{msg}}].$$

We invoke the security of prFE with sampler $\text{Samp}_{\text{prFE}}$ that outputs

$$\left(\begin{array}{l} \text{Functions:} \quad \{F_0[\mathbf{r}^k], F_1[\mathbf{r}^k], F_2[\mathbf{r}^k, \text{dig}_{C^k}]\}_{k \in [Q_{\text{key}}]}, \\ \text{Inputs:} \quad \left\{ \mathbf{x}_{\text{prFE}}^j := \left(\text{dig}_{\mathbf{x}_{\text{pub}}^j}, \mathbf{x}_{\text{priv}}^j, \text{Lst}, \text{sd}_0^j, \text{sd}_1^j, \text{sd}_2^j \right) \right\}_{j \in [Q_{\text{msg}}]}, \\ \text{Auxiliary Information:} \quad \text{aux}_{\text{prFE}} := \left(\text{aux}, \left\{ \mathbf{x}_{\text{pub}}^j \right\}_{j \in [Q_{\text{msg}}]}, \text{Lobf}_{\text{off}}, \{C^k, \mathbf{r}^k\}_{k \in [Q_{\text{key}}]} \right) \end{array} \right)$$

By the security guarantee of prFE with sampler $\text{Samp}_{\text{prFE}}$, Equation (5.44) holds if

$$\left(\begin{array}{l} \left\{ \begin{array}{l} F_0[\mathbf{r}^k](\mathbf{x}_{\text{prFE}}^j) = \left\{ \text{lab}_{i, \mathbf{x}_{\text{priv}, i-L_{\text{pub}}}}^{j,k} \right\}_{i \in [L_{\text{pub}}+1, L_{\text{pub}}+L_{\text{priv}}]} \\ F_1[\mathbf{r}^k](\mathbf{x}_{\text{prFE}}^j) = \text{Lobf}_{\text{on},1}^{j,k} \\ F_2[\mathbf{r}^k, \text{dig}_{C^k}](\mathbf{x}_{\text{prFE}}^j) = \text{Lobf}_{\text{on},2}^{j,k} \end{array} \right\}_{j \in [Q_{\text{msg}}], k \in [Q_{\text{key}}]} \\ \text{aux}, \{ \mathbf{x}_{\text{pub}}^j \}_{j \in [Q_{\text{msg}}]}, \text{Lobf}_{\text{off}}, \{C^k, \mathbf{r}^k\}_{k \in [Q]} \end{array} \right) \\ \approx_c \left(\begin{array}{l} \left\{ \begin{array}{l} \left\{ \delta_{0,i}^{j,k} \right\}_{i \in [L_{\text{priv}}]} \\ \delta_1^{j,k} \\ \delta_2^{j,k} \end{array} \right\}_{j \in [Q_{\text{msg}}], k \in [Q_{\text{key}}]} \\ \text{aux}, \{ \mathbf{x}_{\text{pub}}^j \}_{j \in [Q_{\text{msg}}]}, \text{Lobf}_{\text{off}}, \{C^k, \mathbf{r}^k\}_{k \in [Q_{\text{key}}]} \end{array} \right), \end{array} \right) \quad (5.46)$$

where

$$R_b^{j,k} := \text{PRF}(\text{sd}_b^j, \mathbf{r}^k) \quad \text{for } b = 0, 1, 2,$$

$$(\text{lab}_{i,0}^{j,k}, \text{lab}_{i,1}^{j,k}) = \text{Garble}_{\text{inp},i}(1^\lambda; R_0^{j,k}) \quad \text{for } i \in [L_{\text{pub}} + 1, L_{\text{pub}} + L_{\text{priv}}],$$

$$(\text{Lobf}_{\text{off}}, \text{Lst}) \leftarrow \text{LObfOff}(1^\lambda, 1^S)$$

$$\text{Lobf}_{\text{on},1}^{j,k} = \text{LObfOn}(\text{Lst}, \text{dig}_{\mathbf{x}_{\text{pub}}^j}, E_{\text{pub}}[R_0^{j,k}]; R_1^{j,k})$$

$$\text{Lobf}_{\text{on},2}^{j,k} = \text{LObfOn}(\text{Lst}, \text{dig}_{C^k}, E_{\text{cir}}[R_0^{j,k}]; R_2^{j,k}) \quad \text{for } j \in [Q_{\text{msg}}], k \in [Q_{\text{key}}].$$

Therefore, it suffices to prove Equation (5.46). We observe that it suffices to prove a variant of Equation (5.46) where $R_0^{j,k}, R_1^{j,k}, R_2^{j,k}$ are replaced with independently chosen truly random strings, since (the original version of) Equation (5.46) then follows. This can be seen by the following indistinguishability:

$$\begin{aligned} & \left\{ R_b^{j,k} = \text{PRF}(\text{sd}_b^j, \mathbf{r}^k) : \text{sd}_b^j \leftarrow \{0, 1\}^\lambda, \mathbf{r}^k \leftarrow \{0, 1\}^\lambda \right\}_{b \in \{0,1,2\}, j \in [Q_{\text{msg}}], k \in [Q_{\text{key}}]} \\ & \approx_s \left\{ R_b^{j,k} = \text{PRF}(\text{sd}_b^j, \mathbf{r}^k) : \text{sd}_b^j \leftarrow \{0, 1\}^\lambda, \mathbf{r}^k \leftarrow \{0, 1\}^\lambda \setminus \{\mathbf{r}^1, \dots, \mathbf{r}^{k-1}\} \right\}_{b \in \{0,1,2\}, j \in [Q_{\text{msg}}], k \in [Q_{\text{key}}]} \\ & \approx_c \left\{ R_b^{j,k} \leftarrow \{0, 1\}^\lambda \right\}_{b \in \{0,1,2\}, j \in [Q_{\text{msg}}], k \in [Q_{\text{key}}]}, \end{aligned}$$

where in the second distribution, $\{\mathbf{r}^k\}_{k \in [Q_{\text{key}}]}$ is distributed uniformly at random over $(\{0, 1\}^\lambda)^{Q_{\text{key}}}$ with the constraint that there is no collision among them. We observe that the first indistinguishability holds since there is a colliding pair in $\{\mathbf{r}^k\}_{k \in [Q_{\text{key}}]}$ with probability at most $Q_{\text{key}}^2/2^\lambda$ in the first distribution and the second indistinguishability holds by the security of PRF, since each $R_b^{j,k}$ is generated by fresh pair of seed and input. We then invoke the security of LprIO with sampler $\text{Samp}_{\text{LprIO}}$ that outputs

$$\left(\begin{array}{l} \text{Inputs:} \quad \left\{ \begin{array}{l} X^{j,k} := \{X_i^{j,k} := (i, C_i^k)\}_{i \in [C^k]}, \\ \bar{X}^{j,k} := \{\bar{X}_i^{j,k} := (i, \mathbf{x}_{\text{pub},i}^j)\}_{i \in [L_{\text{pub}}]} \end{array} \right\}_{j \in [Q_{\text{msg}}], k \in [Q_{\text{key}}]} \\ \text{Functions:} \quad \left\{ E^{j,k} := E_{\text{cir}}[R_0^{j,k}], \bar{E}^{j,k} := E_{\text{pub}}[R_0^{j,k}] \right\}_{j \in [Q_{\text{msg}}], k \in [Q_{\text{key}}]}, \\ \text{Auxiliary} \\ \text{Information:} \quad \text{aux}_{\text{LprIO},1} := \left\{ \text{lab}_{i, \mathbf{x}_{\text{priv},i-L_{\text{pub}}}}^{j,k} \right\}_{i \in [L_{\text{pub}}+1, L_{\text{pub}}+L_{\text{priv}}], j \in [Q_{\text{msg}}], k \in [Q_{\text{key}}]}, \\ \text{aux}_{\text{LprIO},2} := (\text{aux}, \{\mathbf{x}_{\text{pub}}^j\}_{j \in [Q_{\text{msg}}]}, \{C^k, \mathbf{r}^k\}_{k \in [Q_{\text{key}}]}) \end{array} \right),$$

where we consider $2Q_{\text{msg}}Q_{\text{key}}$ inputs and functions. For simplifying the notations, we use two indices $j \in [Q_{\text{msg}}]$ and $k \in [Q_{\text{key}}]$ and consider barred and unbarred symbols to represent different variables instead of using the single index for inputs and functions.

In the rest of the proof, we prove

$$\left(\left\{ E^{j,k}(X_i^{j,k}) = \tilde{C}_i^{j,k} \right\}_{j,k,i \in [C^k]}, \left\{ \bar{E}^{j,k}(\bar{X}_i^{j,k}) = \text{lab}_{i, \mathbf{x}_{\text{pub},i}^j}^{j,k} \right\}_{j,k,i \in [L_{\text{pub}}]}, \text{aux}_{\text{LprIO},1}, \text{aux}_{\text{LprIO},2} \right)$$

$$\approx_c \left(\left\{ \gamma_i^{j,k} \right\}_{j,k,i \in [|C^k|]}, \left\{ \tilde{\gamma}_i^{j,k} \right\}_{j,k,i \in [L_{\text{pub}}]}, \alpha, \text{aux}_{\text{LprIO},2} \right) \quad (5.47)$$

where $\gamma_i^{j,k} \leftarrow \mathcal{CT}_{\text{bGC}}^i$ for $j \in [Q_{\text{msg}}], k \in [Q_{\text{key}}], i \in [|C^k|]$, $\tilde{\gamma}_i^{j,k} \leftarrow \{0,1\}^\lambda$ for $j \in [Q_{\text{msg}}], k \in [Q_{\text{key}}], i \in [L_{\text{pub}}]$, and $\alpha \leftarrow \{0,1\}^{\lambda L_{\text{priv}} Q_{\text{msg}} Q_{\text{key}}}$. Note that the length of α is the same as that of $\text{aux}_{\text{LprIO}}$. Here $\mathcal{CT}_{\text{bGC}}^i$ is the co-domain of Garble_i algorithm. This suffices to conclude the proof, since Equation (5.47) implies Equation (5.46) by the security of LprIO and Theorem 5.23.

To prove Equation (5.47), we introduce the following sequence of hybrids.

Hyb₁. This is the LHS distribution of Equation (5.47). By unrolling the definition of the functions and rearranging the terms, we can see that this is equivalent to the following distribution:

$$\left(\left\{ \left\{ \text{lab}_{i,\mathbf{x}_i^j}^{j,k} \right\}_{i \in [L]}, \tilde{C}^{j,k} \right\}_{j \in [Q_{\text{msg}}], k \in [Q_{\text{key}}]}, \text{aux}, \{\mathbf{x}_{\text{pub}}^j\}_{j \in [Q_{\text{msg}}]}, \{C^k, \mathbf{r}^k\}_{k \in [Q_{\text{key}}]} \right)$$

where $(\{\text{lab}_{i,b}^{j,k}\}_{i \in [L], b \in \{0,1\}}, \tilde{C}^{j,k}) \leftarrow \text{Garble}(1^\lambda, C^k; R_0^{j,k})$, $\mathbf{r}^k \leftarrow \{0,1\}^\lambda$ for $j \in [Q_{\text{msg}}]$ and $k \in [Q_{\text{key}}]$, $(\text{aux}, \{\mathbf{x}_{\text{pub}}^j\}_j, \{C^k\}_k) \leftarrow \text{Samp}(1^\lambda)$, and \mathbf{x}_i^j is the i -th bit of $\mathbf{x}^j = (\mathbf{x}_{\text{pub}}^j, \mathbf{x}_{\text{priv}}^j)$. Note that here, we merge the process of generating $\{\tilde{C}_i^{j,k}\}_i$, $\{\text{lab}_{i,\mathbf{x}_{\text{pub},i}}^{j,k}\}_i$, and $\{\text{lab}_{i,\mathbf{x}_{\text{priv},i}}^{j,k}\}_i$ into a single process of running $\text{Garble}(1^\lambda, C^k; R_0^{j,k})$. This does not change the distribution due to the decomposability of bGC , since the former processes are run on input the common randomness $R_0^{j,k}$ for each j and k .

Hyb₂. This hybrid is same as the previous one except that we compute the labels and the garbled circuits using the simulation algorithm. Namely, the view of the adversary in this hybrid is

$$\left(\left\{ \left\{ \text{lab}_i^{j,k} \right\}_{i \in [L]}, \tilde{C}^{j,k} \right\}_{j \in [Q_{\text{msg}}], k \in [Q_{\text{key}}]}, \text{aux}, \{\mathbf{x}_{\text{pub}}^j\}_{j \in [Q_{\text{msg}}]}, \{C^k, \mathbf{r}^k\}_{k \in [Q_{\text{key}}]} \right)$$

where we compute $(\tilde{C}^{j,k}, \{\text{lab}_i^{j,k}\}_{i \in [L]}) \leftarrow \text{bGC.Sim}(1^\lambda, 1^L, C^k(\mathbf{x}^j))$ for all $j \in$

$[Q_{\text{msg}}], k \in [Q_{\text{key}}]$. Due to the simulation security of **bGC**, this hybrid is computationally indistinguishable from the previous one.

Hyb₃. This hybrid is same as the previous one except that we input random strings into the simulator of the blind garbled circuit. Namely, we compute $(\tilde{C}^{j,k}, \{\text{lab}_i^{j,k}\}_{i \in [L]})$ as $(\tilde{C}^{j,k}, \{\text{lab}_i^{j,k}\}_{i \in [L]}) \leftarrow \text{bGC.Sim}(1^\lambda, 1^L, \Delta^{j,k})$, where $\Delta^{j,k} \leftarrow \{0, 1\}$ for all j and k . We can see that this hybrid is indistinguishable from the previous one by Equation (5.45).

Hyb₄. This hybrid is same as the previous one except that we replace the output of the simulator for **bGC** with random strings. By the blindness of **bGC**, this game is indistinguishable from the previous hybrid.

By rearranging the terms, we can see that the distribution in **Hyb₄** is equivalent to that of the RHS of Equation (5.47). We therefore have that the LHS and RHS of Equation (5.47). This completes the proof of Theorem 5.22. ■

5.5.4 Reducing the Dependency on Private Input Length

The sizes of the master public key, ciphertexts, and the secret keys of our construction in Section 5.5.2 are all independent from the length of \mathbf{x}_{pub} and C . However, they still depend on the length of \mathbf{x}_{priv} . Here, we show a simple conversion that removes this dependency from the sizes of the master public key and secret key. The size of the ciphertext inherently depends on the length of the private input since it should be hidden, but we can make this dependency minimal if we start from the scheme with the ciphertext size being independent of the length of the public input. In particular, the size of the ciphertext only additively depends on the length of \mathbf{x}_{priv} . By applying the conversion in this section to our construction in Section 5.5.2, we obtain a construction of partial-hiding FE with the optimal parameter size.

Building Blocks. We use the following ingredients for our construction.

1. A secret key encryption scheme $\text{SKE} = (\text{SKE.Setup}, \text{SKE.Enc}, \text{SKE.Dec})$ with the message space $\{0, 1\}^{L_{\text{priv}}}$ with pseudorandom ciphertext space as per Definition 5.1. We denote the ciphertext space of the scheme by $\mathcal{CT}_{\text{SKE}}$ and the key space of the scheme by \mathcal{K}_{SKE} . Without loss of generality, we assume $\mathcal{K}_{\text{SKE}} = \{0, 1\}^\lambda$. Furthermore, we assume that the ciphertext space of SKE is $\mathcal{CT}_{\text{SKE}} = \{0, 1\}^{L_{\text{priv}} + \lambda}$. Such a construction can be obtained by using PRF for example.
2. A PHprFE scheme $\text{PHprFE} = (\text{Setup}, \text{KeyGen}, \text{Enc} = (\text{EncOff}, \text{EncOn}), \text{Dec})$ whose sizes of the master public key, ciphertext, and secret key are all $\text{poly}(\lambda, L_{\text{priv}})$. We can construct such a scheme assuming LWE and evasive LWE as is shown in Section 5.5. We denote the online part of the ciphertext space by $\mathcal{CT}_{\text{PHprFE}} = \{0, 1\}^{\ell_{\text{PHprFE}}^{\text{cton}}}$.

We describe the new scheme $\text{PHprFE}' = (\text{Setup}', \text{KeyGen}', \text{Enc}' = (\text{EncOff}', \text{EncOn}'), \text{Dec}')$ for circuit class $C : \{0, 1\}^{L_{\text{pub}}} \times \{0, 1\}^{L_{\text{priv}}} \rightarrow \{0, 1\}$ in the following.

$\text{Setup}'(1^\lambda, 1^{L_{\text{pub}}}, 1^{L_{\text{priv}}})$. It does the following.

- Run $\text{Setup}(1^\lambda, 1^{L_{\text{pub}} + L_{\text{priv}} + \lambda}, 1^\lambda) \rightarrow (\text{mpk}, \text{msk})$.
- Output the master public key mpk and the master secret key msk .

$\text{KeyGen}'(\text{msk}, C)$. It does the following.

- Define the circuit C' as follows.
On input $(\text{SKE.sk}, \mathbf{x}_{\text{pub}}, \text{SKE.ct})$, output

$$C'(\mathbf{x}_{\text{pub}}, \text{SKE.ct}, \text{SKE.sk}) = C(\mathbf{x}_{\text{pub}}, \text{SKE.Dec}(\text{SKE.sk}, \text{SKE.ct})).$$

- It runs $\text{KeyGen}(\text{msk}, C') \rightarrow \text{sk}_{C'}$ and outputs $\text{sk}_{C'}$.

$\text{Enc}'(\text{mpk}, \mathbf{x} = (\mathbf{x}_{\text{pub}}, \mathbf{x}_{\text{priv}}))$. The encryption algorithm does the following. We divide the algorithm into the following two steps.

$\text{EncOff}'(\text{mpk})$. It runs $\text{EncOff}(\text{mpk}) \rightarrow (\text{ct}_{\text{off}}, \text{st})$ and outputs the offline part of the ciphertext ct_{off} and state st .

$\text{EncOn}'(\text{st}, \mathbf{x} = (\mathbf{x}_{\text{pub}}, \mathbf{x}_{\text{priv}}))$. It does the following.

- Run $\text{SKE.Setup}(1^\lambda) \rightarrow \text{SKE.sk}$.
- Run $\text{SKE.Enc}(\text{SKE.sk}, \mathbf{x}_{\text{priv}}) \rightarrow \text{SKE.ct}$.
- Set $\mathbf{y} := (\mathbf{x}_{\text{pub}}, \text{SKE.ct})$ and $\mathbf{z} := \text{SKE.sk}$ and run $\text{EncOn}(\text{mpk}, (\mathbf{y}, \mathbf{z})) \rightarrow \text{ct}_{\text{on}}$.
- Output $\text{ct}_{\text{on}}' := (\text{SKE.ct}, \text{ct}_{\text{on}})$.

The final output of $\text{Enc}'(\text{mpk}, \mathbf{x} = (\mathbf{x}_{\text{pub}}, \mathbf{x}_{\text{priv}}))$ is $\text{ct}' := (\text{ct}_{\text{off}}, \text{SKE.ct}, \text{ct}_{\text{on}})$.

$\text{Dec}'(\text{mpk}, \mathbf{x}_{\text{pub}}, \text{sk}_{C'}, C, \text{ct}')$. It does the following.

- Parse $\text{ct}' \rightarrow (\text{ct}_{\text{off}}, \text{SKE.ct}, \text{ct}_{\text{on}})$ and set $\mathbf{y} := (\mathbf{x}_{\text{pub}}, \text{SKE.ct})$ and $\text{ct} := (\text{ct}_{\text{off}}, \text{ct}_{\text{on}})$.
- Define C' as in the key generation algorithm.
- Run $\text{Dec}(\text{mpk}, \mathbf{y}, \text{sk}_{C'}, C', \text{ct}) \rightarrow w$ and output w .

Correctness. We make the following observations.

- From the perfect correctness of underlying PHprFE scheme, with public input $\mathbf{y} = (\mathbf{x}_{\text{pub}}, \text{SKE.ct})$ and private input SKE.sk , we have

$$\begin{aligned} \text{Dec}(\text{mpk}, \mathbf{y}, \text{sk}_{C'}, C', \text{ct}) &= C'(\mathbf{y}, \text{SKE.sk}) \\ &= C'(\mathbf{x}_{\text{pub}}, \text{SKE.ct}, \text{SKE.sk}) \\ &= C(\mathbf{x}_{\text{pub}}, \text{SKE.Dec}(\text{SKE.sk}, \text{SKE.ct})) \quad (\text{by definition of } C') \end{aligned}$$

- Next, from the perfect correctness of the SKE scheme, we have

$$\text{SKE.Dec}(\text{SKE.sk}, \text{SKE.ct}) = \text{SKE.Dec}(\text{SKE.sk}, \text{SKE.Enc}(\text{SKE.sk}, \mathbf{x}_{\text{priv}})) = \mathbf{x}_{\text{priv}}$$

Thus $\text{Dec}(\text{mpk}, \mathbf{y}, \text{sk}_{C'}, C', \text{ct}) = C(\mathbf{x}_{\text{pub}}, \text{SKE.Dec}(\text{SKE.sk}, \text{SKE.ct})) = C(\mathbf{x}_{\text{pub}}, \mathbf{x}_{\text{priv}})$ and hence the correctness.

Efficiency. Here, we show $|\text{mpk}| = \text{poly}(\lambda)$, $|\text{sk}_{C'}| = \text{poly}(\lambda)$, $|\text{ct}'| = \text{poly}(\lambda) + L_{\text{priv}}$.

Given that the length of the private input for the underlying PHprFE is λ , it is straightforward to see the first two equations above hold by the efficiency of PHprFE. To bound the length of the ciphertext ct' , we first observe that $\text{ct}' = (\text{ct}_{\text{off}}, \text{SKE.ct}, \text{ct}_{\text{on}})$, where $\text{ct} := (\text{ct}_{\text{off}}, \text{ct}_{\text{on}})$ constitutes a ciphertext of the underlying PHprFE encrypting (\mathbf{y}, \mathbf{z}) . We have $|\text{ct}| = \text{poly}(\lambda, |\mathbf{z}|) = \text{poly}(\lambda)$ and $|\text{SKE.ct}| = \text{poly}(\lambda) + L_{\text{priv}}$. Therefore, the last equation above follows as well.

We therefore have the following theorem. The security of the scheme is proven in Theorem 5.27.

Theorem 5.25. *Assuming LWE and evasive LWE assumptions, there exists a partially hiding pseudorandom FE scheme, for circuit class $\mathcal{C} = \{C : \{0, 1\}^{L_{\text{pub}}} \times \{0, 1\}^{L_{\text{priv}}} \rightarrow \{0, 1\}\}$, that satisfies reusable security (as per Definition 5.20) whose sizes of the master public key and the secret key are fixed polynomial $\text{poly}(\lambda)$. Furthermore, the size of the ciphertext is $L_{\text{priv}} + \text{poly}(\lambda)$.*

Optimal prFE. Using Remark 24, we get a prFE scheme as a special case of PHprFE, for $\mathbf{x}_{\text{pub}} = \perp$ and $\mathbf{x}_{\text{priv}} = \mathbf{x}$. Next, observing the fact that (1) with $\mathbf{x}_{\text{pub}} = \perp$ the security of PHprFE (Definition 5.19) is equivalent to single-challenge security of prFE (Definition 5.13, with $Q_{\text{msg}} = 1$) and (2) Definition 5.19 is implied by Definition 5.20, we get a prFE scheme with optimal parameters. We formalise this in the following theorem.

Theorem 5.26. *Assuming LWE and evasive LWE assumptions, there exists a prFE scheme, for circuit class $\mathcal{C} = \{C : \{0, 1\}^{L_{\text{inp}}} \rightarrow \{0, 1\}\}$, that satisfies security (as per Definition 5.13, with $Q_{\text{msg}} = 1$) and efficiency $|\text{mpk}| = \text{poly}(\lambda)$, $|\text{sk}_C| = \text{poly}(\lambda)$, $|\text{ct}| = L_{\text{inp}} + \text{poly}(\lambda)$.*

Security The following theorem asserts the security of the construction.

Theorem 5.27. *The above construction PHprFE' satisfies reusable security as per Definition 5.20 if so does PHprFE and SKE is secure as per Definition 5.1.*

Proof. Consider a sampler $\text{Samp}_{\text{PHprFE}'}$ that generates the following:

1. **Key Queries.** It issues Q_{key} key queries $C_1, \dots, C_{Q_{\text{key}}}$.
2. **Ciphertext Queries.** It issues messages $\mathbf{x}^1 = (\mathbf{x}_{\text{pub}}^1, \mathbf{x}_{\text{priv}}^1), \dots, \mathbf{x}^{Q_{\text{msg}}} = (\mathbf{x}_{\text{pub}}^{Q_{\text{msg}}}, \mathbf{x}_{\text{priv}}^{Q_{\text{msg}}})$.
3. **Auxiliary Information.** It outputs the auxiliary information $\text{aux}_{\mathcal{A}}$.

To prove the security as per Definition 5.20, we prove

$$\left(\text{mpk}, \text{aux}, \{C_k\}_{k \in [Q_{\text{msg}}]}, \left\{ \text{sk}_k := \text{sk}_{C'_k} \right\}_{k \in [Q_{\text{msg}}]}, \text{Lobf}_{\text{off}}, \left\{ \mathbf{x}_{\text{pub}}^j, \text{SKE.ct}^j, \text{ct}_{\text{on}}^j \right\}_{j \in [Q_{\text{msg}}]} \right) \approx_c \left(\text{mpk}, \text{aux}, \{C_k\}_{k \in [Q_{\text{msg}}]}, \left\{ \text{sk}_k := \text{sk}_{C'_k} \right\}_{k \in [Q_{\text{msg}}]}, \text{Lobf}_{\text{off}}, \left\{ \mathbf{x}_{\text{pub}}^j, \gamma^j, \delta^j \right\}_{j \in [Q_{\text{msg}}]} \right) \quad (5.48)$$

where $\gamma^j \leftarrow \mathcal{CT}_{\text{SKE}}$ and $\delta^j \leftarrow \mathcal{CT}_{\text{PHprFE}}$ for $j \in [Q_{\text{msg}}]$, assuming we have

$$(1^\lambda, \text{aux}, \{\mathbf{x}_{\text{pub}}^j, C_k, C_k(\mathbf{x}^j)\}_{j \in [Q_{\text{msg}}], k \in [Q_{\text{key}}]}) \approx_c (1^\lambda, \text{aux}, \{\mathbf{x}_{\text{pub}}^j, C_k, \Delta_k^j \leftarrow \{0, 1\}\}_{j \in [Q_{\text{msg}}], k \in [Q_{\text{key}}]}) \quad (5.49)$$

where

$$\begin{aligned} & \left(\{C_k\}_{k \in [Q_{\text{key}}]}, \{\mathbf{x}^j = (\mathbf{x}_{\text{pub}}^j, \mathbf{x}_{\text{priv}}^j)\}_{j \in [Q_{\text{msg}}]}, \text{aux} \in \{0, 1\}^* \right) \leftarrow \text{Samp}_{\text{PHprFE}'}(1^\lambda), \\ & \text{SKE.sk} \leftarrow \text{SKE.Setup}(1^\lambda), \quad \text{SKE.ct}^j \leftarrow \text{SKE.Enc}(\text{SKE.sk}, \mathbf{x}_{\text{priv}}^j), \\ & (\text{mpk}, \text{msk}) \leftarrow \text{Setup}(1^\lambda, 1^{L_{\text{pub}} + L_{\text{priv}} + \lambda}, 1^\lambda), \\ & \text{sk}_k \leftarrow \text{KeyGen}(\text{msk}, C'_k) \text{ for } k \in [Q_{\text{key}}], \text{ where } C'_k \text{ is defined from } C_k \text{ as in the construction,} \\ & (\text{ct}_{\text{off}}, \text{st}) \leftarrow \text{EncOff}(\text{mpk}), \\ & \text{ct}_{\text{on}}^j \leftarrow \text{EncOn}(\text{st}, \mathbf{x}_{\text{priv}}^j) \text{ for } j \in [Q_{\text{msg}}]. \end{aligned}$$

We invoke the security of PHprFE with sampler $\text{Samp}_{\text{PHprFE}}$ that outputs

$$\left(\begin{array}{ll} \text{Functions:} & \{C'_k\}_{k \in [Q_{\text{key}}]}, \\ \text{Inputs:} & \{X_j := (X_{\text{pub}}^j := (\mathbf{x}_{\text{pub}}^j, \text{SKE.ct}^j), X_{\text{priv}}^j := \text{SKE.sk}^j)\}_{j \in [Q_{\text{msg}}]}, \\ \text{Auxiliary Information:} & \text{aux}_{\text{prFE}} := (\text{aux}, \{C_k\}_{k \in [Q_{\text{key}}]}) \end{array} \right)$$

By the security guarantee of PHprFE with sampler $\text{Samp}_{\text{PHprFE}}$,

$$\left(\begin{array}{l} \text{mpk, aux, } \{C_k\}_{k \in [Q_{\text{key}}]}, \\ \{\text{sk}_k := \text{sk}_{C'_k}\}_{k \in [Q_{\text{key}}]}, \\ \text{Lobf}_{\text{off}}, \{X_{\text{pub}}^j = (\mathbf{x}_{\text{pub}}^j, \text{SKE.ct}^j), \text{ct}_{\text{on}}^j\}_{j \in [Q_{\text{msg}}]} \end{array} \right) \approx_c \left(\begin{array}{l} \text{mpk, aux, } \{C_k\}_{k \in [Q_{\text{key}}]}, \\ \{\text{sk}_k := \text{sk}_{C'_k}\}_{k \in [Q_{\text{key}}]}, \\ \text{Lobf}_{\text{off}}, \{X_{\text{pub}}^j = (\mathbf{x}_{\text{pub}}^j, \text{SKE.ct}^j), \delta^j\}_{j \in [Q_{\text{msg}}]} \end{array} \right) \quad (5.50)$$

holds if

$$\left(\text{aux}, \{C_k, X_{\text{pub}}^j, C'_k(X^j)\}_{j \in [Q_{\text{msg}}], k \in [Q_{\text{key}}]} \right) \approx_c \left(\text{aux}, \{C_k, X_{\text{pub}}^j, \Delta_k^j\}_{j \in [Q_{\text{msg}}], k \in [Q_{\text{key}}]} \right). \quad (5.51)$$

We first observe that Equation (5.50) implies Equation (5.48), since we can invoke the security of SKE to conclude that

$$\left(\begin{array}{l} \text{mpk, aux, } \{C_k\}_{k \in [Q_{\text{key}}]}, \\ \{\text{sk}_k := \text{sk}_{C'_k}\}_{k \in [Q_{\text{key}}]}, \\ \text{Lobf}_{\text{off}}, \{\mathbf{x}_{\text{pub}}^j, \text{SKE.ct}^j, \delta^j\}_{j \in [Q_{\text{msg}}]} \end{array} \right) \approx_c \left(\begin{array}{l} \text{mpk, aux, } \{C_k\}_{k \in [Q_{\text{key}}]}, \\ \{\text{sk}_k := \text{sk}_{C'_k}\}_{k \in [Q_{\text{key}}]}, \\ \text{Lobf}_{\text{off}}, \{\mathbf{x}_{\text{pub}}^j, \gamma^j, \delta^j\}_{j \in [Q_{\text{msg}}]} \end{array} \right)$$

holds by noting that SKE.sk^j is used only for computing SKE.ct^j and not used anywhere else.

Therefore, it suffices to prove Equation (5.51) to conclude the proof. We have

$$\begin{aligned} & \left(\text{aux}, \{C_k, X_{\text{pub}}^j = (\mathbf{x}_{\text{pub}}^j, \text{SKE.ct}^j), C'_k(X^j)\}_{j,k} \right) = \left(\text{aux}, \{C_k, \mathbf{x}_{\text{pub}}^j, \text{SKE.ct}^j, C_k(\mathbf{x}^j)\}_{j,k} \right) \\ & \approx_c \left(\text{aux}, \{C_k, \mathbf{x}_{\text{pub}}^j, \gamma^j \leftarrow \mathcal{CT}_{\text{SKE}}, C_k(\mathbf{x}^j)\}_{j,k} \right) \end{aligned}$$

$$\begin{aligned}
&\approx_c \left(\text{aux}, \left\{ C_k, \mathbf{x}_{\text{pub}}^j, \gamma^j \leftarrow C\mathcal{T}_{\text{SKE}}, \Delta_k^j \leftarrow \{0, 1\} \right\}_{j,k} \right) \\
&\approx_c \left(\text{aux}, \left\{ C_k, \mathbf{x}_{\text{pub}}^j, \text{SKE.ct}^j, \Delta_k^j \leftarrow \{0, 1\} \right\}_{j,k} \right)
\end{aligned}$$

where the first line follows from the definition of C'_k and X^j , the second from the security of SKE noting that SKE.sk^j is used only for computing SKE.ct^j and not used anywhere else, the third from Equation (5.49), noting that adding random string $\{\gamma^j\}_j$ to the distributions in Equation (5.49) does not make the task of distinguishing the distributions any easier, and the fourth from the security of SKE again. This proves Equation (5.51) and therefore completes the proof. \blacksquare

5.5.5 Handling Longer Output

So far we have only considered the case where C is a circuit that outputs a single-bit string. Here, we discuss more general case where the output of the circuit C is longer. To handle such circuits in the construction, we only change the key generation algorithm. To generate a secret key for $C : \{0, 1\}^L \rightarrow \{0, 1\}^{\text{out}}$, we first consider circuits $\{C_j\}_{j \in [\text{out}]}$, where C_j is the circuit that outputs the j -th bit of C 's output and then generate secret keys for $\{C_j\}_j$. It is then easy to see that the all bits of $C(\mathbf{x})$ can be recovered by using the secret keys for $\{C_j\}_j$ by decrypting a ciphertext encrypting \mathbf{x} . This does not change the size of the master public key and ciphertext, but makes the secret key linearly dependent on the output length of C . Thus, we get the following theorems.

Theorem 5.28. *Assuming LWE and evasive LWE assumptions, there exists a PHprFE scheme, for circuit class $C = \{C : \{0, 1\}^{L_{\text{pub}}} \times \{0, 1\}^{L_{\text{priv}}} \rightarrow \{0, 1\}^{\text{out}}\}$, that satisfies reusable security (as per Definition 5.20) and efficiency*

$$|\text{mpk}| = \text{poly}(\lambda), \quad |\text{sk}_C| = \text{out} \cdot \text{poly}(\lambda), \quad |\text{ct}| = L_{\text{priv}} + \text{poly}(\lambda).$$

Theorem 5.29. *Assuming LWE and evasive LWE assumptions, there exists a prFE scheme, for circuit class $C = \{C : \{0, 1\}^{L_{\text{inp}}} \rightarrow \{0, 1\}\}$, that satisfies security (as per*

Definition 5.13, with $Q_{\text{msg}} = 1$) and efficiency

$$|\text{mpk}| = \text{poly}(\lambda), \quad |\text{sk}_C| = \text{out} \cdot \text{poly}(\lambda), \quad |\text{ct}| = L_{\text{inp}} + \text{poly}(\lambda).$$

5.6 KP-ABE/PE FOR UNBOUNDED DEPTH CIRCUITS WITH OPTIMAL PARAMETERS

Here we construct KP-ABE and KP-PE schemes supporting unbounded depth circuits and achieving optimal parameters.

5.6.1 Construction of KP-ABE with Optimal Parameters

In this section we construct an ABE scheme $\text{kpABE} = (\text{Setup}, \text{KeyGen}, \text{Enc}, \text{Dec})$ for message space $\{0, 1\}$ and circuit family C_ℓ consisting of circuits with input space $\{0, 1\}^L$ and output space $\{0, 1\}$. For the purpose of application in Section 5.7, we consider a decomposed encryption algorithm as $\text{Enc} = (\text{EncOff}, \text{EncOn})$ in our construction where $\text{EncOff}(\text{mpk}) \rightarrow (\text{ct}_{\text{off}}, \text{st})$, $\text{EncOn}(\text{st}, \mathbf{x}, \mu) \rightarrow \text{ct}_{\text{on}}$ and Enc outputs $(\text{ct}_{\text{off}}, \text{ct}_{\text{on}})$.

Building Blocks. We require the following building blocks for our construction.

1. A PHprFE scheme $\text{PHprFE} = \text{PHprFE} \cdot (\text{Setup}, \text{KeyGen}, \text{Enc} = (\text{EncOff}, \text{EncOn}), \text{Dec})$ from Section 5.5.4 for circuit family $\{C_{\text{prm}} = \{C : \mathcal{X}_{\text{pub}} \times \mathcal{X}_{\text{priv}} \rightarrow \mathcal{Y}\}\}_{\text{prm}}$ where $\mathcal{X}_{\text{pub}} = \{0, 1\}^\ell$, $\mathcal{X}_{\text{priv}} = \{0, 1\}^{\lambda+1}$ and $\mathcal{Y} = \{0, 1\}$. We denote the online part of the ciphertext space by $\text{PHprFE} \cdot \mathcal{CT}_{\text{on}}$. We require the PHprFE scheme to satisfy reusable security (Definition 5.20). As we show in Theorem 5.25, such a PHprFE scheme can be constructed assuming LWE and evasive LWE assumptions.
2. A PRF scheme $\text{PRF} : \{0, 1\}^\lambda \times \{0, 1\}^\lambda \rightarrow \{0, 1\}$. It is known that PRF can be constructed from one-way functions.

Now, we describe our construction.

$\text{Setup}(1^\lambda, 1^L)$. The setup algorithm does the following.

- Generate $(\text{PHprFE.msk}, \text{PHprFE.mpk}) \leftarrow \text{PHprFE.Setup}(1^\lambda, (1^\ell, 1^{\lambda+1}))$.

- Output $\text{msk} = \text{PHprFE.msk}$ and $\text{mpk} = \text{PHprFE.mpk}$.

$\text{KeyGen}(\text{msk}, C) \rightarrow \text{sk}_C$. The key generation algorithm does the following.

- Parse $\text{msk} = \text{PHprFE.msk}$.
- Sample $\mathbf{r} \leftarrow \{0, 1\}^\lambda$ and define the circuit $C[\mathbf{r}]$, with \mathbf{r} hardwired, as follows.
On input $(\mathbf{x}, \mu, \text{sd})$,

$$C[\mathbf{r}](\mathbf{x}, \mu, \text{sd}) = \begin{cases} \mu & \text{if } C(\mathbf{x}) = 0 \\ \text{PRF}(\text{sd}, \mathbf{r}) & \text{otherwise.} \end{cases}$$

- Compute $\text{PHprFE.sk} \leftarrow \text{PHprFE.KeyGen}(\text{PHprFE.msk}, C[\mathbf{r}])$.
- Output $\text{sk}_C = (\text{PHprFE.sk}, \mathbf{r})$.

$\text{Enc}(\text{mpk}, \mathbf{x}, \mu)$. The encryption algorithm works as follows.

$\text{EncOff}(\text{mpk})$. The offline phase of encryption does the following.

- Parse $\text{mpk} = \text{PHprFE.mpk}$.
- Compute
 $(\text{PHprFE.ct}_{\text{off}}, \text{PHprFE.st}) \leftarrow \text{PHprFE.EncOff}(\text{PHprFE.mpk})$.
- Output $\text{ct}_{\text{off}} = \text{PHprFE.ct}_{\text{off}}$ and $\text{st} = \text{PHprFE.st}$.

$\text{EncOn}(\text{st}, \mathbf{x}, \mu)$. The online phase of encryption does the following.

- Parse $\text{st} = \text{PHprFE.st}$.
- Sample $\text{sd} \leftarrow \{0, 1\}^\lambda$, set $X_{\text{pub}} = \mathbf{x}$ and $X_{\text{priv}} = (\mu, \text{sd})$.
- Compute $\text{PHprFE.ct}_{\text{on}} \leftarrow \text{PHprFE.Enc}(\text{PHprFE.st}, X_{\text{pub}}, X_{\text{priv}})$.
- Output $\text{ct}_{\text{on}} := \text{PHprFE.ct}_{\text{on}}$.

Output $\text{ct} := (\text{ct}_{\text{off}}, \text{ct}_{\text{on}})$.

$\text{Dec}(\text{mpk}, \text{sk}_C, C, \text{ct}, \mathbf{x})$. The decryption algorithm does the following.

- Parse $\text{mpk} = \text{PHprFE.mpk}$, $\text{sk}_C = (\text{PHprFE.sk}, \mathbf{r})$ and $\text{ct} = (\text{ct}_{\text{off}}, \text{ct}_{\text{on}}) = (\text{PHprFE.ct}_{\text{off}}, \text{PHprFE.ct}_{\text{on}})$.
- Compute
 $y = \text{PHprFE.Dec}(\text{PHprFE.mpk}, \mathbf{x}, \text{PHprFE.sk}, C[\mathbf{r}], \text{PHprFE.ct})$,
where $C[\mathbf{r}]$ is as defined in the key generation algorithm.
- Output y .

Correctness. For $\text{sk}_C = (\text{PHprFE.sk}, \mathbf{r})$ and $\text{ct} = (\text{ct}_{\text{off}}, \text{ct}_{\text{on}}) = (\text{PHprFE.ct}_{\text{off}}, \text{PHprFE.ct}_{\text{on}}) = \text{PHprFE.ct}$, we have

$$\begin{aligned} \text{PHprFE.Dec}(\text{PHprFE.mpk}, \mathbf{x}, \text{PHprFE.sk}, C[\mathbf{r}], \text{PHprFE.ct}) &= C[\mathbf{r}](\mathbf{x}, \mu, \text{sd}) \\ &= \begin{cases} \mu & \text{if } C(\mathbf{x}) = 1 \\ \text{PRF}(\text{sd}, \mathbf{r}) & \text{otherwise.} \end{cases} \end{aligned}$$

from the correctness of PHprFE scheme. Now, if $C(\mathbf{x}) = 1$, then from the definition of $C[\mathbf{r}]$, we get $C[\mathbf{r}](\mathbf{x}, \mu, \text{sd}) = \mu$ and hence the decryption outputs $y = \mu$ correctly.

Efficiency. Instantiating the PHprFE scheme from Section 5.5.4 with $|\text{PHprFE.mpk}| = \text{poly}(\lambda)$, $|\text{PHprFE.sk}_C| = \text{poly}(\lambda)$, $|\text{PHprFE.ct}| = \text{poly}(\lambda) + |\mathbf{x}_{\text{priv}}|$, our kpABE scheme satisfies $|\text{mpk}| = \text{poly}(\lambda)$, $|\text{sk}_C| = \text{poly}(\lambda)$, $|\text{ct}| = \text{poly}(\lambda)$. We formalise this instantiation using the following theorem. The security is proved in Section 5.6.2.

Theorem 5.30. *Under the LWE and Evasive LWE assumption, there exists very selectively secure KP-ABE scheme supporting circuits $\{C : \{0, 1\}^\ell \rightarrow \{0, 1\}\}$ with unbounded depth and single bit message space with*

$$|\text{mpk}| = \text{poly}(\lambda), |\text{sk}_C| = \text{poly}(\lambda), |\text{ct}| = \text{poly}(\lambda). \quad (5.52)$$

5.6.2 Security of KP-ABE

For our application in Section 5.7, we introduce the following security notion.

Definition 5.21 (VerSel-INDr Reusable Security). A kpABE scheme for circuit family

$C_\ell = \{C : \{0, 1\}^\ell \rightarrow \{0, 1\}\}$ is said to satisfy **VerSel-IND** reusable security if for all stateful PPT adversary \mathcal{A} , the following holds

$$\Pr \left[\begin{array}{l} (\text{aux}_{\mathcal{A}}, C^1, \dots, C^{Q_{\text{key}}}, \mathbf{x}^1, \dots, \mathbf{x}^{Q_{\text{msg}}}, \mu) \leftarrow \mathcal{A}(1^\lambda); \\ (\text{mpk}, \text{msk}) \leftarrow \text{Setup}(1^\lambda, 1^\ell); \\ \beta' = \beta : (\text{ct}_{\text{off}}, \text{st}) \leftarrow \text{EncOff}(\text{mpk}); \\ \{\text{ct}_{\text{on},0}^j \leftarrow \text{EncOn}(\text{st}, \mathbf{x}^j, \mu), \text{ct}_{\text{on},1}^j \leftarrow C\mathcal{T}_{\text{on}}\}_{j \in [Q_{\text{msg}}]}, \beta \leftarrow \{0, 1\}; \\ \beta' \leftarrow \mathcal{A}(\text{aux}_{\mathcal{A}}, \text{mpk}, \{C^k, \text{sk}_{C^k}\}_{k \in Q_{\text{key}}}, \text{ct}_{\text{off}}, \{\mathbf{x}^j, \text{ct}_{\text{on},\beta}^j\}_{j \in [Q_{\text{msg}}]}) \end{array} \right] \leq \frac{1}{2} + \text{negl}(\lambda)$$

where $C\mathcal{T}_{\text{on}}$ is the ciphertext space of EncOn . We require that for all key queries $C^1, \dots, C^{Q_{\text{key}}}$ and challenge attribute queries $\mathbf{x}^1, \dots, \mathbf{x}^{Q_{\text{msg}}}$ we have $C^k(\mathbf{x}^j) = 0$.

Remark 28. We note that the above security definition implies more standard **VerSel-IND** security for ABE. This can be seen by considering the case of $Q_{\text{msg}} = 1$ and recalling that $\text{ct} = (\text{ct}_{\text{off}}, \text{ct}_{\text{on}})$. The above security definition in this special case implies that the message carrying part of the ciphertext is pseudorandom in **VerSel-IND** security game. This immediately implies **VerSel-IND** security.

We prove the above security of our scheme using the following theorem.

Theorem 5.31. *Assume the PHprFE scheme satisfies reusable security (Definition 5.20) w.r.t. the sampler class containing the sampler Samp as defined in Eq. 5.54 and PRF is secure. Then the above construction of kpABE scheme is secure (Definition 5.21).*

Proof. Consider a PPT adversary \mathcal{A} that outputs $\text{coins}_{\mathcal{A}}, C^1, \dots, C^{Q_{\text{key}}}, \mathbf{x}^1, \dots, \mathbf{x}^{Q_{\text{msg}}}, \mu$. To prove reusable security as per Definition 5.21, we show

$$\left(\begin{array}{l} \text{coins}_{\mathcal{A}}, \text{mpk} = \text{PHprFE.mpk}, \\ \{\text{sk}^k = (\text{PHprFE.sk}^k, \mathbf{r}^k), C^k\}_{k \in [Q_{\text{key}}]}, \\ \text{PHprFE.ct}_{\text{off}}, \{\mathbf{x}^j, \text{PHprFE.ct}_{\text{on}}^j\}_{j \in [Q_{\text{msg}}]} \end{array} \right) \approx_c \left(\begin{array}{l} \text{coins}_{\mathcal{A}}, \text{mpk} = \text{PHprFE.mpk}, \\ \{\text{sk}^k = (\text{PHprFE.sk}^k, \mathbf{r}^k), C^k\}_{k \in [Q_{\text{key}}]}, \\ \text{PHprFE.ct}_{\text{off}}, \{\mathbf{x}^j, \delta^j \leftarrow \text{PHprFE.C}\mathcal{T}_{\text{on}}\}_{j \in [Q_{\text{msg}}]} \end{array} \right) \quad (5.53)$$

Also for all the key queries $C^1, \dots, C^{Q_{\text{key}}}$ and challenge attribute queries $\mathbf{x}^1, \dots, \mathbf{x}^{Q_{\text{msg}}}$

issued by the adversary, we have $C^k(\mathbf{x}^j) = 0$.

We invoke the security of PHprFE scheme with sampler Samp that outputs

$$\left(\begin{array}{ll} \text{circuits:} & \{C^k[\mathbf{r}^k]\}_{k \in [Q_{\text{key}}]}, \\ \text{Inputs:} & \{X_{\text{pub}}^j = \mathbf{x}^j, X_{\text{priv}}^j = (\mu, \mathbf{sd}^j)\}_{j \in [Q_{\text{msg}}]}, \\ \text{Auxiliary Information:} & \text{aux}_{\mathcal{A}} = (\text{coins}_{\mathcal{A}}, C^1, \dots, C^{Q_{\text{key}}}, \mathbf{r}^1, \dots, \mathbf{r}^{Q_{\text{key}}}, \text{coins}_{\mathcal{A}}) \end{array} \right) \quad (5.54)$$

Using the guarantee of PHprFE scheme with sampler Samp we have that

$$\left(\begin{array}{l} \text{aux}_{\mathcal{A}}, \text{PHprFE.mpk}, \{C^k[\mathbf{r}^k]\}_{k \in [Q_{\text{key}}]}, \\ \{ \text{PHprFE.sk}^k \}_{k \in [Q_{\text{key}}]}, \\ \text{PHprFE.ct}_{\text{off}}, \{X_{\text{pub}}^j, \text{PHprFE.ct}_{\text{on}}^j\}_{j \in [Q_{\text{msg}}]} \end{array} \right) \approx_c \left(\begin{array}{l} \text{aux}_{\mathcal{A}}, \text{PHprFE.mpk}, \{C^k[\mathbf{r}^k]\}_{k \in [Q_{\text{key}}]}, \\ \{ \text{PHprFE.sk}^k \}_{k \in [Q_{\text{key}}]}, \\ \text{PHprFE.ct}_{\text{off}}, \{X_{\text{pub}}^j, \delta^j\}_{j \in [Q_{\text{msg}}]} \end{array} \right) \quad (5.55)$$

$$\begin{aligned} \text{if } & \left(\text{aux}_{\mathcal{A}}, \{C^k[\mathbf{r}^k]\}_{k \in [Q_{\text{key}}]}, \{X_{\text{pub}}^j\}_{j \in [Q_{\text{msg}}]}, \{C^k[\mathbf{r}^k](X_{\text{pub}}^j, X_{\text{priv}}^j)\}_{k \in [Q_{\text{key}}], j \in [Q_{\text{msg}}]} \right) \\ & \approx_c \left(\text{aux}_{\mathcal{A}}, \{C^k[\mathbf{r}^k]\}_{k \in [Q_{\text{key}}]}, \{X_{\text{pub}}^j\}_{j \in [Q_{\text{msg}}]}, \{\Delta_{k,j} \leftarrow \{0, 1\}^\lambda\}_{k \in [Q_{\text{key}}], j \in [Q_{\text{msg}}]} \right) \end{aligned} \quad (5.56)$$

where $X_{\text{pub}}^j = \mathbf{x}^j, X_{\text{priv}}^j = (\mu, \mathbf{sd}^j)$ for $\mathbf{sd}^j \leftarrow \{0, 1\}^\lambda, j \in [Q_{\text{msg}}]$,

$(\text{PHprFE.mpk}, \text{PHprFE.msk}) \leftarrow \text{PHprFE.Setup}(1^\lambda, (1^\ell, 1^{\lambda+1}))$,

$\text{PHprFE.sk}^k \leftarrow \text{PHprFE.KeyGen}(\text{PHprFE.msk}, C^k[\mathbf{r}^k])$, for $\mathbf{r}^k \leftarrow \{0, 1\}^\lambda$,

$(\text{PHprFE.ct}_{\text{off}}, \text{PHprFE.st}) \leftarrow \text{PHprFE.EncOff}(\text{PHprFE.mpk})$,

$\text{PHprFE.ct}_{\text{on}}^j \leftarrow \text{PHprFE.Enc}(\text{PHprFE.st}, X_{\text{pub}}^j, X_{\text{priv}}^j)$,

$\delta^j \leftarrow \text{PHprFE.CT}_{\text{on}}$ for $j \in [Q_{\text{msg}}]$ where $\text{PHprFE.CT}_{\text{on}}$ is the ciphertext space of PHprFE.EncOn .

First we note that rearranging the terms of Equation (5.55), it is same as distribution in Equation (5.53). Thus, to prove Equation (5.53), it suffices to prove Equation (5.56).

Equation (5.56) holds from the security of the underlying PRF scheme. To see this note

that

$$C^k[\mathbf{r}^k](X_{\text{pub}}^j, X_{\text{priv}}^j) = C^k[\mathbf{r}^k](\mathbf{x}^j, \mu, \text{sd}^j) = \begin{cases} \mu & \text{if } C^k(\mathbf{x}^j) = 1 \\ \text{PRF}(\text{sd}^j, \mathbf{r}^k) & \text{otherwise.} \end{cases}$$

By the admissibility of the adversary into the kpABE security game, we have $C^k(\mathbf{x}^j) = 0$ for all $k \in [Q_{\text{key}}], j \in [Q_{\text{msg}}]$. So, we get

$$\begin{aligned} C^k[\mathbf{r}^k](\mathbf{x}^j, \mu, \text{sd}^j) &= \text{PRF}(\text{sd}^j, \mathbf{r}^k) \\ &\approx_c \Delta_{j,k} \leftarrow \{0, 1\} \end{aligned}$$

where the last equation follows from the security of the PRF scheme. ■

5.6.3 Predicate Encryption with Optimal Parameters

In this section we sketch out the construction of a PE scheme $\text{PE} = (\text{Setup}, \text{KeyGen}, \text{Enc}, \text{Dec})$ supporting unbounded depth circuits and achieving optimal parameters. The construction is same as that in Section 5.6.1 with the following changes.

1. In the building blocks we use a PHprFE scheme with $\mathcal{X}_{\text{pub}} = \{\perp\}, \mathcal{X}_{\text{priv}} = \{0, 1\}^{\ell+\lambda+1}$ and $\mathcal{Y} = \{0, 1\}$.
2. In $\text{KeyGen}(\text{msk}, C)$ algorithm the circuit $C[\mathbf{r}]$ is defined as follows.
On input $(\perp, \mathbf{x}, \mu, \text{sd})$,

$$C[\mathbf{r}](\perp, \mathbf{x}, \mu, \text{sd}) = \begin{cases} \mu & \text{if } C(\mathbf{x}) = 1 \\ \text{PRF}(\text{sd}, \mathbf{r}) & \text{otherwise.} \end{cases}$$

3. In EncOn algorithm we set $\mathbf{x}_{\text{pub}} = \perp$ and $\mathbf{x}_{\text{priv}} = (\mathbf{x}, \mu, \text{sd})$.
4. In $\text{Dec}(\text{mpk}, \text{sk}_C, C, \text{ct})$ ⁷, we compute $y = \text{PHprFE.Dec}(\text{PHprFE.mpk}, \perp, \text{PHprFE.sk}, C[\mathbf{r}], \text{PHprFE.ct})$.

It is easy to see that the above PE scheme satisfies correctness. The PE scheme also satisfies VerSel-INDr reusable security as defined in Definition 5.21 with a similar security proof as in Section 5.6.2 with the above specified changes. Note that since the

⁷Here we do not give \mathbf{x} as an input.

online part of the encryption $\text{EncOn}(\text{st}, \mathbf{x}, \mu)$ encodes both the attribute and the message μ , replacing it with a random string hides the information of \mathbf{x} and μ and thus it satisfies the security requirement for a PE scheme. As for the efficiency, since we encode \mathbf{x} into the private part of the input, the ciphertext size is $|\mathbf{x}| + \text{poly}(\lambda)$. Summarizing the above discussion, we have the following theorem.

Theorem 5.32. *Under the LWE and Evasive LWE assumption, there exists very selectively secure KP-PE scheme supporting circuits $\{C : \{0, 1\}^\ell \rightarrow \{0, 1\}\}$ with unbounded depth with and single bit message space with*

$$|\text{mpk}| = \text{poly}(\lambda), |\text{sk}_C| = \text{poly}(\lambda), |\text{ct}| = \text{poly}(\lambda) + |\mathbf{x}|. \quad (5.57)$$

where $\mathbf{x} \in \{0, 1\}^\ell$.

5.6.4 Extending the Message Space

Our construction of ABE and PE only allows us to encrypt a single-bit message. Here, we explain how to extend the message space while maintaining the optimal parameter size. In both cases, it suffices to encrypt message of length λ , since this allows us to employ the hybrid encryption approach, where we encrypt a secret key $\text{SKE.sk} \in \{0, 1\}^\lambda$ of an SKE scheme and then use this to encrypt the message.

In the case of ABE, we simply encrypt each bit of SKE.sk , which blows up the ciphertext size by a factor of λ , but this still results in the optimal parameter size of Equation (5.52). In the case of PE, this approach leads to a ciphertext of size $\text{poly}(\lambda) + \lambda|\mathbf{x}|$, which ruins the optimal ciphertext size of only additively depending on the length of the attribute. Instead, we change the construction of PE from PHprFE by setting $\mathbf{x}_{\text{pub}} = \perp$ and $\mathbf{x}_{\text{priv}} = (\mathbf{x}, \text{SKE.sk}, \text{sd})$ and considering a circuit $C[i, \mathbf{r}]$ for $i \in [\lambda]$ that is defined as

$$C[i, \mathbf{r}](\perp, \mathbf{x}, \text{SKE.sk}, \text{sd}) = \begin{cases} \text{SKE.sk}_i & \text{if } C(\mathbf{x}) = 1 \\ \text{PRF}(\text{sd}, \mathbf{r}_i) & \text{otherwise} \end{cases},$$

where SKE.sk_i is the i -th bit of SKE.sk . For generating a secret key of PE for circuit

C , we generate PHprFE secret keys for $C[i, \mathbf{r}_i]$ for all $i \in [\lambda]$ with freshly chosen \mathbf{r}_i . This allows the decryptor to recover SKE.sk in a bit-by-bit manner. It is not difficult to see that the construction achieves optimal parameter size of Equation (5.57) and still maintains the security.

5.7 COMPILING KP-ABE TO CP-ABE USING prFE

In this section we give a compiler that converts kpABE scheme for circuits of unbounded depth to a cpABE scheme for circuits of unbounded depth using a prFE scheme for pseudorandom functionality. In particular, if we start from a kpABE scheme with optimal parameter size, the resulting cpABE scheme achieves the optimal parameter size as well. By combining the kpABE and cpABE, we obtain an ABE scheme for Turing machines from LWE and evasive LWE. This improves [13] in terms of the assumption, which requires the non-standard tensor circular LWE assumption additionally.

5.7.1 Construction

Building Blocks. We require the following building blocks for our construction.

1. A key-policy ABE scheme $\text{kpABE} = (\text{kpABE.Setup}, \text{kpABE.KeyGen}, \text{kpABE.EncOff}, \text{kpABE.EncOn}, \text{kpABE.Dec})$ for circuit class $C_{L(\lambda)} = \{C : \{0, 1\}^\ell \rightarrow \{0, 1\}\}$ consisting of circuits with input length $L(\lambda)$ that satisfies VerSel-INDr reusable security (Definition 5.21). We require that the scheme satisfies optimal parameter size, namely, the size of the master public key, secret keys, and ciphertexts are all fixed polynomial. We denote the online ciphertext space of kpABE scheme by $C\mathcal{T}_{\text{on}} := \{0, 1\}^{\ell_{\text{on}}}$, online ciphertext size by ℓ_{on} . We also assume that the randomness used by kpABE.EncOn is of length λ without loss of generality. If it requires longer randomness, we can derive it by a PRF. We can instantiate such kpABE scheme by our construction in Section 5.6.1.
2. A FE scheme for pseudorandom functionality $\text{prFE} = (\text{prFE.Setup}, \text{prFE.KeyGen}, \text{prFE.Enc}, \text{prFE.Dec})$ for circuit class $C = \{C : \{0, 1\}^L \rightarrow \{0, 1\}^{\ell_{\text{on}}}\}$. Here, the depth of the circuit is unbounded, but the input and the output lengths are fixed. For our compiler, we set $L(\lambda) = \text{poly}(\lambda)$ for some fixed polynomial $\text{poly}(\cdot)$. We require the size of the master public key and the ciphertext to be fixed polynomial in λ and the size of the secret key to be $\ell_{\text{on}} \cdot \text{poly}(\lambda)$. Such a construction can be obtained by applying the conversion described in Section 5.5.5 to our construction of PHprFE in Section 5.5.4. We use $C\mathcal{T}_{\text{prFE}}$ to denote the ciphertext space of the scheme and

$\text{prm} = (1^L, 1^{\ell_{\text{on}}})$ to denote the parameters describing the circuit class.

3. A pseudorandom function $\text{PRF} : \{0, 1\}^\lambda \times \{0, 1\}^\lambda \rightarrow \{0, 1\}^\lambda$. It is known that PRF can be constructed from one-way functions.

Now, we describe our compiler for constructing a ciphertext-policy ABE scheme $\text{cpABE} = (\text{Setup}, \text{KeyGen}, \text{Enc}, \text{Dec})$ for circuits of unbounded depth with attribute length L .

$\text{Setup}(1^\lambda, 1^L) \rightarrow (\text{cpABE.mpk}, \text{cpABE.msk})$. The setup algorithm does the following.

- Run $(\text{prFE.mpk}, \text{prFE.msk}) \leftarrow \text{prFE.Setup}(1^\lambda, \text{prm})$.
- Set $\text{cpABE.mpk} = \text{prFE.mpk}$ and $\text{cpABE.msk} = \text{prFE.msk}$. Output $(\text{cpABE.mpk}, \text{cpABE.msk})$.

$\text{KeyGen}(\text{cpABE.msk}, \mathbf{x}) \rightarrow \text{cpABE.sk}_\mathbf{x}$. The key generation algorithm does the following.

- Parse $\text{cpABE.msk} = \text{prFE.msk}$ and sample $\mathbf{r} \leftarrow \{0, 1\}^\lambda$.
- Define circuit $F[\mathbf{x}, \mathbf{r}]$, with \mathbf{x}, \mathbf{r} hardwired, as follows.
On input $(\text{kpABE.st}, \text{sd}, \mu)$:
 - Compute and output $\text{kpABE.ct}_{\text{on}}$
where $\text{kpABE.ct}_{\text{on}} = \text{kpABE.EncOn}(\text{kpABE.st}, \mathbf{x}, \mu; \text{PRF}(\text{sd}, \mathbf{r}))$.
- Compute $\text{prFE.sk} \leftarrow \text{prFE.KeyGen}(\text{prFE.msk}, F[\mathbf{x}, \mathbf{r}])$.
- Output $\text{cpABE.sk}_\mathbf{x} := (\mathbf{r}, \text{prFE.sk})$.

$\text{Enc}(\text{cpABE.mpk}, C, \mu) \rightarrow \text{cpABE.ct}$. The encryption algorithm does the following.

- Parse $\text{cpABE.mpk} = \text{prFE.mpk}$ and sample a PRF key $\text{sd} \leftarrow \{0, 1\}^\lambda$.
- Generate $(\text{kpABE.mpk}, \text{kpABE.msk}) \leftarrow \text{kpABE.Setup}(1^\lambda, 1^L)$.
- Generate $(\text{kpABE.ct}_{\text{off}}, \text{kpABE.st}) \leftarrow \text{kpABE.EncOff}(\text{kpABE.mpk})$.
- Compute $\text{prFE.ct} \leftarrow \text{prFE.Enc}(\text{prFE.mpk}, (\text{kpABE.st}, \text{sd}, \mu))$.

- Compute $\text{kpABE.sk}_C \leftarrow \text{kpABE.KeyGen}(\text{kpABE.msk}, C)$.
- Output $\text{cpABE.ct} := (\text{prFE.ct}, \text{kpABE.mpk}, \text{kpABE.ct}_{\text{off}}, \text{kpABE.sk}_C)$.

$\text{Dec}(\text{cpABE.mpk}, \text{cpABE.sk}_x, \text{cpABE.ct}, C)$. The decryption algorithm does the following.

- Parse $\text{cpABE.mpk} = \text{prFE.mpk}$, $\text{cpABE.sk}_x = \text{prFE.sk}$, and $\text{cpABE.ct} = (\text{prFE.ct}, \text{kpABE.mpk}, \text{kpABE.ct}_{\text{off}}, \text{kpABE.sk}_C)$.
- Compute $y = \text{prFE.Dec}(\text{prFE.mpk}, \text{prFE.sk}, F[x, r], \text{prFE.ct})$.
- Compute $\text{kpABE.Dec}(\text{kpABE.mpk}, \text{kpABE.sk}_C, C, (\text{kpABE.ct}_{\text{off}}, y), x)$ and output

Correctness. We prove the correctness of our scheme using the following theorem.

Theorem 5.33. *Assume kpABE is perfectly correct. Then the above construction of cpABE scheme is correct.*

Proof. From the correctness of prFE scheme, with probability 1 we have

$$y = F[x, r](\text{kpABE.st}, \text{sd}, \mu) = \text{kpABE.ct}_{\text{on}}$$

where $\text{kpABE.ct}_{\text{on}} = \text{kpABE.EncOn}(\text{kpABE.st}, x, \mu; \text{PRF}(\text{sd}, r))$. Next, from the correctness of kpABE , if $C(x) = 1$, it follows that $\text{kpABE.Dec}(\text{kpABE.mpk}, \text{kpABE.sk}_C, C, \text{kpABE.ct}, x) = \mu$, where $\text{kpABE.ct} = (\text{kpABE.ct}_{\text{off}}, \text{kpABE.ct}_{\text{on}})$. ■

Efficiency. Our cpABE scheme satisfies

$$|\text{cpABE.mpk}| = \text{poly}(\lambda), |\text{cpABE.sk}_x| = \text{poly}(\lambda), |\text{cpABE.ct}| = \text{poly}(\lambda).$$

To see the above we make the following observations:

1. Instantiating kpABE scheme as in Section 5.6.1, we have
 $|\text{kpABE.mpk}| = \text{poly}(\lambda)$, $|\text{kpABE.sk}_C| = \text{poly}(\lambda)$,
 $|\text{kpABE.ct}_{\text{off}}| = |\text{kpABE.ct}_{\text{on}}| = \text{poly}(\lambda)$
2. Instantiating prFE scheme as in Theorem 5.29, we have $|\text{prFE.mpk}| = \text{poly}(\lambda)$,
 $|\text{prFE.sk}| = \ell_{\text{on}} \text{poly}(\lambda) = \text{poly}(\lambda)$, $|\text{prFE.ct}| = L_{\text{inp}} + \text{poly}(\lambda)$, where L_{inp} is
the input length of the prFE scheme. In our construction $L_{\text{inp}} = |\text{kpABE.st}| +$
 $|\text{sd}| + |\mu| = \text{poly}(\lambda)$.

Using the above we note that

- $|\text{cpABE.mpk}| = |\text{prFE.mpk}| = \text{poly}(\lambda)$.
- $|\text{cpABE.sk}_{\mathbf{x}}| = |\mathbf{r}| + |\text{prFE.sk}| = \lambda + \text{poly}(\lambda) = \text{poly}(\lambda)$.
- $|\text{cpABE.ct}| = |\text{prFE.ct}| + |\text{kpABE.mpk}| + |\text{kpABE.ct}_{\text{off}}| + |\text{kpABE.sk}_C| =$
 $\text{poly}(\lambda)$.

We formalise the instantiation using the following theorem.

Theorem 5.34. *Under the LWE and Evasive LWE assumption, there exists very selectively secure CP-ABE scheme supporting circuits with unbounded depth $\{C : \{0, 1\}^\ell \rightarrow \{0, 1\}\}$ and single bit message space with*

$$|\text{mpk}| = \text{poly}(\lambda), |\text{sk}_{\mathbf{x}}| = \text{poly}(\lambda), |\text{ct}| = \text{poly}(\lambda) \quad (5.58)$$

where $\mathbf{x} \in \{0, 1\}^\ell$.

5.7.2 Security

We prove the security of our scheme via the following theorem.

Theorem 5.35. *If the prFE scheme is secure (Definition 5.13, with $Q_{\text{msg}} = 1$), with respect to the sampler class containing the sampler Samp as defined in Eq. 5.59, kpABE scheme satisfies VerSel-INDr reusable security (Definition 5.21) then the construction of cpABE satisfies VerSel-IND security.*

Proof. Suppose the adversary \mathcal{A} with randomness $\text{coins}_{\mathcal{A}}$ queries for $C, \mu, \mathbf{x}_1, \dots, \mathbf{x}_Q$.

We want to prove

$$\left(\begin{array}{l} \text{coins}_{\mathcal{A}}, \text{cpABE.mpk} = \text{prFE.mpk}, \\ \{F[\mathbf{x}_k, \mathbf{r}_k]\}_{k \in [Q]}, \text{cpABE.sk}_{x_k} = \{\text{prFE.sk}_k\}_{k \in [Q]}, \\ \text{cpABE.ct} = (\text{prFE.ct}, \text{kpABE.mpk}, \text{kpABE.ct}_{\text{off}}, \text{kpABE.sk}_C) \end{array} \right) \\ \approx_c \left(\begin{array}{l} \text{coins}_{\mathcal{A}}, \text{cpABE.mpk} = \text{prFE.mpk}, \\ \{F[\mathbf{x}_k, \mathbf{r}_k]\}_{k \in [Q]}, \text{cpABE.sk}_{x_k} = \{\text{prFE.sk}_k\}_{k \in [Q]}, \\ \text{cpABE.ct} = (\Delta, \text{kpABE.mpk}, \text{kpABE.ct}_{\text{off}}, \text{kpABE.sk}_C) \end{array} \right)$$

where $\Delta \leftarrow C\mathcal{T}_{\text{prFE}}$, assuming we have $C(\mathbf{x}_k) = 1$ for all the key queries $\mathbf{x}_1, \dots, \mathbf{x}_Q$ and the challenge circuit C issued by the adversary. Observe that the equation on the left is the view of the adversary in the real world and that on the right is the view of the adversary in the ideal world⁸. Here $F[\mathbf{x}_k, \mathbf{r}_k]$ denotes the functions corresponding to k -th key query \mathbf{x}_k as defined in the KeyGen algorithm, $\text{prFE.sk}_k \leftarrow \text{prFE.KeyGen}(\text{prFE.msk}, F[\mathbf{x}_k, \mathbf{r}_k])$ for $k \in [Q]$, $(\text{kpABE.ct}_{\text{off}}, \text{kpABE.st}) \leftarrow \text{kpABE.EncOff}(\text{kpABE.mpk})$, $\text{prFE.ct} \leftarrow \text{prFE.Enc}(\text{prFE.mpk}, (\text{kpABE.st}, \text{sd}, \mu))$ for $\text{sd} \leftarrow \{0, 1\}^\lambda$ and $\text{kpABE.sk}_C \leftarrow \text{kpABE.KeyGen}(\text{kpABE.msk}, C)$.

We invoke the security of prFE with sampler Samp that outputs

$$\left(\begin{array}{l} \text{Functions:} \quad \{F[\mathbf{x}_k, \mathbf{r}_k]\}_{k \in [Q]}, \\ \text{Input:} \quad (\text{kpABE.st}, \text{sd}, \mu), \\ \text{Auxiliary} \\ \text{Information:} \quad \text{aux} = (\{\mathbf{x}_k, \mathbf{r}_k\}_{k \in [Q]}, \text{kpABE.mpk}, \text{kpABE.sk}_C, \text{kpABE.ct}_{\text{off}}, \text{coins}_{\mathcal{A}}) \end{array} \right) \quad (5.59)$$

By the security guarantee of prFE with sampler Samp we have that

$$\text{if } (\text{aux}, \{F[\mathbf{x}_k, \mathbf{r}_k], F[\mathbf{x}_k, \mathbf{r}_k](\text{kpABE.st}, \text{sd}, \mu)\}_{k \in [Q]}) \approx_c \left(\text{aux}, \{F[\mathbf{x}_k, \mathbf{r}_k], \delta_k \leftarrow \{0, 1\}^{\ell_{\text{on}}}\}_{k \in [Q]} \right)$$

⁸Note that the information about μ is encoded in prFE.ct.

$$\text{then } \left(\begin{array}{c} \text{prFE.mpk, aux, } \{F[\mathbf{x}_k, \mathbf{r}_k], \text{prFE.sk}_k\}_{k \in [Q]} \\ \text{prFE.ct} \leftarrow \text{prFE.Enc}(\text{prFE.mpk}, (\text{kpABE.st}, \text{sd}, \mu)) \end{array} \right) \\ \approx_c \left(\begin{array}{c} \text{prFE.mpk, aux, } \{F[\mathbf{x}_k, \mathbf{r}_k], \text{prFE.sk}_k\}_{k \in [Q]} \\ \Delta \leftarrow C\mathcal{T}_{\text{prFE}} \end{array} \right),$$

where one can see that the latter equation is equivalent to Section 5.7.2. Thus to prove Section 5.7.2, it suffices to prove

$$\left(\text{aux, } \{\mathbf{x}_k, \mathbf{r}_k\}_{k \in [Q]}, \{\text{kpABE.ct}_{\text{on},k} = F[\mathbf{x}_k, \mathbf{r}_k](\text{kpABE.st}, \text{sd}, \mu)\}_{k \in [Q]}, \text{kpABE.sk}_C \right) \\ \approx_c \left(\text{aux, } \{\mathbf{x}_k, \mathbf{r}_k\}_{k \in [Q]}, \{\delta_k \leftarrow \{0, 1\}^{\ell_{\text{on}}}\}_{k \in [Q]}, \text{kpABE.sk}_C \right).$$

For clarity in the further steps of security proof we write the equation on L.H.S. of the above equation as

$$\left(\begin{array}{c} \text{coins}_{\mathcal{A}}, \{\mathbf{x}_k, \mathbf{r}_k\}_{k \in [Q]}, \text{kpABE.mpk}, \text{kpABE.sk}_C, \\ \text{kpABE.ct}_{\text{off}}, \{\text{kpABE.ct}_{\text{on},k} := \text{kpABE.EncOn}(\text{kpABE.st}, \mathbf{x}_k, \mu; \text{PRF}(\text{sd}, \mathbf{r}_k))\}_{k \in [Q]} \end{array} \right), \quad (5.60)$$

where we unroll aux and $F[\mathbf{x}_k, \mathbf{r}_k](\text{kpABE.st}, \text{sd}, \mu)$ based on their definitions. We prove the pseudorandomness of $\{\text{kpABE.ct}_{\text{on},k}\}_k$ in Equation (5.60) via the following sequence of hybrids.

Hyb₀. This is the distribution in Equation (5.60).

Hyb₁. This hybrid is same as the previous one except that we output a failure symbol if the set $\{\mathbf{r}_k\}_{k \in [Q]}$ contains a collision. We prove that the probability with which there occurs a collision is negligible in λ . To prove this it suffices to show that there is no $k, k' \in [Q]$ such that $k \neq k'$ and $\mathbf{r}_k = \mathbf{r}_{k'}$. The probability of this happening can be bounded by $Q^2/2^\lambda$ by taking the union bound with respect to all the combinations of k, k' . Thus the probability of outputting the failure symbol is

$Q^2/2^\lambda$ which is $\text{negl}(\lambda)$.

Hyb₂. This hybrid is same as the previous hybrid except that we change all the PRF values $\{\text{PRF}(\text{sd}, \mathbf{r}_k)\}_k$ to truly random values $\{R_k \leftarrow \{0, 1\}^\lambda\}_k$. By the change introduced in the previous hybrid, $\text{PRF}(\text{sd}, \cdot)$ is invoked on fresh input for each k . Therefore, we can replace $\{\text{PRF}(\text{sd}, \mathbf{r}_k)\}_k$ with truly random $\{R_k\}_k$ by the security of PRF without being noticed by the adversary. We now consider the following distribution:

$$\left(\begin{array}{c} \text{coins}_{\mathcal{A}}, \{\mathbf{x}_k, \mathbf{r}_k\}_{k \in [Q]}, \text{kpABE.mpk}, \text{kpABE.sk}_C, \\ \text{kpABE.ct}_{\text{off}}, \{\text{kpABE.ct}_{\text{on},k} \leftarrow \text{kpABE.EncOn}(\text{kpABE.st}, \mathbf{x}_k, \mu)\}_{k \in [Q]} \end{array} \right),$$

Hyb₃. In this hybrid we invoke the reusable security of kpABE scheme to switch $\{\text{kpABE.ct}_{\text{on},k}\}_k$ to be random strings in $\{0, 1\}^{\ell_{\text{on}}}$.

We claim that an adversary \mathcal{A} who can distinguish Hyb₂ and Hyb₃ can be used to break reusable security of kpABE. The reduction \mathcal{B} works as follows.

- \mathcal{A} sends $\text{coins}_{\mathcal{A}}, C, \mathbf{x}_1, \dots, \mathbf{x}_Q, \mu$ to the reduction.
- \mathcal{B} sends $(C, \{\mathbf{x}_k\}_k, \mu)$ to the kpABE challenger. The challenger does the following.
 - Generates $(\text{kpABE.mpk}, \text{kpABE.msk}) \leftarrow \text{kpABE.Setup}(1^\lambda, 1^\ell)$.
 - Computes $\text{kpABE.sk}_C \leftarrow \text{kpABE.KeyGen}(1^\lambda, C)$ and $(\text{kpABE.ct}_{\text{off}}, \text{kpABE.st}) \leftarrow \text{kpABE.EncOff}(\text{kpABE.mpk})$.
 - Computes $\text{kpABE.ct}_{\text{on},k}^0 \leftarrow \text{kpABE.Enc}(\text{kpABE.st}, \mathbf{x}_k, \mu)$ and $\text{kpABE.ct}_{\text{on},k}^1 \leftarrow \{0, 1\}^{\ell_{\text{on}}}$ for all $k \in [Q]$.
 - Samples a bit $\beta \leftarrow \{0, 1\}$ and returns $(\text{kpABE.mpk}, \text{kpABE.sk}_C, \text{kpABE.ct}_{\text{off}}, \{\text{kpABE.ct}_{\text{on},k}^\beta\}_{k \in [Q]})$ to \mathcal{B} .
- \mathcal{B} returns $(\text{coins}_{\mathcal{A}}, \{\mathbf{x}_k, \mathbf{r}_k\}_{k \in [Q]}, \text{kpABE.mpk}, \text{kpABE.sk}_C, \text{kpABE.ct}_{\text{off}}, \{\text{kpABE.ct}_{\text{on},k}^\beta\}_{k \in [Q]})$ to \mathcal{A} .

- \mathcal{A} outputs a guess bit β' . \mathcal{B} outputs the same bit as its guess.

We note that if the challenger samples $\beta = 0$, then \mathcal{B} simulates Hyb_2 with adversary else it simulates Hyb_3 .

Admissibility of \mathcal{B} . Observe that by the admissibility of cpABE , \mathcal{A} sends challenge queries $C, \mathbf{x}_1, \dots, \mathbf{x}_Q, \mu$ such that $C(\mathbf{x}_k) = 0$ for all $k \in [Q]$. Thus the query $(C, \{\mathbf{x}_k\}_k, \mu)$ sent by \mathcal{B} to the kpABE challenger satisfies $C(\mathbf{x}_k) = 0$ for all $k \in [Q]$. This establishes the admissibility of \mathcal{B} .

We observe that the view of the adversary in Hyb_3 is the same as the R.H.S of Section 5.7.2 and hence the proof. ■

Implications to ABE for Turing Machines. We note that our unbounded CP-ABE scheme Theorem 5.34 and unbounded KP-ABE scheme Theorem 5.30 can be used to instantiate ABE for Turing Machines ([13]).

Corollary 5.35.1. *Under the LWE and evasive LWE assumptions, there exists a very selectively secure ABE for TM with*

$$|\text{mpk}| = \text{poly}(\lambda), \quad |\text{sk}| = |M| \cdot \text{poly}(\lambda), \quad |\text{ct}| = |\mathbf{x}| \cdot t \cdot \text{poly}(\lambda)$$

where the Turing machine M runs on input x for time step t .

[13] uses the LWE, evasive LWE and circular tensor LWE assumptions for their construction with $|\text{mpk}| = \text{poly}(\lambda)$, $|\text{sk}| = \text{poly}(|M|, \lambda)$, $|\text{ct}| = \text{poly}(\lambda, |\mathbf{x}|, t)$.

APPENDIX

5.A BLIND BATCH ENCRYPTION FROM LWE

In this section, we construct a blind batch encryption scheme from LWE that is required for the construction of the laconic pseudorandom poly-domain obfuscation scheme in Section 5.4.

5.A.1 Basic Scheme from LWE

Here, we provide the construction of basic blind batch encryption under LWE assumption. The construction is a slight variant of the hash encryption scheme proposed in [84], where we modify the construction so that it has perfect correctness and satisfies our notion of strong blindness (Definition 5.9). The construction here does not satisfy the efficiency requirement required in Section 5.4. However, we can bootstrap the construction to satisfy the requirement using the conversion in Section 5.A.2.

In the following, we use the rounding function $\lfloor \cdot \rfloor_2 : \mathbb{Z}_q \rightarrow \mathbb{Z}_2$ defined as $\lfloor x \rfloor_2 \stackrel{\text{def}}{=} \left\lfloor \frac{2}{q} \cdot x \right\rfloor \bmod 2$.

Setup($1^\lambda, 1^N$). The setup algorithm does the following.

- Sample $\mathbf{u}_{j,b} \leftarrow \mathbb{Z}_q^n$ for $j \in [N]$ and $b \in \{0, 1\}$.
- Output $\text{crs} := \{\mathbf{u}_{j,b}\}_{j \in [N], b \in \{0,1\}}$.

Gen($\text{crs}, X \in \{0, 1\}^N$). The generation algorithm does the following.

- Parse $\text{crs} = \{\mathbf{u}_{j,b}\}_{j \in [N], b \in \{0,1\}}$.
- Compute $\mathbf{h} := \sum_{j \in [N]} \mathbf{u}_{j,X_j} \in \mathbb{Z}_q^n$ where $X_j \in \{0, 1\}$ is the j -th bit of X .
- Output \mathbf{h} .

SingleEnc($\text{crs}, \mathbf{h}, i, (\mu_0, \mu_1)$). The single encryption algorithm, for $i \in [N]$ and messages $(\mu_0, \mu_1) \in \{0, 1\}^2$, does the following.

- Parse $\text{crs} = \{\mathbf{u}_{j,b}\}_{j \in [N], b \in \{0,1\}}$.
- Sample $\mathbf{s} \leftarrow \mathbb{Z}_q^n$, $e_{j,b} \leftarrow \mathcal{D}_{\mathbb{Z}, \sigma}$ for $j \in [N] \setminus \{i\}$ and $b \in \{0, 1\}$, and $e'_{i,0}, e'_{i,1} \leftarrow \mathcal{D}_{\mathbb{Z}, \gamma}$.
- Compute $c_{j,b} := \mathbf{s}^\top \mathbf{u}_{j,b} + e_{j,b}$ for $j \in [N] \setminus \{i\}$ and $b \in \{0, 1\}$.
- Compute
$$c_{i,b} := \lfloor \mathbf{s}^\top (\mathbf{h} - \mathbf{u}_{i,b}) + e'_{i,b} \rfloor_2 \oplus \mu_b$$
for $b \in \{0, 1\}$.

- Check whether $\mathbf{s}^\top(\mathbf{h} - \mathbf{u}_{i,b}) \in [q/4 - B, q/4 + B] \cup [3q/4 - B, 3q/4 + B]$ holds for $b = 0$ or $b = 1$. If so, replace each of $\{c_{j,b}\}_{j \in [N], b \in \{0,1\}}$ with \perp and set $c_{i,b} = \mu_b$ for $b \in \{0, 1\}$ and some $B > 0$ ⁹. Otherwise, do nothing.
- Output $\mathbf{ct} := \{c_{j,b}\}_{j \in [N], b \in \{0,1\}}$.

SingleDec($\mathbf{crs}, X, i, \mathbf{ct}$). The single decryption algorithm, for $X \in \{0, 1\}^N$ and $i \in [N]$, does the following.

- Parse $\mathbf{crs} = \{\mathbf{u}_{j,b}\}_{j \in [N], b \in \{0,1\}}$, $\mathbf{ct} \rightarrow (\{c_{j,b}\}_{j \in [N], b \in \{0,1\}})$ where $c_{j,b} \in \mathbb{Z}_q \cup \{\perp\}$ for $j \in [N] \setminus \{i\}$ and $c_{i,b} \in \{0, 1\}$ for $b \in \{0, 1\}$.
- If $c_{j,b} = \perp$ for some j, b , set $\mu' := c_{i,X_i}$.
- Otherwise, compute

$$\mu' := c_{i,X_i} \oplus \left[\sum_{j \in [N] \setminus \{i\}} c_{j,X_j} \right]_2.$$

- Output μ' .

Remark 29. We note that the \mathbf{crs} in the above scheme is a uniformly random bit string.

Parameter. We set q to be a power of 2, so that a random string over \mathbb{Z}_q can be interpreted as a binary random string. We set γ super-polynomially larger than σ so that the smudging is possible.

$$q = 2^{5\lambda}, \quad B = 2^{3\lambda}, \quad \sigma = 2^{2\lambda}/N, \quad \gamma = 2^{2\lambda}\lambda^{\omega(1)}$$

Succinctness. We can see that the hash is of fixed length and thus the construction is fully succinct.

⁹We introduce this step to remove the correctness error.

Correctness. With the above set parameters, we show that our scheme achieves perfect correctness.

- If $c_{j,b} = \perp$ for some $j \in [N] \setminus \{i\}$ and $b \in \{0, 1\}$, then $\mu' := c_{i,X_i} = \mu_{X_i}$ with probability 1 and hence the correctness.
- If $c_{j,b} \neq \perp$ for any $j \in [N] \setminus \{i\}$ and $b \in \{0, 1\}$, we compute

$$\begin{aligned}
\mu' &= c_{i,X_i} \oplus \left\lfloor \sum_{j \in [N] \setminus \{i\}} c_{j,X_j} \right\rfloor_2 \\
&= \left\lfloor \mathbf{s}^\top (\mathbf{h} - \mathbf{u}_{i,X_i}) + e'_{i,X_i} \right\rfloor_2 \oplus \mu_{X_i} \oplus \left\lfloor \sum_{j \in [N] \setminus \{i\}} \mathbf{s}^\top \mathbf{u}_{j,X_j} + e_{j,X_j} \right\rfloor_2 \\
&= \left\lfloor \mathbf{s}^\top \left(\sum_{j \in [N]} \mathbf{u}_{j,X_j} - \mathbf{u}_{i,X_i} \right) + e'_{i,X_i} \right\rfloor_2 \oplus \mu_{X_i} \oplus \left\lfloor \sum_{j \in [N] \setminus \{i\}} \mathbf{s}^\top \mathbf{u}_{j,X_j} + e_{j,X_j} \right\rfloor_2 \\
&= \left\lfloor \sum_{j \in [N] \setminus \{i\}} \mathbf{s}^\top \mathbf{u}_{j,X_j} + e'_{i,X_i} \right\rfloor_2 \oplus \mu_{X_i} \oplus \left\lfloor \sum_{j \in [N] \setminus \{i\}} \mathbf{s}^\top \mathbf{u}_{j,X_j} + e_{j,X_j} \right\rfloor_2 \\
&= \mu_{X_i}
\end{aligned}$$

where the last equality follows from our parameter setting, $|e'_{i,X_i}| < B$ and $\sum_{j \in [N] \setminus \{i\}} |e_{j,X_j}| \leq B$ and the guarantee that $\sum_{j \in [N] \setminus \{i\}} \mathbf{s}^\top \mathbf{u}_{j,X_j}$ is sufficiently far, from the second to the last step of the encryption algorithm, from the "unsafe zone" where small noise can change the rounded value (i.e., $[q/4 - B, q/4 + B] \cup [3q/4 - B, 3q/4 + B]$).

Security. Next, we prove the security of our scheme.

Theorem 5.36. *The above construction satisfies SingleEnc security (Definition 5.8) under the LWE assumption.*

Proof. We consider the following sequence of hybrids.

Hyb₀. Real game.

Hyb₁. Change the encryption algorithm so that it never sets $c_{j,b} = \perp$ (i.e., we erase the branch of the computation introduced to eliminate the possibility of the decryption

error). Since $\mathbf{s}^\top (\mathbf{h} - \mathbf{u}_{i,0})$ and $\mathbf{s}^\top (\mathbf{h} - \mathbf{u}_{i,1})$ are distributed uniformly at random over \mathbb{Z}_q , the probability of \perp being output is bounded by $4B/q$. Therefore, this hybrid is statistically indistinguishable from the previous one.

Hyb₂. In this hybrid we compute c_{i,X_i} as

$$c_{i,X_i} = \left[\sum_{j \in [N] \setminus \{i\}} c_{j,X_j} + e'_{i,X_i} \right]_2 \oplus \mu_{X_i}.$$

By the smudging lemma, this hybrid is statistically close to the previous hybrid.

To see this, we note the following.

- In Hyb₁ we have

$$c_{i,X_i} := \left[\mathbf{s}^\top (\mathbf{h} - \mathbf{u}_{i,X_i}) + e'_{i,X_i} \right]_2 \oplus \mu_{X_i}$$

- Substituting $\mathbf{h} = \sum_{j \in [N]} \mathbf{u}_{j,X_j}$, we get

$$c_{i,X_i} = \left[\mathbf{s}^\top \sum_{j \in [N] \setminus \{i\}} \mathbf{u}_{j,X_j} + e'_{i,X_i} \right]_2 \oplus \mu_{X_i}.$$

- From our parameter setting we have $|\sum_{j \in [N] \setminus \{i\}} e_{j,X_j}| \leq \lambda^{\omega(1)} \sum_{j \in [N] \setminus \{i\}} |e_{j,X_j}| \leq |e'_{i,X_i}|$ where $e_{j,b} \in \mathcal{D}_{\mathbb{Z},\sigma}$ and $e'_{i,X_i} \in \mathcal{D}_{\mathbb{Z},\gamma}$. Thus using noise flooding (Theorem 2.3) and Section 5.A.1 we have

$$\begin{aligned} c_{i,X_i} &= \left[\mathbf{s}^\top \sum_{j \in [N] \setminus \{i\}} \mathbf{u}_{j,X_j} + \sum_{j \in [N] \setminus \{i\}} e_{j,X_j} + e'_{i,X_i} \right]_2 \oplus \mu_{X_i} \\ &= \left[\sum_{j \in [N] \setminus \{i\}} (\mathbf{s}^\top \mathbf{u}_{j,X_j} + e_{j,X_j}) + e'_{i,X_i} \right]_2 \oplus \mu_{X_i} \\ &= \left[\sum_{j \in [N] \setminus \{i\}} c_{j,X_j} + e'_{i,X_i} \right]_2 \oplus \mu_{X_i} \end{aligned}$$

with overwhelming probability.

Hyb₃. To compute $c_{i,1-X_i}$, we set it as

$$c_{i,1-X_i} = \lfloor \mathbf{s}^\top (\mathbf{h} - \mathbf{u}_{i,1-X_i}) + e_{i,1-X_i} + e'_{i,1-X_i} \rfloor_2 \oplus \mu_{1-X_i},$$

where $e_{i,1-X_i} \leftarrow \mathcal{D}_{\mathbb{Z},\sigma}$. By the smudging lemma, this hybrid is statistically close to the previous hybrid.

To see this, we note that initially we have

$$c_{i,1-X_i} = \lfloor \mathbf{s}^\top (\mathbf{h} - \mathbf{u}_{i,1-X_i}) + e'_{i,1-X_i} \rfloor_2 \oplus \mu_{1-X_i}$$

where $e'_{i,1-X_i} \in \mathcal{D}_{\mathbb{Z},\gamma}$. Next, from our parameter setting, we have $\lambda^{\omega(1)} |e_{i,1-X_i}| \leq |e'_{i,1-X_i}|$ where $e_{i,1-X_i} \leftarrow \mathcal{D}_{\mathbb{Z},\sigma}$. Thus using noise flooding (Theorem 2.3), we can write Section 5.A.1 as

$$c_{i,1-X_i} = \lfloor \mathbf{s}^\top (\mathbf{h} - \mathbf{u}_{i,1-X_i}) + e_{i,1-X_i} + e'_{i,1-X_i} \rfloor_2 \oplus \mu_{1-X_i}.$$

Hyb₄. Replace $\{c_{j,b}\}_{j \neq i, b \in \{0,1\}}$ and $\mathbf{s}^\top (\mathbf{h} - \mathbf{u}_{i,1-X_i}) + e_{i,1-X_i}$ computed for $c_{i,1-X_i}$ with random elements in \mathbb{Z}_q . This game is computationally indistinguishable from the previous one by LWE.

Hyb₅. Now, $c_{i,1-X_i}$ is sampled as $c_{i,1-X_i} \leftarrow \{0,1\}$. This game is statistically close to the previous game, since the value inside the round function is already uniformly random over \mathbb{Z}_q . ■

Theorem 5.37. *The above construction satisfies strong blindness as per Definition 5.9 under the LWE assumption.*

Proof. The proof is almost the same as that of Theorem 5.36, where we consider the same sequence of the hybrids. We can skip Hyb₂, since c_{i,X_i} is already random by the definition of the blindness security game. In Hyb₅, the ciphertext is a random string.

Therefore, we can conclude the proof of the theorem. ■

5.A.2 Bootstrapping the Basic Scheme

Our construction of BBE in Section 5.A.1 has large CRS size and single ciphertext size, both of which linearly depend on N . However, we require a BBE scheme whose sizes of these parameters are independent of N in Section 5.4. Here, we obtain a scheme with the required properties from LWE by applying the conversion from [58] to our construction in Section 5.A.1.

Theorem 5.38 (Adapted from Appendix A.2 of [58]). *Assuming that there exists a $1/2$ -succinct BBE scheme (defined in Definition 5.7) whose CRS size is $\text{poly}(\lambda, N)$ and single-ciphertext size is $\text{poly}(\lambda, N)$. Then, the construction can be converted into a fully-succinct BBE scheme with CRS size $\text{poly}(\lambda, \log N)$ and single-ciphertext size $\text{poly}(\lambda, \log N)$. Furthermore, the conversion preserves strong blindness property.*

Sketch of proof. Since the above theorem is not explicitly shown in [58], we provide an explanation on how to extract the above theorem from their result. There, they show a construction of fully-succinct BBE scheme starting from a $1/2$ -succinct BBE scheme. They only provide the description of the encryption algorithm, but it is easy to extract a description of single encryption algorithm from it. For the reference for the readers, we sketch their construction and explain how to obtain the single encryption algorithm out of it.

To setup the system, they generate $d = \log(N/\lambda)$ number of CRSes, all of which are generated by $\text{Setup}(1^\lambda, 1^{2^\lambda})$. Each CRS string is assigned to each layer of the tree. To compute a hash value on input $X \in \{0, 1\}^N$, they consider a Merkle tree of depth d . The leaves of the tree consist of N/λ nodes and X is evenly split and assigned to the corresponding node. Then, we assign hash values to the internal nodes of the tree starting from the layer of the tree right above the leaves to the root. To define a hash value h_v associated to a node v , we hash the values associated with its children, namely, $h_{v||0}$ and

$h_{v\|1}$, where we use the CRS corresponding to that layer. The final output of the hash (i.e., $\text{Gen}(\text{crs}, X)$) is the hash value h_ϵ assigned to the root node ϵ .

We then explain how the encryption algorithm works. For each node v , they generate subct_1 part of the underlying BBE ciphertext and a garbled circuit.¹⁰ We denote the former by $\text{subct}_{v,1}$ and the latter \tilde{C}_v . For the garbled circuit corresponding to the root node, we additionally provide the input labels lab_{h_ϵ} that corresponds to h_ϵ . The garbled circuits associated with internal nodes are obtained by garbling a circuit that takes as input h_v and outputs subct_2 part of the BBE ciphertext that encrypts the labels of the garbled circuits corresponding to its children under the public key h_v . For the leaf nodes, the garbled circuits encrypt the messages instead of the labels.

To decrypt a ciphertext, for each leaf v , we traverse the hash tree from the root to v and obtain the message corresponding to X_v as follows. We first obtain $\text{subct}_{\epsilon,2}$ by evaluating \tilde{C}_ϵ on labels lab_{h_ϵ} . Then, combined with $\text{subct}_{\epsilon,1}$, this recovers the entire BBE ciphertext corresponding to the root node ϵ that encrypts the labels of the garbled circuits of the next layer. This BBE ciphertext can be decrypted by using the string that concatenates h_0 and h_1 . This in particular gives the labels $\text{lab}_{h_{v_1}}$ corresponding to h_{v_1} , where v_1 is the first bit of v . This then allows us to recover $\text{subct}_{v_1,2}$ by evaluating the garbled circuit. We traverse down the tree in this way until we reach at the leaf node v , where we recover the corresponding message.

We then explain how we define the single encryption algorithm. To encrypt a message for a position i , we consider the leaf node v corresponding to the index i . We then run the encryption algorithm and remove the ciphertext components that are not necessary for traversing down the tree to the leaf node v . Namely, we only include lab_{h_ϵ} and $\text{subct}_{w,1}, \tilde{C}_w$ for all w that is an ancestor of v in the ciphertext. We can see that these components are sufficient for the decryption to work correctly.

¹⁰We refer to Remark 17 for the explanation on subct_1 and subct_2 .

Now, we can see that the construction achieves CRS size of $\text{poly}(\lambda, \log N)$, since there are $d = \log(N/\lambda)$ number of CRSes of the base BBE. Similarly, the single-ciphertext size is of $\text{poly}(\lambda, \log N)$, since there are d number of garbled circuits and $\text{subct}_{v,1}$, each of which is of fixed polynomial size. We also observe that the conversion preserves the strong blindness. In [58], they show that the ciphertext is pseudorandom except for $\{\text{subct}_{w,1}\}_{w:w \text{ is an ancestor of } v}$. If the underlying BBE satisfies strong blindness, $\text{subct}_{w,1}$ is an empty string for all w . This indicates that the entire ciphertext of the final scheme is pseudorandom, as desired. This holds even if we consider single-ciphertext, since the single ciphertext is obtained by removing some part of the ciphertext. ■

By applying the conversion to our construction in Section 5.A.1, we obtain the following theorem.

Theorem 5.39. *There exists a fully-succinct BBE scheme with CRS size of $\text{poly}(\lambda, \log N)$ and single-ciphertext size of $\text{poly}(\lambda, \log N)$ from LWE where N denotes the length of the keys supported by the scheme*

5.B BOOTSTRAPPING AB-LFE TO KP-ABE WITH UNBOUNDED DEPTH USING prFE

In this section, we show a formal description of the construction of **kpABE** for unbounded circuits using **1ABE** sketched in Section 5.1.4. This provides an alternative pathway to obtain **kpABE** for unbounded depth circuits different from that given in Section 5.6. Instead of **1ABE**, our construction is formally described using **abLFE**, but this turns out to be almost equivalent, as discussed later in this section. We can instantiate the **abLFE** by the recent construction by [122] or blind garbled circuit [58]. The former instantiation leads to more efficient construction than the latter, but it introduces an additional assumption of circular LWE in addition to LWE and evasive LWE. Compared with Section 5.6, the constructions obtained here are simpler, though their parameter sizes are sub-optimal. Importantly, we do not rely on the result regarding **pPRIO** from

our companion paper [14] in this section.

5.B.1 Attribute Based Laconic Functional Encryption

Syntax. An attribute based laconic function evaluation (abLFE) scheme for a circuit class $\{C_{\text{prm}} = \{C : \mathcal{X}_{\text{prm}} \rightarrow \{0, 1\}\}_{\text{prm}}$ for a parameter $\text{prm} = \text{prm}(\lambda)$ and a message space \mathcal{M} consists of four algorithms (crsGen , Compress , Enc , Dec) defined as follows.

$\text{crsGen}(1^\lambda, \text{prm}) \rightarrow \text{crs}$. The generation algorithm takes as input the security parameter 1^λ and circuit parameters prm and outputs a uniformly sampled common reference string crs .

$\text{Compress}(\text{crs}, C) \rightarrow \text{digest}$. The compress algorithm takes as input the common random string crs and a circuit $C \in \mathcal{C}$ and outputs a digest digest .

$\text{Enc}(\text{crs}, \text{digest}, (\mathbf{x}, \mu)) \rightarrow \text{ct}$. The encryption algorithm takes as input the common random string crs , a digest digest , an attribute $\mathbf{x} \in \mathcal{X}_{\text{prm}}$ and a message $\mu \in \mathcal{M}$ and outputs a ciphertext ct .

$\text{Dec}(\text{crs}, C, \text{ct}) \rightarrow \mu/\perp$. The decryption algorithm takes as input the common random string crs , a circuit C , digest and a ciphertext ct and outputs a message $\mu \in \mathcal{M}$ or \perp .

Definition 5.22 (Correctness). An abLFE scheme for circuit family \mathcal{C}_{prm} is correct if for all prm , $C \in \mathcal{C}_{\text{prm}}$, $\mathbf{x} \in \mathcal{X}_{\text{prm}}$ such that $C(\mathbf{x}) = 1$, and for all messages $\mu \in \mathcal{M}$,

$$\Pr \left[\begin{array}{l} \text{crs} \leftarrow \text{crsGen}(1^\lambda, \text{prm}), \\ \text{digest} = \text{Compress}(\text{crs}, C), \\ \text{ct} \leftarrow \text{Enc}(\text{crs}, \text{digest}, (\mathbf{x}, \mu)) : \\ \text{Dec}(\text{crs}, C, \text{ct}) \neq \mu \end{array} \right] = \text{negl}(\lambda)$$

where the probability is taken over the coins of Setup , KeyGen , and Enc .

Definition 5.23 (Pseudorandom Ciphertext Security). For a abLFE scheme and an adversary \mathcal{A} , we define the experiment for security $\text{Expt}_{\beta, \mathcal{A}}^{\text{abLFE}}(1^\lambda)$ as follows.

1. Run \mathcal{A} to receive circuit parameters prm . Run $\text{crs} \leftarrow \text{crsGen}(1^\lambda, \text{prm})$ and send crs to \mathcal{A} .
2. \mathcal{A} chooses $C \in \mathcal{C}_{\text{prm}}$, $\mathbf{x} \in \mathcal{X}_{\text{prm}}$ and $\mu \in \mathcal{M}$. Run $\text{digest} = \text{Compress}(\text{crs}, C)$, sample $\beta \leftarrow \{0, 1\}$. If $\beta = 0$, it computes $\text{ct}_0 \leftarrow \text{Enc}(\text{crs}, \text{digest}, (\mathbf{x}, \mu))$ else if $\beta = 1$, it computes $\text{ct}_1 \leftarrow \mathcal{CT}_{\text{abLFE}}$, where $\mathcal{CT}_{\text{abLFE}}$ is the ciphertext space of abLFE . It sends $\text{digest}, \text{ct}_\beta$ to \mathcal{A} .
3. \mathcal{A} outputs a guess bit β' as the output of the experiment.

We define the advantage $\text{Adv}_{\mathcal{A}}^{\text{abLFE}}(\lambda)$ of \mathcal{A} in the above game as

$$\text{Adv}_{\mathcal{A}}^{\text{abLFE}}(\lambda) := \left| \Pr[\text{Expt}_{0, \mathcal{A}}^{\text{abLFE}}(1^\lambda) = 1] - \Pr[\text{Expt}_{1, \mathcal{A}}^{\text{abLFE}}(1^\lambda) = 1] \right|.$$

We say that a abLFE scheme is *adaptive* pseudorandom ciphertext secure if for every *admissible* PPT adversary \mathcal{A} , we have $\text{Adv}_{\mathcal{A}}^{\text{abLFE}}(\lambda) \leq \text{negl}(\lambda)$, where \mathcal{A} is said to be *admissible* if $C(\mathbf{x}) = 0$.

The *selective* (resp. *very selective*) notion of the security requires the adversary \mathcal{A} to choose \mathbf{x} (resp. \mathbf{x}, C) along with prm before it receives crs .

Definition 5.24 (Decomposability). We say that a abLFE scheme for a circuit class $\{\mathcal{C}_{\text{prm}} = \{C : \mathcal{X}_{\text{prm}} \rightarrow \{0, 1\}\}\}_{\text{prm}}$ satisfies decomposability if for any $\text{crs} \leftarrow \text{crsGen}(1^\lambda, \text{prm})$ and $\text{digest} \leftarrow \text{Compress}(\text{crs}, C)$, we have $\text{digest} = \{\text{digest}_i\}_{i \in [q_C]}$ for some polynomial q_C , which may depend on C , and $\text{size}(\text{digest}_i) \leq \text{poly}(\lambda)$. Furthermore, we have that $\text{Enc}(\text{crs}, \text{digest}, (\mathbf{x}, \mu)) = \{\text{Enc}_i(\text{crs}, \text{digest}_i, (\mathbf{x}, \mu))\}_{i \in [q_C]}$ where size of the encryption circuit $\text{size}(\text{Enc}_i(\cdot, \cdot, (\cdot, \cdot))) \leq \text{poly}(\lambda, |\mathbf{x}|)$.

Remark 30. Here, we do not require the digest to be much smaller than the circuit description C , unlike the usual convention in the context of abLFE . This relaxation allows us to instantiate abLFE using blind garbled circuits, which do not have compact digests.

5.B.2 Construction of kpABE with Unbounded Depth

Building Blocks. We require the following building blocks for our construction.

1. An attribute-based laconic function evaluation scheme $\text{abLFE} = (\text{crsGen}, \text{Compress}, \text{Enc}, \text{Dec})$ for circuit class $C_{L(\lambda)}$, consisting of circuits with input length $L(\lambda)$ and with unbounded depth and size. We let $\ell_{\text{ct}}^{\text{abLFE}}$ and $C\mathcal{T}_{\text{abLFE}} = \{0, 1\}^{\ell_{\text{ct}}^{\text{abLFE}}}$ denote the ciphertext length and the ciphertext space of the scheme, respectively. We assume that the abLFE scheme is decomposable (Definition 5.24), i.e., we have $\text{abLFE.Enc} = \{\text{abLFE.Enc}_i\}_{i \in [q_C]}$ for some q_C and use d_{abLFE} to denote the maximum depth of a circuit required to compute $\{\text{abLFE.Enc}_i\}_{i \in [q_C]}$.
2. A FE scheme for pseudorandom functionality $\text{prFE} = (\text{prFE.Setup}, \text{prFE.KeyGen}, \text{prFE.Enc}, \text{prFE.Dec})$ for circuit class $C_{L(\lambda), d_{\text{prFE}}(\lambda), \ell_{\text{ct}}^{\text{abLFE}}}$ consisting of circuits with input length $L(\ell, \lambda)$, maximum depth $d_{\text{prFE}}(\lambda)$ and output length $\ell_{\text{ct}}^{\text{abLFE}}$. We denote by prm the parameters $(1^{L(\lambda)}, 1^{d_{\text{prFE}}(\lambda)}, 1^{\ell_{\text{ct}}^{\text{abLFE}}})$ that specifies the function class being supported. We also denote the ciphertext space of the scheme by $C\mathcal{T}_{\text{prFE}}$.
3. A PRF function $\text{PRF} : \{0, 1\}^\lambda \times \{0, 1\}^\lambda \rightarrow \{0, 1\}^{R_{\text{len}}}$ where R_{len} is the length of randomness used in abLFE.Enc . We assume that PRF can be computed by a circuit of depth at most d_{prFE} .

We assume that uniform sampling from the ciphertext space is possible without any parameter other than the security parameter λ .

Now, we describe our compiler for constructing a key-policy ABE scheme $\text{kpABE} = (\text{Setup}, \text{KeyGen}, \text{Enc}, \text{Dec})$ for circuits of unbounded depth with attribute length $L(\lambda)$. We denote the ciphertext space of the scheme by $C\mathcal{T}_{\text{kpABE}}$. For our construction, we have $C\mathcal{T}_{\text{kpABE}} = C\mathcal{T}_{\text{prFE}}$.

$\text{Setup}(1^\lambda, 1^L) \rightarrow (\text{mpk}, \text{msk})$. The setup algorithm does the following.

- Run $(\text{prFE.msk}, \text{prFE.mpk}) \leftarrow \text{prFE.Setup}(1^\lambda, \text{prm})$ and $\text{crs} \leftarrow \text{abLFE.crsGen}(1^\lambda)$.
- Set $\text{msk} = \text{prFE.msk}$ ¹¹ and $\text{mpk} = (\text{prFE.mpk}, \text{crs})$. Output (msk, mpk) .

¹¹W.L.O.G we assume that msk contains mpk .

$\text{KeyGen}(\text{msk}, C) \rightarrow \text{sk}_C$. The key generation algorithm does the following.

- Parse $\text{msk} = \text{prFE.msk}$ and sample $\mathbf{r} \leftarrow \{0, 1\}^\lambda$.
- Compute $\text{digest} = \text{abLFE.Compress}(\text{crs}, C)$. Parse $\text{digest} = \{\text{digest}_i\}_{i \in [q_C]}$.
- For $i \in [q_C]$, define circuit $F[\text{crs}, \text{digest}_i, \mathbf{r}]$, with $\text{crs}, \text{digest}_i, \mathbf{r}$ hardwired, as follows
On input $(\text{sd}, \mathbf{x}, \mu)$:
 - Compute $\text{abLFE.ct}_i := \text{abLFE.Enc}_i(\text{crs}, \text{digest}_i, (\mathbf{x}, \mu); \text{PRF}(\text{sd}, \mathbf{r}))$.
 - Output abLFE.ct_i .
- For $i \in [q_C]$ compute $\text{prFE.sk}_i \leftarrow \text{prFE.KeyGen}(\text{prFE.msk}, F[\text{crs}, \text{digest}_i, \mathbf{r}])$.
- Output $\text{sk}_C := (\{\text{digest}_i, \text{prFE.sk}_i\}_{i \in [q_C]}, \mathbf{r})$.

$\text{Enc}(\text{mpk}, \mathbf{x}, \mu) \rightarrow \text{ct}$. The encryption algorithm does the following.

- Parse $\text{mpk} = (\text{prFE.mpk}, \text{crs})$ and sample a PRF key $\text{sd} \leftarrow \{0, 1\}^\lambda$.
- Compute $\text{prFE.ct} \leftarrow \text{prFE.Enc}(\text{prFE.mpk}, (\text{sd}, \mathbf{x}, \mu))$.
- Output $\text{ct} := \text{prFE.ct}$.

$\text{Dec}(\text{mpk}, \text{sk}_C, C, \text{ct}, \mathbf{x}) \rightarrow \mathbf{y}$. The decryption algorithm does the following.

- Parse $\text{mpk} = (\text{prFE.mpk}, \text{crs})$, $\text{sk}_C = (\{\text{digest}_i, \text{prFE.sk}_i\}_{i \in [q_C]}, \mathbf{r})$ and $\text{ct} = \text{prFE.ct}$.
- For all $i \in [q_C]$, compute $y_i = \text{prFE.Dec}(\text{prFE.mpk}, \text{prFE.sk}_i, F[\text{crs}, \text{digest}_i, \mathbf{r}], \text{prFE.ct})$.
- Set $\mathbf{y} = (y_1, \dots, y_{q_C})$ and output $\text{abLFE.Dec}(\text{crs}, C, \mathbf{y})$.

Correctness. We prove the correctness of our scheme using the following theorem.

Theorem 5.40. *Assume abLFE and prFE schemes are correct, and PRF is secure.*

Then the above construction of kpABE scheme is correct.

Proof. From the correctness of prFE scheme we have

$$\begin{aligned} y_i &= F[\text{crs}, \text{digest}_i, \mathbf{r}](\text{sd}, \mathbf{x}, \mu) \\ &= \text{abLFE.ct}_i = \text{abLFE.Enc}_i(\text{crs}, \text{digest}_i, (\mathbf{x}, \mu); \text{PRF}(\text{sd}, \mathbf{r})) \end{aligned}$$

for $i \in [q_C]$ with probability 1. Thus we have $\mathbf{y} = \{\text{abLFE.ct}_i\}_{i \in [q_C]} = \text{abLFE.ct}$. Next, by the correctness of abLFE scheme it follows that, if $C(\mathbf{x}) = 1$,

$$\text{abLFE.Dec}(\text{crs}, C, \mathbf{y}) = \text{abLFE.Dec}(\text{crs}, C, \text{abLFE.ct}) = \mu$$

with all but negligible probability. ■

Security. We prove the security of our scheme via the following theorem.

Theorem 5.41. *Assume that the prFE scheme is secure (Definition 5.13) with respect to the sampler class containing the sampler Samp as defined in Eq. 5.62, abLFE scheme satisfies very selective pseudorandom ciphertext security (Definition 5.23). Then our construction of kpABE scheme satisfies VerSel-INDr security (Definition 2.8).*

Proof. Suppose the adversary \mathcal{A} with randomness $\text{coins}_{\mathcal{A}}$ queries for $\mathbf{x}, \mu, C_1, \dots, C_Q$.

To prove the security of kpABE scheme as per Definition 2.8, we show

$$\left(\begin{array}{l} \text{coins}_{\mathcal{A}}, \text{mpk} = (\text{prFE.mpk}, \text{crs}), \\ \{F[\text{crs}, \text{digest}_{k,i}, \mathbf{r}_k]\}_{k \in [Q], i \in [q_{C_k}]}, \\ \text{sk}_{C_k} = \{\text{prFE.sk}_{k,i}\}_{k \in [Q], i \in [q_{C_k}]}, \\ \text{ct} = \text{prFE.ct} \leftarrow \text{prFE.Enc}(\text{prFE.mpk}, (\text{sd}, \mathbf{x}, \mu)) \end{array} \right) \approx_c \left(\begin{array}{l} \text{coins}_{\mathcal{A}}, \text{mpk} = (\text{prFE.mpk}, \text{crs}), \\ \{F[\text{crs}, \text{digest}_{k,i}, \mathbf{r}_k]\}_{k \in [Q], i \in [q_{C_k}]}, \\ \text{sk}_{C_k} = \{\text{prFE.sk}_{k,i}\}_{k \in [Q], i \in [q_{C_k}]}, \\ \text{ct} = \delta \leftarrow \mathcal{CT}_{\text{prFE}} \end{array} \right) \quad (5.61)$$

where $F[\text{crs}, \text{digest}_{k,i}, \mathbf{r}_k]$ denotes the functions corresponding to k -th key query C_k as defined in the KeyGen algorithm and $\text{prFE.sk}_{k,i} \leftarrow \text{prFE.KeyGen}(\text{prFE.msk}, F[\text{crs}, \text{digest}_{k,i}, \mathbf{r}_k])$ for $k \in [Q], i \in [q_{C_k}]$. Also for all the key queries C_1, \dots, C_Q and the challenge attribute \mathbf{x} issued by the adversary, we have $C_k(\mathbf{x}) = 0$.

We invoke the security of prFE with sampler Samp that outputs

$$\left(\begin{array}{ll} \text{Functions:} & \{F[\text{crs}, \text{digest}_{k,i}, \mathbf{r}_k]\}_{k \in [Q], i \in [q_{C_k}]}, \\ \text{Input:} & (\text{sd}, \mathbf{x}, \mu), \\ \text{Auxiliary Information:} & \text{aux} = (\text{coins}_{\mathcal{A}}, \{\text{crs}, C_k, \text{digest}_{k,i}, \mathbf{r}_k\}_{k \in [Q], i \in [q_{C_k}]}) \end{array} \right) \quad (5.62)$$

By the security guarantee of prFE with sampler Samp we have

$$\begin{aligned} & \left(\begin{array}{l} \text{prFE.mpk}, \text{aux}, \{F[\text{crs}, \text{digest}_{k,i}, \mathbf{r}_k], \text{prFE.sk}_{k,i}\}_{k \in [Q], i \in [q_{C_k}]} \\ \text{prFE.ct} \leftarrow \text{prFE.Enc}(\text{prFE.mpk}, (\text{sd}, \mathbf{x}, \mu)) \end{array} \right) \\ & \approx_c \left(\begin{array}{l} \text{prFE.mpk}, \text{aux}, \{F[\text{crs}, \text{digest}_{k,i}, \mathbf{r}_k], \text{prFE.sk}_{k,i}\}_{k \in [Q], i \in [q_{C_k}]} \\ \Delta \leftarrow C\mathcal{T}_{\text{prFE}} \end{array} \right) \end{aligned}$$

if

$$\begin{aligned} & \left(\text{aux}, \{F[\text{crs}, \text{digest}_{k,i}, \mathbf{r}_k], \text{abLFE.ct}_{k,i}\}_{k \in [Q], i \in [q_{C_k}]} \right) \\ & \approx_c \left(\text{aux}, \{F[\text{crs}, \text{digest}_{k,i}, \mathbf{r}_k], (\delta_{k,1}, \dots, \delta_{k,q_{C_k}}) \leftarrow \{0, 1\}^{\ell_{\text{ct}}^{\text{abLFE}}}\}_{k \in [Q]} \right) \quad (5.63) \end{aligned}$$

where

$$\text{abLFE.ct}_{k,i} = F[\text{crs}, \text{digest}_{k,i}, \mathbf{r}_k](\text{sd}, \mathbf{x}, \mu) = \text{abLFE.Enc}_i(\text{crs}, \text{digest}_{k,i}, (\mathbf{x}, \mu); \text{PRF}(\text{sd}, \mathbf{r}_k)) \text{ for } k \in [Q], i \in [q_{C_k}].$$

Thus to prove Equation (5.61), it suffices to prove Equation (5.63). We prove Equation (5.63) via the following sequence of hybrids.

Hyb₀. This is the LHS distribution of Equation (5.63).

$$\left(\text{aux}, \left\{ \begin{array}{l} F[\text{crs}, \text{digest}_{k,i}, \mathbf{r}_k], \\ \text{abLFE.ct}_{k,i} = \text{abLFE.Enc}_i(\text{crs}, \text{digest}_{k,i}, (\mathbf{x}, \mu); \text{PRF}(\text{sd}, \mathbf{r}_k)) \end{array} \right\}_{k \in [Q], i \in [q_{C_k}]} \right).$$

We can rewrite the above distribution as

$$\left(\text{aux}, \{F[\text{crs}, \text{digest}_{k,i}, \mathbf{r}_k], \text{abLFE.ct}_k\}_{k \in [Q], i \in [q_{C_k}]} \right).$$

where $\text{digest}_k = \{\text{digest}_{k,i}\}_{i \in [q_{C_k}]}$ and $\text{abLFE.ct}_k = \{\text{abLFE.ct}_{k,i}\}_{i \in [q_{C_k}]}$ for all $k \in [Q]$.

Hyb₁. This hybrid is same as the previous one except that we output a failure symbol if the set $\{\mathbf{r}_k\}_{k \in [Q]}$, in **aux**, contains a collision. We prove that the probability with which there occurs a collision is negligible in λ . To prove this it suffices to show that there is no $k, k' \in [Q]$ such that $k \neq k'$ and $\mathbf{r}_k = \mathbf{r}_{k'}$. The probability of this happening can be bounded by $Q^2/2^\lambda$ by taking the union bound with respect to all the combinations of k, k' . Thus the probability of outputting the failure symbol is $Q^2/2^\lambda$ which is $\text{negl}(\lambda)$.

Hyb₂. In this hybrid we change all the PRF values computed using **sd** to random. Namely, we replace $\text{PRF}(\text{sd}, \mathbf{r}_k)$ with true randomness R_k . Since PRF is invoked for fresh input for each $k \in [Q]$, this hybrid is indistinguishable from the previous hybrid. We now consider the following distribution:

$$\left(\text{aux}, \{F[\text{crs}, \text{digest}_{k,i}, \mathbf{r}_k], \text{abLFE.ct}_k = \text{abLFE.Enc}(\text{crs}, \text{digest}_k, (\mathbf{x}, \mu))\}_{k \in [Q], i \in [q_{C_k}]} \right)$$

Hyb₃. In this hybrid we invoke the security of **abLFE** scheme to switch abLFE.ct_k to random for all $k \in [Q]$. Namely, the distribution is now:

$$\left(\{F[\text{crs}, \text{digest}_{k,i}, \mathbf{r}_k], \text{abLFE.ct}_k \leftarrow \mathcal{CT}_{\text{abLFE}}\}_{k \in [Q], i \in [q_{C_k}]} \right)$$

By the admissibility of the adversary and very selective pseudorandom ciphertext security of **abLFE**, this hybrid is indistinguishable from the previous one. This completes the proof. ■

Instantiating the **abLFE** scheme as above, and using a **prFE** scheme supporting $d_{\text{prFE}} = \text{poly}(\lambda)$ depth circuits with input length $L = L + \lambda + 1$, we obtain the following theorem.

Theorem 5.42. *Under the LWE assumption and the Evasive LWE assumption, there*

exists a very selectively secure **kpABE** scheme for circuits of unbounded depth and attribute length ℓ with

$$|\text{mpk}| = \ell \cdot \text{poly}(\lambda), \quad |\text{sk}_C| = |C| \cdot \ell \cdot \text{poly}(\lambda), \quad |\text{ct}| = \ell \cdot \text{poly}(\lambda).$$

We note that the **kpABE** scheme instantiated as above has longer secret keys but is not based on any circular assumptions.

5.B.3 Using the **abLFE** from HLL

We can directly instantiate the **abLFE** in Section 5.B.2 with the construction shown by [122] (HLL henceforth). For this instantiation, we do not assume decomposability of **abLFE** since the HLL construction has succinct digest, i.e. $q_C = 1$ for any C and $|\text{digest}| = O(1)$. This instantiation leads to the parameter size of $|\text{mpk}| = \text{poly}(\ell, \lambda)$, $|\text{sk}_C| = \text{poly}(\ell, \lambda)$, and $|\text{ct}| = \text{poly}(\ell, \lambda)$, which is already better than the previous instantiation.

By adding a twist to the construction in Section 5.B.2 exploiting the structural property of HLL, we can improve the secret key size so that its dependency on ℓ can be removed. To do so, we exploit the online-offline structure of the **abLFE.Enc** algorithm which can be split as **abLFE.Enc** = (**abLFE.EncOff**, **abLFE.EncOn**). Here, **abLFE.EncOff** takes as input **crs** and **x** and outputs the offline part of the ciphertext **abLFE.ct_{off}** and short state **st** and **abLFE.EncOn** takes as input **digest** and **st** and outputs online part of the ciphertext **abLFE.ct_{on}**. Formally, HLL proved the following theorem:

Theorem 5.43 ([122]). *Under the circular LWE assumption, there exists a very selectively secure **abLFE** scheme for circuit class $C = \{C : \{0, 1\}^\ell \rightarrow \{0, 1\}^{\ell'}\}$ satisfying*

$$|\text{crs}| = O(\ell, \lambda), \quad |\text{digest}| = O(\lambda), \quad |\text{st}| = O(\lambda), \quad |\text{ct}_{\text{off}}| = O(\ell, \lambda), \quad |\text{ct}_{\text{on}}| = O(\ell', \lambda)$$

We make slight modifications to our construction of **kpABE** scheme to optimize the secret key size. The high level idea is very simple. Instead of letting the **prFE** decryption recover the entire **abLFE** ciphertext, we recover only the online part of it. We then put the

offline part of **abLFE** ciphertext into the ciphertext of **kpABE** so that the decryptor can recover the entire **abLFE** ciphertext during the decryption. This eliminates the necessity of hardwiring **crs** to the **prFE** secret key, since the online part can be computed only from the short state and digest. This leads to the improvement on the efficiency, since the state and digest are of fixed polynomial size, while **crs** is of size $O(\ell)$. Concretely, we modify the **KeyGen**, **Enc**, **Dec** algorithm of Section 5.B.2 as follows.

KeyGen(**msk**, C). This is same as the **KeyGen** algorithm in Section 5.B.2 except the following.

- Define circuit $F[\text{digest}, \mathbf{r}]$ (instead of $F[\text{crs}, \text{digest}_i, \mathbf{r}]$) , with **digest**, **r** hardwired, as follows
On input (**sd**, **st**):
 - Compute $\text{abLFE.ct}_{\text{on}} := \text{abLFE.EncOn}(\text{st}, \text{digest}; \text{PRF}(\text{sd}, \mathbf{r}))$.
 - Output $\text{abLFE.ct}_{\text{on}}$.
- Compute $\text{prFE.sk} \leftarrow \text{prFE.KeyGen}(\text{prFE.msk}, F[\text{digest}, \mathbf{r}])$.
- Output $\text{sk}_C = (\text{digest}, \text{prFE.sk}, \mathbf{r})$.

Enc(**mpk**, **x**, μ). The encryption algorithm does the following.

- Parse $\text{mpk} = (\text{prFE.mpk}, \text{crs})$ and sample a PRF key $\text{sd} \leftarrow \{0, 1\}^\lambda$.
- Compute $(\text{abLFE.ct}_{\text{off}}, \text{st}) \leftarrow \text{abLFE.EncOff}(\text{crs}, (\mathbf{x}, \mu))$.
- Compute $\text{prFE.ct} \leftarrow \text{prFE.Enc}(\text{prFE.mpk}, (\text{sd}, \text{st}))$.
- Output $\text{ct} := (\text{abLFE.ct}_{\text{off}}, \text{prFE.ct})$.

Dec(**mpk**, sk_C , C , **ct**, **x**). The decryption algorithm does the following.

- Parse $\text{mpk} = (\text{prFE.mpk}, \text{crs})$, $\text{sk}_C = (\text{digest}, \text{prFE.sk}, \mathbf{r})$ and $\text{ct} = (\text{abLFE.ct}_{\text{off}}, \text{prFE.ct})$.
- Compute
 $\text{abLFE.ct}_{\text{on}} = \text{prFE.Dec}(\text{prFE.mpk}, \text{prFE.sk}, F[\text{digest}, \mathbf{r}], \text{prFE.ct})$.

- Set $\mathbf{y} = (\text{abLFE.ct}_{\text{off}}, \text{abLFE.ct}_{\text{on}})$ and output $\text{abLFE.Dec}(\text{crs}, C, \mathbf{y})$.

We note that even with the above changes the correctness and security arguments are same as that of Section 5.B.2.

For this instantiation, we have $d_{\text{abLFE}}^{\text{EncOn}} = \text{poly}(\lambda)$ where $d_{\text{abLFE}}^{\text{EncOn}}$ is the maximum depth of a circuit required to compute abLFE.EncOn and hence we use a prFE scheme supporting $d_{\text{prFE}} = \text{poly}(\lambda)$ depth circuits with input length $L = \text{poly}(\lambda)$. We formalise this using the following theorem.

Theorem 5.44. *Under the circular LWE assumption and the Evasive LWE assumption, there exists a very selectively secure kpABE scheme for circuits of unbounded depth and attribute length ℓ with*

$$|\text{mpk}| = \text{poly}(\ell, \lambda), \quad |\text{sk}_C| = \text{poly}(\lambda), \quad |\text{ct}| = \text{poly}(\ell, \lambda).$$

We note that the kpABE scheme instantiated as above has succinct keys and ciphertexts. Our scheme achieves the same parameters as the unbounded depth KP-ABE scheme by [122] but does not make use of the circular evasive LWE assumption as they do.

CHAPTER 6

PSEUDORANDOM MULTI-INPUT FUNCTIONAL ENCRYPTION AND APPLICATIONS

6.1 INTRODUCTION

Multi-Input FE and iO . A compelling extension of functional encryption, introduced by Goldwasser et al. is multi-input FE (MIFE) [101] where multiple parties can independently encrypt their data and the key generator can provide a function key that jointly decrypts all the ciphertexts. In more detail, now we have n parties, each of who independently computes the ciphertext for its data \mathbf{x}_i , for $i \in [n]$, the key generator provides a key for an n -ary function f and decryption allows to recover $f(\mathbf{x}_1, \dots, \mathbf{x}_n)$. Evidently for real world applications of computing on encrypted data, it is highly desirable to be able to compute on data generated by different parties independently – this can enable important statistical functionalities such as running medical research algorithms on encrypted genomic or medical data. As discussed in the original work of MIFE, the notion is more meaningful in the symmetric rather than public key setting, since the latter allows for too much leakage on the challenge message by dint of legitimate combinations with messages chosen by the adversary. In two concurrent, influential works [29, 41] it was shown that single input FE, if it supports sufficiently expressive functionality and satisfies an efficiency property called *compactness*, is powerful enough to generically imply multi-input FE. We remark that for restricted functionalities, the compiler does not apply and multi-input FE for interesting restricted classes have been studied extensively [4, 81, 3, 75, 161, 2, 1, 141, 9, 8, 10, 23].

Aside from the real world applications of MIFE, the work of [29, 41] showed that sufficiently expressive MIFE can be used to construct the powerful notion of *Indistinguishability Obfuscation (iO)*, which seeks to garble circuits while preserving

their input-output behaviour. In more detail, given a circuit C , an obfuscation \tilde{C} preserves the correctness of C so that $C(\mathbf{x}) = \tilde{C}(\mathbf{x})$ for *every* input \mathbf{x} , but hides everything else about C . This security property can be formalized in various ways, and one popular formalization asks that an adversary, given the obfuscation of C_b where b is a random bit and C_0 is functionally *equivalent* to C_1 , cannot predict the value of b with non-negligible advantage.

While non-obvious in the beginning what such an object is useful for, a series of works has shown that iO can be used to instantiate a large number of advanced cryptographic primitives [90, 158, 67, 101, 118, 140, 39, 78, 92, 117, 125, 138]. Following this, there ensued a quest for constructing iO from reasonable cryptographic assumptions, with a large number of works coming closer and closer to the goal [91, 5, 20, 128, 95, 96, 171, 54, 53], until the beautiful work of Jain, Lin and Sahai [129] finally accomplished the goal.

Constructions from Lattice Assumptions. So far, the only constructions of compact FE/MIFE/ iO from standard assumptions rely on pairings (together with LPN variants and/or low depth PRGs) [129, 130, 154] and are hence quantum insecure. Perhaps even more importantly, it is dissatisfying to have only a single pathway to constructing such a central object in the theory of cryptography. Casting around for other assumptions, hopefully with (conjectured) quantum security – the most obvious candidate is some assumption based on lattices. In particular, the Learning With Errors (LWE) assumption has proven to be amazingly versatile and provides elegant solutions to encrypted computation primitives such as fully homomorphic encryption or attribute based encryption [63, 52, 55, 99, 107, 43]. Thus the hope basing iO on LWE is natural, and unsurprisingly, has received a lot of attention. Several exciting candidates for compact FE and iO have been constructed from strengthenings of the LWE assumption [5, 20, 171, 96, 82] – unfortunately, these either rely on some heuristic, or their

underlying assumptions have been broken [120, 126].

Our Approach. Since full fledged FE/*iO* from lattice based assumptions has been elusive despite significant effort, a principled approach for making progress would be to restrict the functionality supported by FE/*iO* so that it is still meaningful for applications while also admitting a construction from some reasonable lattice assumption. A promising candidate assumption that interpolates the safety of LWE and the insecurity of *iO* or multilinear map assumptions is the recently introduced *evasive* LWE [169, 163]. At a very high level, evasive LWE can be seen as a lattice analog of the generic group model, which is popular in the pairings world, in that it restricts the class of attacks that an adversary can mount. Below, we let \underline{X} denote a noisy version of X where the exact value of noise is not important and $\mathbf{B}^{-1}(\mathbf{P})$ denote a short preimage \mathbf{K} (say) such that $\mathbf{BK} = \mathbf{P} \bmod q$. The evasive LWE assumption roughly says that if

$$(\mathbf{A}, \mathbf{B}, \mathbf{P}, \underline{\mathbf{s}^\top \mathbf{A}}, \underline{\mathbf{s}^\top \mathbf{B}}, \underline{\mathbf{s}^\top \mathbf{P}}, \text{aux}) \approx_c (\mathbf{A}, \mathbf{B}, \mathbf{P}, \$, \$, \$, \text{aux})$$

where $\mathbf{A}, \mathbf{B}, \mathbf{P}$ are matrices of appropriate dimensions, \mathbf{s} is a secret vector, aux is some auxiliary information and $\$$ represents random, then

$$(\mathbf{A}, \mathbf{B}, \mathbf{P}, \underline{\mathbf{s}^\top \mathbf{A}}, \underline{\mathbf{s}^\top \mathbf{B}}, \mathbf{B}^{-1}(\mathbf{P}), \text{aux}) \approx_c (\mathbf{A}, \mathbf{B}, \mathbf{P}, \$, \$, \mathbf{B}^{-1}(\mathbf{P}), \text{aux})$$

Evasive LWE has proven to be a very meaningful strengthening of LWE in that it has provided several strong new applications that had been elusive from plain LWE, despite significant research effort over decades – optimal broadcast encryption [169], witness encryption [164], multi-input attribute based encryption [22], optimal broadcast and trace [11], attribute based encryption (ABE) for unbounded depth circuits [122] and ABE for Turing machines [13], to name a few. Evasive LWE is considered plausible in “safe” regimes (discussed in detail in Section 5.1.3). Continuing this agenda, we ask:

*Can we construct MIFE/*iO* for a nontrivial functionality from evasive LWE?*

6.1.1 Our Results

We answer the above question in the affirmative and construct the first MIFE and iO for pseudorandom functionalities, namely functions where the output is pseudorandom for every input *seen by the adversary*. In the context of iO this means that the entire truth table must be pseudorandom, since the adversary can compute the functionality for any input. But for MIFE, it suffices to only restrict the output for the functions queried by the adversary. We denote such an MIFE by mi-prFE.

In more detail, we obtain the following results:

Theorem 6.1 (mi-prFE for constant arity). *Assume evasive LWE and LWE. Then there exists a mi-prFE scheme for arity $n = O(1)$, supporting functions with (fixed) input length L and bounded polynomial depth.*¹

For polynomial arity, we require a strengthening of evasive LWE, which we call non-uniform κ -evasive LWE. This is necessary to handle some delicate technical issues, which also appear in the proof of witness encryption from evasive LWE by [164], to the best of our understanding. We believe that a similar strengthening of evasive LWE should also be required for their proof, for the same reasons as ours. We then show:

Theorem 6.2 (mi-prFE for polynomial arity). *Assume non-uniform κ -evasive LWE, non-uniform sub-exponential PRF, and non-uniform sub-exponential LWE. Then there exists a mi-prFE scheme for arity $n = \text{poly}(\lambda)$, supporting functions with bounded polynomial depth.*

We bootstrap our mi-prFE to obtain the first iO for pseudorandom functionalities, similar to [29, 41], albeit via a different proof of security. We denote this by prIO. In more detail:

Theorem 6.3 (prIO). *Assuming non-uniform κ -evasive LWE, non-uniform sub-exponential PRF, and non-uniform sub-exponential LWE, there exists a prIO*

¹As is shown in [19], evasive LWE does not hold for the general samplers. When we invoke the assumption, we always invoke it with respect to a specific sampler class induced by the respective application. However, for ease of exposition, we will sometimes omit the reference to the specific sampler in the overview. Similar remarks can be applied to prFE and mi-prFE.

scheme for all polynomial sized circuits.

We also define a variant of prIO which only supports polynomial sized domains – we denote this variant by pPRIO . The advantage of considering this restricted variant is that we can base its security on (plain) evasive LWE , rather than its strengthening.

Theorem 6.4. *Assuming evasive LWE and LWE , there exists a secure pPRIO scheme supporting circuits of bounded size.*

Applications.

We obtain several applications from our new tools of mi-prFE and prIO :

1. *Multi Input Predicate Encryption (miPE) for Constant Arity.* Assuming evasive LWE and LWE , there exists a mi-prFE scheme for arity $n = O(1)$, supporting functions of bounded polynomial depth. The only prior work to support miPE with constant arity for function class \mathbf{P} is by Agrawal et al. [22] and uses a strengthening of (non-standard) tensor LWE [168] together with evasive LWE .
2. *Multi Input Predicate Encryption for Polynomial Arity.* Assuming non-uniform κ -evasive LWE , non-uniform sub-exponential PRF , and non-uniform sub-exponential LWE , there exists a miPE scheme for arity $n = \text{poly}(\lambda)$, supporting functions of bounded polynomial depth. MIPE for polynomial arity supporting \mathbf{P} was not known before, to the best of our knowledge.
3. *Two Party ID Based Key Exchange.* Assuming non-uniform κ -evasive LWE , non-uniform sub-exponential PRF , non-uniform sub-exponential LWE , and the DBDH assumption, there exists a secure two-party ID-NIKE scheme. This leads to the first ID-NIKE in the standard model without using multilinear maps or indistinguishability obfuscation.
4. *Instantiating the Random Oracle.* Hohenberger, Sahai and Waters [119] used iO to instantiate the random oracle in several applications. In more detail, they showed selective security of the full-domain hash (FDH) signature based on trapdoor permutations (TDP) [37], the adaptive security of RSA FDH signatures [79], the selective security of BLS signatures, and the adaptive security of BLS signatures [45] in the standard model. Our prIO can be used to instantiate all these applications.

Additional Prior Work. The notion of iO for pseudorandom functionalities was considered implicitly by the work of Mathialagan, Peters and Vaikuntanathan [146] where they used subexponential LWE and evasive LWE to construct adaptively sound

zero-knowledge SNARKs for UP. In a previous version of their work [147], which was in private circulation and shared with us, this notion was defined explicitly and leveraged to obtain unlevelled fully homomorphic encryption. However, an explicit construction of iO for pseudorandom functionalities was not provided.

6.1.2 Technical Overview

Compact FE for Pseudorandom Functionalities. Our starting point is our companion work which provides FE for pseudorandom functionalities in the single input setting [12]. We recall the syntax here: The setup algorithm takes as input the security parameter λ and parameter prm , specifying the parameters of the function class, and outputs (mpk, msk) . The key generation algorithm on input msk and a function $f : \{0, 1\}^L \rightarrow \{0, 1\}^\ell$ outputs a functional secret key sk_f . The encryption algorithm on input mpk and an input message $\mathbf{x} \in \{0, 1\}^L$ outputs a ciphertext ct . The decryption algorithm takes the functional secret key sk_f and ciphertext ct as input and outputs some $\mathbf{y} \in \{0, 1\}^\ell$. Correctness requires that decryption should output $f(\mathbf{x})$ if the setup, encrypt and key generation algorithms were run honestly.

The definition of security, termed **prCT** security, says that so long as the *output* of the functionality is pseudorandom, the *ciphertext* is pseudorandom. In more detail, let Samp be a PPT algorithm that on input 1^λ , outputs

$$(f_1, \dots, f_{Q_{\text{key}}}, x_1, \dots, x_{Q_{\text{msg}}}, \mathbf{aux} \in \{0, 1\}^*)$$

where Q_{key} is the number of key queries, Q_{msg} is the number of message queries.

We define the following advantage functions:

$$\begin{aligned} \text{Adv}_{\mathcal{A}_0}^{\text{PRE}}(\lambda) &\stackrel{\text{def}}{=} \Pr[\mathcal{A}_0(\mathbf{aux}, \{f_i, f_i(x_j)\}_{i \in [Q_{\text{key}}], j \in [Q_{\text{msg}}]}) = 1] \\ &\quad - \Pr[\mathcal{A}_0(\mathbf{aux}, \{f_i, \Delta_{i,j} \leftarrow \mathcal{Y}_{\text{prm}}\}_{i \in [Q_{\text{key}}], j \in [Q_{\text{msg}}]}) = 1] \end{aligned} \tag{6.1}$$

$$\begin{aligned} \text{Adv}_{\mathcal{A}_1}^{\text{POST}}(\lambda) \stackrel{\text{def}}{=} & \Pr[\mathcal{A}_1(\text{mpk}, \text{aux}, \{f_i, \text{ct}_j \leftarrow \text{Enc}(\text{mpk}, x_j), \text{sk}_{f_i}\}_{i \in [Q_{\text{key}}], j \in [Q_{\text{msg}}]}) = 1] \\ & - \Pr[\mathcal{A}_1(\text{mpk}, \text{aux}, \{f_i, \delta_j \leftarrow \mathcal{CT}, \text{sk}_{f_i}\}_{i \in [Q_{\text{key}}], j \in [Q_{\text{msg}}]}) = 1] \end{aligned} \quad (6.2)$$

where $(f_1, \dots, f_{Q_{\text{key}}}, x_1, \dots, x_{Q_{\text{msg}}}, \text{aux} \in \{0, 1\}^*) \leftarrow \text{Samp}(1^\lambda)$, $(\text{mpk}, \text{msk}) \leftarrow \text{Setup}(1^\lambda, \text{prm})$ and \mathcal{CT} is the ciphertext space. We say that a prFE scheme for function family \mathcal{F}_{prm} satisfies prCT security if for every PPT Samp and \mathcal{A}_1 , there exists another PPT \mathcal{A}_0 such that

$$\text{Adv}_{\mathcal{A}_0}^{\text{PRE}}(\lambda) \geq \text{Adv}_{\mathcal{A}_1}^{\text{POST}}(\lambda)/Q(\lambda) - \text{negl}(\lambda) \quad (6.3)$$

and $\text{Time}(\mathcal{A}_0) \leq \text{Time}(\mathcal{A}_1) \cdot Q(\lambda)$ for some polynomial $Q(\cdot)$. The work of [12] obtains the following result:

Theorem 6.5. *Assuming the LWE and evasive LWE assumptions, there exists a secure prFE scheme for function class $\mathcal{F}_{L(\lambda), \ell(\lambda), \text{dep}(\lambda)} = \{f : \{0, 1\}^L \rightarrow \{0, 1\}^\ell\}$ satisfying $|\text{mpk}| = L \cdot \text{poly}(\text{dep}, \lambda)$, $|\text{sk}_f| = \ell \cdot \text{poly}(\text{dep}, \lambda)$, $|\text{ct}| = L \cdot \text{poly}(\text{dep}, \lambda)$ where $\text{dep} = \text{poly}(\lambda)$ is the depth bound on the functions supported by the scheme.*

We remark that their result is actually significantly more general – they construct the broader notion of *partially hiding* FE, where the ciphertext can additionally have a public attribute in addition to a private payload², with optimal parameters and can support circuits of unbounded depth. But bounded depth prFE is sufficient for the applications considered in the present work, so we restrict our attention to this.

Extending to the Multi-Input Setting. We extend the notion of prFE to multi-input setting and define a *secret-key* multi-input FE for pseudorandom functionalities $\text{mi-prFE} = (\text{Setup}, \text{KeyGen}, \text{Enc}_1, \dots, \text{Enc}_n, \text{Dec})$ for n -ary functions as follows: The setup algorithm on input 1^λ , arity 1^n and parameter prm , specifying the parameters

²The astute reader may wonder why regular FE does not imply partially hiding FE (PHFE) just by outputting the public attribute as part of the ciphertext – such an approach would make the ciphertext size grow with the length of public attribute which can be avoided in PHFE.

of the function class, outputs (mpk, msk) . The key generation algorithm on input msk and a function $f : (\mathcal{X}_{\text{prm}})^n \rightarrow \mathcal{Y}_{\text{prm}}$ outputs a functional secret key sk_f . The i -th encryption algorithm on input msk and an input message $x_i \in \mathcal{X}_{\text{prm}}$ outputs a ciphertext ct_i . The decryption algorithm on input secret key sk_f and n ciphertexts $\text{ct}_1, \dots, \text{ct}_n$ (corresponding to inputs x_1, \dots, x_n respectively) outputs some $y \in \mathcal{Y}_{\text{prm}}$. The security has a similar flavor to the single-input setting, where we require that the ct is pseudorandom given the output of the function of encrypted input is pseudorandom—however it requires much care to accommodate the fact that there are exponentially many function evaluations even if only polynomially many input queries per slot are issued. We define the security as follows.

Let $\kappa = \kappa(\lambda)$ be a function in λ and Samp be a PPT algorithm that on input 1^λ , outputs

$$\left(\{f_k\}_{k \in [q_0]}, \{x_1^{j_1}\}_{j_1 \in [q_1]}, \dots, \{x_n^{j_n}\}_{j_n \in [q_n]}, \text{aux} \in \{0, 1\}^* \right)$$

where q_0 is the number of key queries and q_i is the number of encryption queries for the i -th slot. We define the following advantage functions:

$$\begin{aligned} \text{Adv}_{\mathcal{A}_0}^{\text{PRE}}(\lambda) &\stackrel{\text{def}}{=} \Pr[\mathcal{A}_0(1^\kappa, f_1, \dots, f_{q_0}, \{f_k(x_1^{j_1}, \dots, x_n^{j_n})\}_{k \in [q_0], j_1 \in [q_1], \dots, j_n \in [q_n]}, \text{aux}) = 1] \\ &\quad - \Pr[\mathcal{A}_0(1^\kappa, f_1, \dots, f_{q_0}, \{\Delta_{k, j_1, \dots, j_n} \leftarrow \mathcal{Y}_{\text{prm}}\}_{k \in [q_0], j_1 \in [q_1], \dots, j_n \in [q_n]}, \text{aux}) = 1] \\ \text{Adv}_{\mathcal{A}_1}^{\text{POST}}(\lambda) &\stackrel{\text{def}}{=} \Pr[\mathcal{A}_1(\text{mpk}, f_1, \dots, f_{q_0}, \{\text{Enc}_i(\text{msk}, x_i^{j_i})\}_{i \in [n], j_i \in [q_i]}, \text{sk}_{f_1}, \dots, \text{sk}_{f_{q_0}}, \text{aux}) = 1] \\ &\quad - \Pr[\mathcal{A}_1(\text{mpk}, f_1, \dots, f_{q_0}, \{\delta_i^{j_i} \leftarrow \text{Sim}(\text{msk})\}_{i \in [n], j_i \in [q_i]}, \text{sk}_{f_1}, \dots, \text{sk}_{f_{q_0}}, \text{aux}) = 1] \end{aligned}$$

where $(\text{mpk}, \text{msk}) \leftarrow \text{Setup}(1^\lambda, 1^n, \text{prm})$ and $\text{sk}_{f_k} \leftarrow \text{KeyGen}(\text{msk}, f_k)$ for $k \in [q_0]$.

We say that a mi-prFE scheme is secure if for every PPT Samp , \mathcal{A}_1 , and Sim there exists another PPT \mathcal{A}_0 and a polynomial $p(\cdot)$ such that

$$\text{Adv}_{\mathcal{A}_0}^{\text{PRE}}(\lambda) \geq \text{Adv}_{\mathcal{A}_1}^{\text{POST}}(\lambda)/p(\kappa) - \text{negl}(\kappa), \quad \text{Time}(\mathcal{A}_0) \leq p(\kappa) \cdot \text{Time}(\mathcal{A}_1).$$

The parameter κ above is introduced to adjust the strength of the requirement for the

precondition. By default, we require $\kappa \geq \lambda^n$ since the input length to the distinguisher is polynomial in λ^n anyway and this condition should be fulfilled for the above equations to make sense. If we need κ to be larger, this strengthens the requirement for the precondition, as it means we want the distributions in the pre-condition to be indistinguishable against an adversary with a longer running time.³ Ideally, we want κ to be as small as λ^n to make the requirement weaker. Looking ahead, for our construction of mi-prFE scheme supporting polynomial arity $n = \text{poly}(\lambda)$, we require large κ as an artifact of the security proof techniques. In the special case of n being constant, we can achieve $\kappa = \lambda^n$.

Construction. Next, we describe our construction for bounded depth mi-prFE using a bounded depth prFE and a secret-key encryption scheme. Our construction adapts the key idea from [29], of "unrolling" ciphertexts on the fly via recursive decryption. However, since our security notion is quite different, our proof departs significantly from theirs, as we will discuss below.

Specifically, we use n instances of a single-input prFE scheme $\{\text{prFE}_i\}_{i \in [n]}$, with appropriate input lengths, to build a n arity mi-prFE scheme where the i -th encryption algorithm Enc_i outputs the prFE_{i+1} functional secret-key, $\text{prFE}_{i+1}.\text{sk}$, for $i \in [n-1]$ and Enc_n outputs a ciphertext corresponding to prFE_n scheme. Here the $\text{prFE}_{i+1}.\text{sk}$ contains the input \mathbf{x}_i hardcoded within itself, wrapped in an SKE scheme, since prFE does not support function hiding. It computes the ciphertext $\text{prFE}_i.\text{ct}$ for the input $(\text{SKE}.\text{sk}, \mathbf{x}_i, \mathbf{x}_{i+1}, \dots, \mathbf{x}_{n-1}, \mathbf{x}_n)$ decryptable by $\text{prFE}_i.\text{sk}$ which in turn computes the ciphertext $\text{prFE}_{i-1}.\text{ct}$ decryptable by $\text{prFE}_{i-1}.\text{sk}$ and so on. Now, note that the decryption of slot n ciphertext with slot $n-1$ functional secret-key will give us a ciphertext decryptable by functional key at slot $n-2$. Unrolling upto slot 1, we get a ciphertext, $\text{prFE}_1.\text{ct}$, corresponding to prFE_1 scheme. Finally, the key generation algorithm outputs a functional secret-key for prFE_1 which together with $\text{prFE}_1.\text{ct}$ will give us the desired

³Recall that \mathcal{A}_1 is a PPT algorithm. This means that it runs in polynomial time in its input length. Here, κ serves as a "padding", which artificially makes the input longer and allows \mathcal{A}_1 to run in longer time.

output.

In more detail⁴,

1. The setup algorithm generates n instances of prFE scheme $\{\text{prFE}_i.\text{mpk}, \text{prFE}_i.\text{msk}\}$ for appropriate input lengths and a secret key corresponding to SKE scheme SKE.sk . It outputs $\text{mpk} = (\{\text{prFE}_i.\text{mpk}\}_{i \in [n]})$ and $\text{msk} = (\text{SKE.sk}, \{\text{prFE}_i.\text{msk}, \text{prFE}_i.\text{mpk}\}_{i \in [n]})$.
2. The key generation algorithm on input msk and a function $f : (\{0, 1\}^L)^n \rightarrow \{0, 1\}$, computes $\text{prFE}_1.\text{sk}_f \leftarrow \text{prFE}_1.\text{KeyGen}(\text{prFE}_1.\text{msk}, f)$ and outputs $\text{sk}_f := \text{prFE}_1.\text{sk}_f$.
3. The i -th encryption algorithm on input $(\text{msk}, \mathbf{x}_i \in \{0, 1\}^L)$, parses $\text{msk} = (\text{SKE.sk}, \{\text{prFE}_i.\text{msk}, \text{prFE}_i.\text{mpk}\}_{i \in [n]})$ and does as follows. If $i \in [n - 1]$
 - Compute $\text{SKE.ct}_i \leftarrow \text{SKE.Enc}(\text{SKE.sk}, \mathbf{x}_i)$.
 - Define function $F_i := F_i[\text{SKE.ct}_i, \text{prFE}_i.\text{mpk}]$ which on input $(\text{SKE.sk}, \mathbf{x}_{i+1}, \dots, \mathbf{x}_{n-1}, \mathbf{x}_n)$ first computes $\mathbf{x}_i = \text{SKE.Dec}(\text{SKE.sk}, \text{SKE.ct}_i)$ and then computes a $\text{prFE}_i.\text{ct}$ encoding $(\text{SKE.sk}, \mathbf{x}_i, \mathbf{x}_{i+1}, \dots, \mathbf{x}_{n-1}, \mathbf{x}_n)$ if $i \neq 1$ else it computes $\text{prFE}_1.\text{ct}$ encoding $(\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_{n-1}, \mathbf{x}_n)$ ⁵. It outputs $\text{prFE}_i.\text{ct}$.
 - It computes a functional key for F_i using the $i + 1$ -th instance of prFE and outputs it as the i -th ciphertext, i.e., $\text{ct}_i := \text{prFE}_{i+1}.\text{sk} \leftarrow \text{prFE}_{i+1}.\text{KeyGen}(\text{prFE}_{i+1}.\text{msk}, F_i)$.

If $i = n$, it outputs $\text{ct}_n := \text{prFE}_n.\text{ct} \leftarrow \text{prFE}_n.\text{Enc}(\text{prFE}_n.\text{mpk}, (\text{SKE.sk}, \mathbf{x}_n))$.
4. The decryption algorithm on input $\text{sk}_f = \text{prFE}_1.\text{sk}_f$, and ciphertexts $\text{ct}_i = \text{prFE}_{i+1}.\text{sk}$ for $i \in [n - 1]$, and $\text{ct}_n = \text{prFE}_n.\text{ct}$ does the following: (a) Iteratively compute $\text{prFE}_{i-1}.\text{ct}$ for $i \in [2, n]$ by decrypting $\text{prFE}_i.\text{ct}$ with $\text{prFE}_i.\text{sk}$ starting with $i = n$. (b) Compute and output $\mathbf{y} \leftarrow \text{prFE}_1.\text{Dec}(\text{prFE}_1.\text{mpk}, \text{prFE}_1.\text{sk}_f, f, \text{prFE}_1.\text{ct})$.

Correctness follows from the correctness of underlying ingredients. To see this, note that by the correctness of prFE_n and the definition of F_{n-1} , we have $\text{prFE}_n.\text{Dec}(\text{prFE}_n.\text{mpk}, \text{prFE}_n.\text{sk}, F_{n-1}, \text{prFE}_n.\text{ct})$ will output $F_{n-1}(\text{SKE.sk}, \mathbf{x}_n)$ correctly. Next, by the correctness of the SKE scheme, we have $\mathbf{x}_{n-1} = \text{SKE.Dec}(\text{SKE.sk}, \text{SKE}, \text{ct}_{n-1})$ thus $F_{n-1}(\text{SKE.sk}, \mathbf{x}_n) = \text{prFE}_{n-1}.\text{ct}$ where

⁴We omit substantial notation here for the ease of readability.

⁵The randomness for computing the ciphertexts comes from a PRF, which we omit here in the overview.

$\text{prFE}_{n-1}.\text{ct}$ encodes

$(\text{SKE.sk}, \mathbf{x}_{n-1}, \mathbf{x}_n)$. Unrolling as in decryption step (a), we get $\text{prFE}_1.\text{ct}$ which encodes $(\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_{n-1}, \mathbf{x}_n)$. Now, from step (b) of decryption and correctness of prFE_1 scheme we have $\text{prFE}_1.\text{Dec}(\text{prFE}_1.\text{mpk}, \text{prFE}_1.\text{sk}_f, f, \text{prFE}_1.\text{ct}) = f(\mathbf{x}_1, \dots, \mathbf{x}_n)$.

Security. While the key-idea of our construction is adapted from [29], our security proof differs significantly from theirs as we elaborate next. Consider the initial view of an adversary \mathcal{A} which outputs q_0 key queries $\{f_1, \dots, f_{q_0}\}$, q_i input queries for i -th slot $\{x_1^{j_1}\}_{j_1 \in [q_1]}, \dots, \{x_n^{j_n}\}_{j_n \in [q_n]}$ and auxiliary information $\text{aux}_{\mathcal{A}}$

$$\mathcal{D}_{0,0} : \left(\begin{array}{l} \text{aux}_{\mathcal{A}}, \{\text{prFE}_i.\text{mpk}\}_{i \in [n]}, \{f_k, \text{sk}_{f_k} = \text{prFE}_1.\text{sk}_{f_k}\}_{k \in [q_0]} \\ \left\{ \text{ct}_i^{j_i} = \text{SKE.ct}_i^{j_i}, \text{prFE}_{i+1}.\text{sk}^{j_i} \right\}_{\substack{i \in [n-1], \\ j_i \in [q_i]}}, \\ \left\{ \text{ct}_n^{j_n} = \text{prFE}_n.\text{ct}^{j_n} \right\}_{j_n \in [q_n]} \end{array} \right)$$

To prove security we design a simulator Sim as follows: On input $\text{msk} = (\text{SKE.sk}, \{\text{prFE}_i.\text{msk}, \text{prFE}_i.\text{mpk}\}_{i \in [n]})$

- for $i \in [n-1]$, it first samples a random SKE ciphertext γ_i from the ciphertext space of SKE scheme, defines $F_i[\gamma_i, \text{prFE}_i.\text{mpk}]$ as in the construction and outputs $\text{ct}_i = \text{prFE}_{i+1}.\text{sk} \leftarrow \text{prFE}_{i+1}.\text{KeyGen}(\text{prFE}_{i+1}.\text{msk}, F_i[\gamma_i, \text{prFE}_i.\text{mpk}])$.
- for $i = n$, it outputs a randomly sampled ct_n from the ciphertext space of prFE_n scheme.

Given the above simulator it suffices to show that Section 6.1.2 is indistinguishable from the following distribution

$$\mathcal{D}_{0,1} : \left(\begin{array}{l} \text{aux}_{\mathcal{A}}, \{\text{prFE}_i.\text{mpk}\}_{i \in [n]}, \{f_k, \text{sk}_{f_k} = \text{prFE}_1.\text{sk}_{f_k}\}_{k \in [q_0]} \\ \left\{ \text{ct}_i^{j_i} = \gamma_i^{j_i} \leftarrow \mathcal{CT}_{\text{SKE}}, \text{prFE}_{i+1}.\text{sk}^{j_i} \right\}_{\substack{i \in [n-1], \\ j_i \in [q_i]}}, \\ \left\{ \delta^{j_n} \leftarrow \mathcal{CT}_{\text{prFE}_n} \right\}_{j_n \in [q_n]} \end{array} \right)$$

At a high level, the security proof proceeds as follows. To prove the pseudorandomness of $\{\text{prFE}_n.\text{ct}^{j_n}\}_{j_n}$, we show that the decryption results of these ciphertexts using the secret keys $\{\text{prFE}_n.\text{sk}^{j_{n-1}}\}_{j_{n-1}}$ are all pseudorandom. This allows us to invoke the security of

prFE_n . The decryption results of the above ciphertexts using the secret keys are ciphertexts $\{\text{prFE}_{n-1}.\text{ct}^{j_{n-1}, j_n}\}_{j_{n-1}, j_n}$ of prFE_{n-1} and we would like to prove the pseudorandomness of them. We again consider the decryption results of the ciphertexts using the secret keys $\{\text{prFE}_{n-1}.\text{sk}^{j_{n-2}}\}_{j_{n-2}}$. This process continues until we reach the point where we have to prove the pseudorandomness of $\{\text{prFE}_1.\text{ct}^{j_1, \dots, j_n}\}_{j_1, \dots, j_n}$, where each ciphertext encodes $(\mathbf{x}_1^{j_1}, \dots, \mathbf{x}_n^{j_n})$. By invoking the security of prFE_1 once again, we can conclude that it suffices to show that $\{f_k(\mathbf{x}_1^{j_1}, \dots, \mathbf{x}_n^{j_n})\}_{k, j_1, \dots, j_n}$ are pseudorandom even given SKE ciphertexts encrypting each $\mathbf{x}_i^{j_i}$, where the latter is dealt as auxiliary information throughout the process of the above recursive invocations of prFE security. We then invoke the security of SKE to erase the information of $\mathbf{x}_i^{j_i}$ from the SKE ciphertexts. This allows us to conclude, since the pseudorandomness of $\{f_k(\mathbf{x}_1^{j_1}, \dots, \mathbf{x}_n^{j_n})\}_{k, j_1, \dots, j_n}$ directly follows from the precondition.

A bit more formally, to prove Section 6.1.2 \approx Section 6.1.2 we begin by invoking the security of prFE_n with sampler that provides inputs $\{(\text{SKE.sk}, \mathbf{x}_n^{j_n})\}_{j_n \in [q_n]}$, functions $\{F_{n-1}^{j_{n-1}}[\text{SKE.ct}_{n-1}^{j_{n-1}}, \text{prFE}_{n-1}.\text{mpk}]\}_{j_{n-1} \in [q_{n-1}]}$, and all the remaining components of Section 6.1.2 as auxiliary information. Now, from the security guarantee of prFE_n with the sampler, we know that to prove $\text{prFE}.\text{ct}_n^{j_n}$ is pseudorandom it suffices to show function output

$$\begin{aligned} & F_{n-1}^{j_{n-1}}[\text{SKE.ct}_{n-1}^{j_{n-1}}, \text{prFE}_{n-1}.\text{mpk}](\text{SKE.sk}, \mathbf{x}_n^{j_n}) \\ &= \text{prFE}_{n-1}.\text{Enc}(\text{prFE}_{n-1}.\text{mpk}, (\text{SKE.sk}, \mathbf{x}_{n-1}^{j_{n-1}}, \mathbf{x}_n^{j_n})) = \text{prFE}_{n-1}.\text{ct}^{j_{n-1}, j_n} \end{aligned}$$

is pseudorandom for all $j_{n-1} \in [q_{n-1}]$, $j_n \in [q_n]$. Thus it suffices to show

$$\begin{aligned} \mathcal{D}_{1,0} : & \left(\begin{array}{c} \text{aux}_{\mathcal{A}}, \{\text{prFE}_i.\text{mpk}\}_{i \in [n-1]}, \{f_k, \text{sk}_{f_k} = \text{prFE}_1.\text{sk}_{f_k}\}_{k \in [q_0]}, \\ \left\{ \text{SKE.ct}_i^{j_i} \right\}_{\substack{i \in [n-1], \\ j_i \in [q_i]}}, \left\{ \text{prFE}_{i+1}.\text{sk}^{j_i} \right\}_{\substack{i \in [n-2], \\ j_i \in [q_i]}}, \left\{ \text{prFE}_{n-1}.\text{ct}^{j_{n-1}, j_n} \right\}_{\substack{j_{n-1} \in [q_{n-1}], \\ j_n \in [q_n]}} \end{array} \right) \\ & \approx \end{aligned}$$

$$\mathcal{D}_{1,1} : \left(\text{aux}_{\mathcal{A}}, \{\text{prFE}_i.\text{mpk}\}_{i \in [n-1]}, \{f_k, \text{sk}_{f_k} = \text{prFE}_1.\text{sk}_{f_k}\}_{k \in [q_0]}, \right. \\ \left. \left\{ \gamma_i^{j_i} \right\}_{\substack{i \in [n-1], \\ j_i \in [q_i]}}, \left\{ \text{prFE}_{i+1}.\text{sk}^{j_i} \right\}_{\substack{i \in [n-2], \\ j_i \in [q_i]}}, \left\{ \delta^{j_{n-1}, j_n} \right\}_{\substack{j_{n-1} \in [q_{n-1}], \\ j_n \in [q_n]}} \right)$$

where $\gamma_i^{j_i} \leftarrow \mathcal{CT}_{\text{SKE}}$, $\delta^{j_{n-1}, j_n} \leftarrow \mathcal{CT}_{\text{prFE}_{n-1}}$, and $\mathcal{CT}_{\text{SKE}}$ and $\mathcal{CT}_{\text{prFE}_{n-1}}$ denotes the ciphertext space of the SKE scheme and prFE_{n-1} scheme, respectively. Recursively invoking the security of prFE_i for $i = n-1, \dots, 2$, it suffices to show the following:

$$\mathcal{D}_{n-1,0} : \left(1^\kappa, \text{aux}_{\mathcal{A}}, \text{prFE}_1.\text{mpk}, \{f_k, \text{sk}_{f_k} = \text{prFE}_1.\text{sk}_{f_k}\}_{k \in [q_0]}, \right. \\ \left. \left\{ \text{SKE.ct}_i^{j_i} \right\}_{\substack{i \in [n-1], \\ j_i \in [q_i]}}, \left\{ \text{prFE}_1.\text{ct}^{j_1, \dots, j_n} \right\}_{j_1 \in [q_1], \dots, j_n \in [q_n]} \right) \\ \approx \\ \mathcal{D}_{n-1,1} : \left(1^\kappa, \text{aux}_{\mathcal{A}}, \text{prFE}_1.\text{mpk}, \{f_k, \text{sk}_{f_k} = \text{prFE}_1.\text{sk}_{f_k}\}_{k \in [q_0]}, \right. \\ \left. \left\{ \gamma_i^{j_i} \leftarrow \mathcal{CT}_{\text{SKE}} \right\}_{\substack{i \in [n-1], \\ j_i \in [q_i]}}, \left\{ \delta^{j_1, \dots, j_n} \leftarrow \mathcal{CT}_{\text{prFE}_1} \right\}_{j_1 \in [q_1], \dots, j_n \in [q_n]} \right).$$

Finally, applying the security of prFE_1 once again, we can see that it suffices to show the following:

$$\mathcal{D}_{n,0} : \left(1^\kappa, \text{aux}_{\mathcal{A}}, \left\{ f_k, f_k(\mathbf{x}_1^{j_1}, \dots, \mathbf{x}_n^{j_n}) \right\}_{k, j_1, \dots, j_n}, \left\{ \text{SKE.ct}_i^{j_i} \leftarrow \text{SKE.Enc}(\text{SKE.sk}, \mathbf{x}_i^{j_i}) \right\}_{i, j_i} \right) \\ \approx_c \\ \mathcal{D}_{n,1} : \left(1^\kappa, \text{aux}_{\mathcal{A}}, \left\{ f_k, \Delta_k^{j_1, \dots, j_n} \leftarrow \{0, 1\} \right\}_{k, j_1, \dots, j_n}, \left\{ \gamma_i^{j_i} \leftarrow \mathcal{CT}_{\text{SKE}} \right\}_{i, j_i} \right) \quad (6.4)$$

Here, 1^κ appearing in the above distributions is introduced for compensating the blow up of the size of the adversary caused by the multiple invocations of the prFE security. The detail is not important here and we refer to Section 6.3 for the detail.

To prove Equation (6.4), we first invoke the security of SKE scheme to show the pseudorandomness of SKE ciphertexts. This erases the information of $\mathbf{x}_i^{j_i}$ from $\text{SKE.ct}_i^{j_i}$. Next, we use the fact that the functionality supported by our scheme is pseudorandom to argue that the function values $f_k(\mathbf{x}_1^{j_1}, \dots, \mathbf{x}_n^{j_n})$ are pseudorandom.

Subtleties in the proof. The high level overview presented above hides many important details, and indeed, as stated is not secure. There are two important subtleties that arise when making the formal argument, and these are so significant that they require us to strengthen the underlying evasive LWE assumption. Moreover, to the best of our understanding, these issues also arise in prior work [164] and fixing them there also requires to strengthen the evasive LWE assumption as we do below.

We note that at a high level, the overall structure of our security proof above is similar to that of witness encryption in [164]. In both proofs, the main step considers parameterized distributions $\{\mathcal{D}_{h,b}\}_{h \in [n,0], b \in \{0,1\}}$ and shows that $\mathcal{D}_{0,0}$ and $\mathcal{D}_{0,1}$ are indistinguishable if $\mathcal{D}_{n,0} \approx_c \mathcal{D}_{n,1}$ even with subexponentially small advantage against subexponential time adversary. To show this claim, [164] uses the evasive LWE assumption, while we use the security of the prFE instances, which in turn is reduced to the evasive LWE assumption. While this difference stems simply from the fact that we introduce the intermediate primitive of prFE to construct mi-prFE instead of directly constructing it from evasive LWE, there are more fundamental differences as well. In particular, we identify certain subtle issues in the proof by [164] and fix these by strengthening the assumptions. We elaborate on this below. To begin, we formally define the evasive LWE assumption.

Evasive LWE. Let **Samp** be a PPT algorithm that outputs $(\mathbf{S}, \mathbf{P}, \text{aux})$ on input 1^λ . For PPT adversaries \mathcal{A}_0 and \mathcal{A}_1 , we define the following advantage functions:

$$\mathcal{A}_{\text{Adv}_0}^{\text{PRE}}(\lambda) \stackrel{\text{def}}{=} \Pr[\text{Adv}_0(\mathbf{B}, \mathbf{S}\mathbf{B} + \mathbf{E}, \mathbf{S}\mathbf{P} + \mathbf{E}', \text{aux}) = 1] - \Pr[\text{Adv}_0(\mathbf{B}, \mathbf{C}_0, \mathbf{C}', \text{aux}) = 1] \quad (6.5)$$

$$\mathcal{A}_{\text{Adv}_1}^{\text{POST}}(\lambda) \stackrel{\text{def}}{=} \Pr[\text{Adv}_1(\mathbf{B}, \mathbf{S}\mathbf{B} + \mathbf{E}, \mathbf{K}, \text{aux}) = 1] - \Pr[\text{Adv}_1(\mathbf{B}, \mathbf{C}_0, \mathbf{K}, \text{aux}) = 1] \quad (6.6)$$

where $(\mathbf{S}, \mathbf{P}, \text{aux}) \leftarrow \text{Samp}(1^\lambda)$, $\mathbf{B}, \mathbf{C}_0, \mathbf{C}'$ are uniform matrices of appropriate dimensions, \mathbf{E}, \mathbf{E}' are low norm Gaussian error matrices, and $\mathbf{K} \leftarrow \mathbf{B}^{-1}(\mathbf{P})$.

We say that the *evasive* LWE (EvLWE) assumption holds if for every PPT Samp and Adv_1 , there exists another PPT Adv_0 and a polynomial $Q(\cdot)$ such that

$$\mathcal{A}_{\text{Adv}_0}^{\text{PRE}}(\lambda) \geq \mathcal{A}_{\text{Adv}_1}^{\text{POST}}(\lambda)/Q(\lambda) - \text{negl}(\lambda) \text{ and } \text{Time}(\mathcal{A}_0) \leq Q(\lambda) \cdot \text{Time}(\mathcal{A}_1). \quad (6.7)$$

Issue 1: On the multiplicative invocation of evasive LWE. To prove $\mathcal{D}_{0,0} \approx_c \mathcal{D}_{0,1}$, [164] assumes that there exists an adversary \mathcal{A}_0 that distinguishes them with non-negligible advantage ϵ and polynomial time t for the sake of contradiction. They then invoke the evasive LWE assumption with respect to an appropriately defined sampler Samp_1 to conclude that there exists a distinguishing adversary \mathcal{A}_1 against $\mathcal{D}_{1,0}$ and $\mathcal{D}_{1,1}$. This process continues multiple times, where they invoke evasive LWE with respect to the security parameter $\lambda_j := 2^j \lambda$ and an adversary \mathcal{A}_j for the j -th invocation to obtain another adversary \mathcal{A}_{j+1} , where \mathcal{A}_k is a distinguisher against $\mathcal{D}_{k,0}$ and $\mathcal{D}_{k,1}$. Denoting the distinguishing advantage against $\mathcal{D}_{j,0}$ and $\mathcal{D}_{j,1}$ of \mathcal{A}_j by ϵ_j , we have $\epsilon_{j+1} \geq \epsilon_j / \text{poly}_j(\lambda_j)$, where poly_j is a polynomial that is determined by the sampler Samp_j . Finally, they obtain a distinguisher \mathcal{A}_n against $\mathcal{D}_{n,0}$ and $\mathcal{D}_{n,1}$, where $\epsilon_n = \epsilon / \text{poly}_1(\lambda_1) \text{poly}_2(\lambda_2) \cdots \text{poly}_n(\lambda_n)$ and the running time being $\text{poly}_1(\lambda_1) \text{poly}_2(\lambda_2) \cdots \text{poly}_n(\lambda_n)$. They derive the conclusion by saying

$$\text{poly}_1(\lambda_1) \text{poly}_2(\lambda_2) \cdots \text{poly}_n(\lambda_n) = \text{poly}(\lambda_1 \cdots \lambda_n) = \text{poly}(2^{n^2} \lambda^n) \quad (6.8)$$

and setting the parameters so that there is no adversary of this running time and distinguishing advantage against $\mathcal{D}_{n,0}$ and $\mathcal{D}_{n,1}$. However, a subtlety is that $\text{poly}_1(\lambda_1) \text{poly}_2(\lambda_2) \cdots \text{poly}_n(\lambda_n) = \text{poly}(\lambda_1 \cdots \lambda_n)$ is not necessarily true. For example, one can consider the setting where we have $\text{poly}_j(\lambda) = \lambda^{2^j}$. This example may look a bit artificial, but it does not contradict the evasive LWE assumption, since j is treated as a constant asymptotically. In words, the issue arises from the fact that even if each

polynomial has a constant exponent, the maximum of the exponents can be arbitrarily large function in λ , when we consider non-constant number of polynomials. In this setting, \mathcal{A}_n 's distinguishing advantage is too small to derive the contradiction.

The above issue occurs due to the invocations of evasive LWE super-constant times. To resolve the problem, we consider non-uniform sampler $\{\text{Samp}_{h^*}\}_{h^*}$ that hardwires the "best" index h^* and invoke the evasive LWE⁶ only with respect to this sampler. In more detail, to argue that the final distributions $\mathcal{D}_{n,0}$ and $\mathcal{D}_{n,1}$ are indistinguishable, it is required that *all* pairs of distributions $\mathcal{D}_{j,0}$ and $\mathcal{D}_{j,1}$ be indistinguishable. Then, to arrive at a contradiction, it suffices to find even one intermediate pair (for index h^*) which is distinguishable – we invoke evasive LWE with respect to that. We avoid the problem of incurring security loss of polynomial with arbitrarily large exponent, since we invoke the evasive LWE only once w.r.t a single sampler in the proof. However, this solution entails the strengthening of the assumption where we consider non-uniform samplers. We believe that the same strengthening of the assumption is required for the proof in [164] as well.

Issue 2: On the additive term of evasive LWE. We believe there is another subtle issue that arises in the proof by [164]. To focus on the issue, we ignore the first issue discussed above and assume $\text{poly}_j(\lambda) = \lambda^c$ holds for some fixed $c \in \mathbb{N}$ that does not depend on j , which makes Equation (6.8) correct. In the proof of [164] (and in our explanation above), we implicitly ignore the negligible additive term when applying evasive LWE. Namely, when we apply the assumption with respect to \mathcal{A}_j , the lower bound for the advantage ϵ_{j+1} of \mathcal{A}_{j+1} should be $\epsilon_{j+1} \geq \epsilon_j / \text{poly}(\lambda_j) - \text{negl}(\lambda_j)$ rather than $\epsilon_{j+1} \geq \epsilon_j / \text{poly}(\lambda_j)$. This does not cause any difference when we consider the setting where ϵ_j is non-negligible in λ_j . However, for larger j , ϵ_j is negligible function in the security parameter λ_j . Concretely, the lower bound on ϵ_{n-1} obtained by ignoring the

⁶To be precise, in our context, we invoke the security of prFE rather than evasive LWE. However, since the same issue arises both in our context and the context of witness encryption [164] and the security of prFE is eventually reduced to the evasive LWE in our context, we intentionally do not distinguish the invocation of evasive LWE and prFE security here.

additive term is $\epsilon/\text{poly}(2^{(n-1)^2}\lambda^{n-1})$. If we apply evasive LWE once more with respect to $\lambda_n = 2^n\lambda^n$ to complete the proof, we have $\epsilon_n \geq \epsilon_{n-1}/\text{poly}(\lambda_n) - \text{negl}(\lambda_n)$. The RHS of the inequality may be negative, since $\epsilon_{n-1}/\text{poly}(\lambda_n)$ is some specific negligible function in λ_n and this may be smaller than the second term $\text{negl}(\lambda_n)$. Therefore, what we can derive here is the trivial bound $\epsilon_n \geq 0$, which is not enough for our purpose. To fix this issue, we introduce additional parameter κ and then modify the assumption so that the additive term is negligibly small in κ , which is set much larger than λ . Again, we believe that the same strengthening of the assumption is required for the proof in [164] as well. We define the strengthened version of evasive LWE next.

Non-Uniform κ -Evasive LWE. Let $\text{Samp} = \{\text{Samp}_\lambda\}_\lambda$ be a non-uniform sampler that takes as input 1^λ and outputs $(\mathbf{S}, \mathbf{P}, \text{aux})$ as above. For non-uniform adversaries $\mathcal{A}_0 = \{\mathcal{A}_{0,\lambda}\}_\lambda$ and $\mathcal{A}_1 = \{\mathcal{A}_{1,\lambda}\}_\lambda$, we define the advantage functions $\mathcal{A}_{\text{Adv}_0}^{\text{PRE}}(\lambda)$ and $\mathcal{A}_{\text{Adv}_1}^{\text{POST}}(\lambda)$ as in Equation (6.5) and Equation (6.6), respectively. For a function $\kappa := \kappa(\lambda)$ of the security parameter λ , we say that the non-uniform κ -evasive LWE assumption holds if for every non-uniform sampler Samp and a non-uniform adversary Adv_1 whose sizes are polynomial in λ' and κ respectively for $\lambda' := \lambda'(\lambda) < \kappa(\lambda)$, there exists another non-uniform adversary Adv_0 and a polynomial $Q(\cdot)$ such that

$$\mathcal{A}_{\text{Adv}_0}^{\text{PRE}}(\lambda) \geq \mathcal{A}_{\text{Adv}_1}^{\text{POST}}(\lambda)/Q(\lambda') - \text{negl}(\kappa) \quad \text{and} \quad \text{Size}(\mathcal{A}_0) \leq Q(\lambda') \cdot \text{Size}(\mathcal{A}_1). \quad (6.9)$$

There are two main differences from the standard **EvLWE** assumption: (i) we consider non-uniform sampler and adversaries, (ii) it is parameterized by κ and the additive term $\text{negl}(\lambda)$ that appears in Equation (6.7) is replaced by $\text{negl}(\kappa)$ in Equation (6.9). These changes are introduced so that our **prFE** can satisfy a stronger notion of security, as discussed below. Note that in the case λ' is superpolynomial in λ , $\text{Samp}(1^\lambda)$ outputs \mathbf{S} , \mathbf{P} , and aux whose sizes are polynomial in λ' and thus superpolynomial in λ . We require the above assumption with $\kappa = 2^{\text{poly}(\lambda)}$ in our construction.

prFE with Non-Uniform κ Security. We now describe a stronger notion of security from prFE, which we need for our mi-prFE compiler.

We say that a prFE scheme is secure in the non-uniform κ setting if for every Samp and an adversary \mathcal{A}_1 whose sizes are polynomial in λ' and κ respectively for $\lambda' := \lambda'(\lambda) < \kappa(\lambda)$, there exists another adversary \mathcal{A}_0 such that

$$\mathcal{A}_{\mathcal{A}_0}^{\text{PRE}}(\lambda) \geq \mathcal{A}_{\mathcal{A}_1}^{\text{POST}}(\lambda)/Q(\lambda') - \text{negl}(\kappa) \quad (6.10)$$

where the advantage functions $\mathcal{A}_{\text{Adv}_0}^{\text{PRE}}(\lambda)$ and $\mathcal{A}_{\text{Adv}_1}^{\text{POST}}(\lambda)$ are defined as in Equation (6.1) and Equation (6.2) respectively, and $\text{Size}(\mathcal{A}_0) \leq \text{Size}(\mathcal{A}_1) \cdot Q(\lambda')$ for some polynomial $Q(\cdot)$.

The new security notion differs from the previous one in two ways: (i) it considers non-uniform adversaries instead of uniform adversaries, (ii) it is parameterized by κ and the additive term $\text{negl}(\lambda)$ that appears in Equation (6.3) is replaced by $\text{negl}(\kappa)$ in Equation (6.10). By taking κ asymptotically larger than λ (e.g., $\kappa := \lambda^\lambda$), we can make the additive term $\text{negl}(\kappa)$ much smaller than $\text{negl}(\lambda)$.

Fortunately, we can show, with some careful adjustments, that the prFE scheme from [12] also satisfies this stronger notion of security. In Section 6.A.1, we prove the following theorem.

Theorem 6.6. *Let $\kappa = 2^{\lambda^c}$ for some constant c . Assuming non-uniform κ -evasive LWE, subexponentially secure PRF against non-uniform adversary, and non-uniform sub-exponential LWE, there exists a prFE scheme for function class $\mathcal{F}_{L(\lambda), \ell(\lambda), \text{dep}(\lambda)} = \{f : \{0, 1\}^L \rightarrow \{0, 1\}^\ell\}$ satisfying κ -prCT security as per Definition 6.3 with efficiency*

$$|\text{mpk}| = L \cdot \text{poly}(\text{dep}, \lambda), \quad |\text{sk}_f| = \ell \cdot \text{poly}(\text{dep}, \lambda), \quad |\text{ct}| = L \cdot \text{poly}(\text{dep}, \lambda).$$

where $\text{dep} = \text{poly}(\lambda)$ is the depth bound on the functions supported by the scheme.

With the above modifications, we are ready to state our final theorem:

Theorem 6.7 (mi-prFE for poly arity). *Let $\kappa = \lambda^{n^2 \log \lambda}$. Assume non-uniform κ -evasive LWE, non-uniform sub-exponential PRF, and non-uniform sub-exponential LWE. Then there exists a mi-prFE scheme for arity $n = \text{poly}(\lambda)$, supporting functions with input length L and bounded polynomial depth $d = d(\lambda)$, and satisfying κ -security (Definition 6.10) with efficiency $|\text{mpk}| = \text{poly}(n, L, d, \Lambda, \lambda)$, $|\text{sk}_f| = \text{poly}(d, \lambda)$, $|\text{ct}_1| = nL \text{poly}(\text{dep}, \lambda)$, $|\text{ct}_i| = \text{poly}(n, L, d, \Lambda, \lambda)$ for $i \in [2, n]$ where $\Lambda := (n^2 \lambda)^{1/\delta}$.*

Since the aforementioned issues in the proof only arise when evasive LWE is applied super-constant number of times, we do not need the stronger version of evasive LWE to support constant arity. Hence we get:

Theorem 6.8 (mi-prFE for constant arity). *Let $\kappa = \lambda^n$. Assume evasive LWE and LWE. Then there exists a mi-prFE scheme for arity $n = O(1)$, supporting functions with input length L and bounded polynomial depth $d = d(\lambda)$, and satisfying κ -security (Definition 6.10) with efficiency $|\text{mpk}| = \text{poly}(n, L, d, \lambda)$, $|\text{sk}_f| = \text{poly}(d, \lambda)$, $|\text{ct}_1| = nL \cdot \text{poly}(\text{dep}, \lambda)$, $|\text{ct}_i| = \text{poly}(n, L, d, \lambda)$ for $i \in [2, n]$.*

***iO* for Pseudorandom Functionalities.** We now use our multi-input FE to construct *iO*, à la AJ/BV [29, 41] for the same functionality. We begin with the definition of *iO* for pseudorandom functionalities, which we refer to as **prIO**. Recall that this notion was first defined by Mathialagan, Peters and Vaikuntanathan [146]. The syntax of **prIO** is as in regular *iO*, where we have (i) an obfuscation algorithm which takes as input the security parameter λ and a circuit C and outputs an obfuscated circuit \tilde{C} , (ii) an evaluation algorithm takes as input an obfuscated circuit \tilde{C} and an input x . It outputs $y = C(x)$. We also require that the evaluation time of the obfuscated circuit be only polynomially slower than the run time of the circuit C on x .

Our notion of security however, differs from the standard notion considered in the literature of *iO* and is specified as follows. For the security parameter $\lambda = \lambda(\lambda)$, let Samp

be a PPT algorithm that on input 1^λ , outputs

$$(C_0, C_1, \text{aux} \in \{0, 1\}^*)$$

where $C_0 : \{0, 1\}^n \rightarrow \{0, 1\}^m$ and $C_1 : \{0, 1\}^n \rightarrow \{0, 1\}^m$ have the same description size. We then require that

$$\begin{aligned} \text{If } (1^\kappa, \{C_0(x)\}_{x \in \{0, 1\}^n}, \text{aux}) &\approx_c (1^\kappa, \{\Delta_x \leftarrow \{0, 1\}^m\}_{x \in \{0, 1\}^n}, \text{aux}) \\ &\approx_c (1^\kappa, \{C_1(x)\}_{x \in \{0, 1\}^n}, \text{aux}) \end{aligned} \quad (6.11)$$

$$\text{then } (iO(1^\lambda, C_0), \text{aux}) \approx_c (iO(1^\lambda, C_1), \text{aux}) \quad (6.12)$$

where the parameter κ above is introduced to adjust the strength of the requirement for the precondition, similarly to the case of mi-prFE. Roughly speaking, the above security definition says that the obfuscations of two circuits with pseudorandom truth tables are indistinguishable.

Our construction follows the blueprint of the multi-input FE to iO conversion by Ananth and Jain [29]. Briefly, the obfuscation of a circuit $C : \{0, 1\}^n \rightarrow \{0, 1\}^m$ using a $(n + 1)$ input mi-prFE scheme is

$$\{\text{mi-prFE.sk}_U, \text{mi-prFE.ct}_{1,0}, \text{mi-prFE.ct}_{1,1}, \dots, \text{mi-prFE.ct}_{n,0}, \text{mi-prFE.ct}_{n,1}, \text{mi-prFE.ct}_{n+1,C}\}$$

where mi-prFE.sk_U is the mi-prFE functional key corresponding to the universal circuit U such that $U(x_1, \dots, x_n, C) = C(x_1, \dots, x_n)$, $\text{mi-prFE.ct}_{i,b}$ for $i \in [n]$, $b \in \{0, 1\}$ denotes the i -th slot mi-prFE ciphertext encoding bit b , and $\text{mi-prFE.ct}_{n+1,C}$ denotes the $(n + 1)$ -th slot mi-prFE ciphertext encoding the circuit C . The evaluation algorithm on input $\mathbf{x} = (x_1, \dots, x_n)$ runs $\text{mi-prFE.Dec}(\text{mi-prFE.sk}_U, \text{mi-prFE.ct}_{1,x_1}, \dots, \text{mi-prFE.ct}_{n,x_n}, \text{mi-prFE.ct}_{n+1,C})$.

Correctness as well as security follow from those of mi-prFE. This leads to the following theorem, shown in Section 6.5.

Theorem 6.9. *Let $\kappa = \lambda^{n^2 \log \lambda}$. Assuming non-uniform κ -evasive LWE, non-uniform*

sub-exponential PRF,, and non-uniform sub-exponential LWE, there exists a prIO scheme for circuits.

iO for Pseudorandom Functionalities with Polynomial Domains. We also define a variant of prIO which only supports polynomial sized domains. We denote this variant by pPRIO. The advantage of considering this restricted variant is that we can base the security of the construction on (plain) evasive LWE, rather than non-uniform κ -variant of it. For applications we consider stronger variant of the primitive where we decompose the obfuscation algorithm into online-offline parts and consequently define a reusable security for this variant. In more detail, we show the following:

Theorem 6.10. *Assuming evasive LWE and LWE, there exists a secure pPRIO scheme supporting circuits of bounded size $L = \text{poly}(\lambda)$ with $|\text{Obf}_{\text{off}}| = \text{poly}(L, \lambda)$, $|\text{Obf}_{\text{on}}| = \text{poly}(L, \lambda)$, where Obf_{off} and Obf_{on} refer to the offline and online part of the obfuscated program, respectively.*

We refer the reader to Section 6.6 for details. As discussed in Section 6.1, we use this construction in our companion paper [12] to construct ABE for optimal parameters.

Applications. We show that our new tool of MIFE/iO for pseudorandom functionalities is quite powerful and yields several interesting applications.

Multi-Input Predicate Encryption for Circuits. A multi-input predicate encryption (miPE) scheme [25] for n -ary functions $\text{miPE} = (\text{Setup}, \text{KeyGen}, \text{Enc}_1, \dots, \text{Enc}_n, \text{Dec})$ generalizes single input predicate encryption [109] to support multiple encryptors, who each encrypt their data with independently chosen randomness. In miPE, the setup algorithm on input 1^λ , arity 1^n and parameter prm , specifying the parameters of the function class, outputs (mpk, msk) . The key generation algorithm on input msk and a function

$f : (\mathcal{X}_{\text{prm}})^n \rightarrow \mathcal{Y}_{\text{prm}}$ outputs a functional secret key sk_f . The i -th encryption algorithm on input msk , an attribute $x_i \in \mathcal{X}_{\text{prm}}$ and a message $\mu_i \in \{0, 1\}$ outputs a ciphertext ct_i . The decryption algorithm on input secret key sk_f and n ciphertexts $\text{ct}_1, \dots, \text{ct}_n$ (corresponding to inputs $(x_1, \mu_1) \dots, (x_n, \mu_n)$ respectively) outputs a string $\mu' \in \{0, 1\}^n \cup \perp$.

We prove the following security guarantee for miPE: Consider an adversary \mathcal{A} which outputs q_0 key queries $\{f_0, \dots, f_{q_0}\}$, q_i pairs of attribute-message queries for i -th slot $\{(x_1^{j_1}, \mu_1^{j_1})\}_{j_1 \in [q_1]}, \dots, \{(x_n^{j_n}, \mu_n^{j_n})\}_{j_n \in [q_n]}$ and auxiliary information $\text{aux}_{\mathcal{A}}$. We say that a miPE scheme is secure if the following holds for the adversary \mathcal{A}

$$\begin{aligned} & \left(\text{mpk}, \{f_k, \text{sk}_{f_k}\}_{k \in [q_0]}, \left\{ \text{ct}_i^{j_i} \leftarrow \text{Enc}_i(\text{msk}, x_i^{j_i}, \mu_i^{j_i}) \right\}_{i \in [n], j_i \in [q_i]}, \text{aux}_{\mathcal{A}} \right) \\ & \approx_c \left(\text{mpk}, \{f_k, \text{sk}_{f_k}\}_{k \in [q_0]}, \left\{ \delta_i^{j_i} \leftarrow \text{Sim}(\text{msk}) \right\}_{i \in [n], j_i \in [q_i]}, \text{aux}_{\mathcal{A}} \right) \end{aligned}$$

given $f_k(x_1^{j_1}, \dots, x_n^{j_n}) = 0$ for every $i \in [n], j_i \in [q_i]$, and $k \in [q_0]$, where $(\text{mpk}, \text{msk}) \leftarrow \text{Setup}(1^\lambda, 1^n, \text{prm})$ and $\text{sk}_{f_k} \leftarrow \text{KeyGen}(\text{msk}, f_k)$.

Previously, the works of [25, 87] defined the notion of multi-input predicate encryption and provided the first constructions for specific functionalities. The follow-up work of Agrawal et al. [22] supported the most general functionality – it allowed to compute arbitrary predicates in \mathbf{P} on vector $(\mathbf{x}_1, \dots, \mathbf{x}_n)$ where \mathbf{x}_i is encrypted by party $i \in [n]$ and k is a constant. The miPE of this work is obtained by first building a multi-input ABE, which is then compiled into miPE using a generic compiler by [25]. Their multi-input ABE for constant arity is quite complex and leverages intricate algebraic properties of the underlying building blocks. Moreover, the limitation for a constant k seems inherent to their techniques, since the parameters grow exponentially in n . In contrast, our construction is extremely simple and can bootstrap a single-input PE scheme to a polynomial-input one generically by simply using an mi-prFE to generate PE ciphertext using randomness jointly chosen by the encryptors. The PE must have pseudorandom ciphertext so as to be suitable for the compiler but this is a relatively mild property and

readily satisfied by known constructions [109]. In more detail, our construction works as follows.

- The setup generates a mi-prFE instance (mi-prFE.msk, mi-prFE.mpk) and a single-input PE instance (PE.msk, PE.mpk). It outputs $\text{msk} = (\text{mi-prFE.msk}, \text{PE.msk})$ and $\text{mpk} = (\text{mi-prFE.mpk}, \text{PE.mpk})$.
- The i -th slot encryption algorithm on input $(\text{msk}, \mathbf{x}_i, \mu_i)$ generates an i -th slot mi-prFE ciphertext $\text{mi-prFE.ct}_i \leftarrow \text{mi-prFE.Enc}_i(\text{mi-prFE.msk}, (\mathbf{x}_i, \mu_i))$. It outputs $\text{ct}_i = \text{mi-prFE.ct}_i$.
- The key generator on input msk and a function f generates a single-input PE functional secret key $\text{PE.sk}_f \leftarrow \text{PE.KeyGen}(\text{PE.msk}, f)$. It also generates a mi-prFE key, mi-prFE.sk_F , for function $F[\text{PE.mpk}]$ that, on input n attribute-message pairs $(\mathbf{x}_1, \mu_1), \dots, (\mathbf{x}_n, \mu_n)$, generates a single-input PE ciphertext w.r.t. attribute $\mathbf{x} = (\mathbf{x}_1, \dots, \mathbf{x}_n)$ and message $\mu = (\mu_1, \dots, \mu_n)$. It outputs $\text{sk}_f = (\text{PE.sk}_f, \text{mi-prFE.sk}_F)$.
- The decryption algorithm first runs the mi-prFE decryption using mi-prFE.sk_F and $\{\text{ct}_i = \text{mi-prFE.ct}_i\}_{i \in [n]}$ to compute the single-input PE ciphertext, PE.ct , encoding message $\mu = (\mu_1, \dots, \mu_n)$ w.r.t attribute $\mathbf{x} = (\mathbf{x}_1, \dots, \mathbf{x}_n)$. Finally it performs PE decryption using PE.sk_f and PE.ct .

Correctness and security follow readily from those of the underlying building blocks.

Please see Section 6.4 for details. In summary, we obtain the following theorems.

Theorem 6.11 (miPE for poly arity). *Let $\kappa = \lambda^{n^2 \log \lambda}$. Assuming non-uniform κ -evasive LWE, non-uniform sub-exponential PRF, and non-uniform sub-exponential LWE, there exists a miPE scheme for arity $n = \text{poly}(\lambda)$, supporting functions of bounded polynomial depth.*

Using mi-prFE scheme for constant arity allows to relax the assumption to normal evasive LWE.

Theorem 6.12 (miPE for constant arity). *Assuming evasive LWE and LWE, there exists a mi-prFE scheme for arity $n = O(1)$, supporting functions of bounded polynomial depth.*

Two Party ID based Non-Interactive Key Exchange. Next, we provide a construction of two party ID based non-interactive key exchange (ID-NIKE) scheme. The construction is the same as the ID-based NIKE system by Sakai, Ohgishi, and Kasahara [159] except

that the hash function is replaced with an obfuscation of a PRF, which can be supported by our prIO . In more detail, we show:

Theorem 6.13. *Let $\kappa = \lambda^{n^2 \log \lambda}$. Assuming non-uniform κ -evasive LWE (Assumption 6.14), non-uniform sub-exponential PRF, non-uniform sub-exponential LWE, and the DBDH assumption, there exists a secure ID-NIKE scheme.*

This leads to the first construction of ID-NIKE without multilinear maps [88] or indistinguishability obfuscation in the standard model. Please see Section 6.8 for details.

Instantiating the Random Oracle. In an elegant work [119], Hohenberger, Sahai and Waters posed the following question: “Can we instantiate the random oracle with an actual family of hash functions for existing cryptographic schemes in the random oracle model, such as Full Domain Hash signatures?” They then demonstrated that the selective security of the full-domain hash (FDH) signature based on trapdoor permutations (TDP) [37], the adaptive security of RSA FDH signatures [79], the selective security of BLS signatures, and the adaptive security of BLS signatures [45] can be proven in the standard model by carefully instantiating the underlying hash function by IO for each application.

We show in Section 6.7.1, that the random oracle in the FDH signature can be instantiated using prIO instead of full-fledged $i\mathcal{O}$. Similarly, we can instantiate the random oracle in selectively secure BLS signatures with prIO , following a strategy similar to that in [119]. At a high level, these proofs follow those in the random oracle model (ROM), where we use $i\mathcal{O}$ to obfuscate a derandomized version of the simulator for the hash function in ROM-based proofs. In these settings, the truth table of the simulated hash function is pseudorandom, allowing us to follow the same proof strategy using prIO .

For adaptively secure RSA FDH and BLS signatures, the situation is different. In these cases, Hohenberger et al. adopt an alternative proof strategy that deviates from the high level strategy of obfuscating the simulator for the proof in the ROM. This is due to the fact that the original proofs [45, 79] are incompatible with the conditions required for

using iO , where the truth table of the hash functions must remain unchanged across game hops. To be compatible with iO , they introduce a structure for the hash function, making its truth table no longer pseudorandom. This prevents us from replacing the hash function with prIO following their approach.

To instantiate the hash function with prIO , we revert to the original ROM-based proof strategy [45, 79]. Unlike the iO -based approach, prIO -based proof does not require the truth table of the hash function to remain unchanged across game hops; it only requires the truth table to be pseudorandom. This relaxed condition enables the use of the original ROM security proofs. Please see Section 6.7 for details.

Organization of the Chapter We provide the preliminaries used in this work in Section 6.2. In Section 6.3, we define the notion of multi-input FE for pseudorandom functionalities (mi-prFE) and construct a (bounded-depth) mi-prFE using a single-input FE scheme for pseudorandom functionalities (prFE) – a tool from our companion paper [12]. In Section 6.A, we recall the construction of prFE from [12] and prove that it achieves strengthened security notion of non-uniform κ -prCT security which is required for the mi-prFE compiler in Section 6.3. In Section 6.4, we give the construction for multi-input predicate encryption scheme for polynomial arity. In Section 6.5, we give the definitions for indistinguishability obfuscation for pseudorandom functionalities (prIO) for circuits and construct this using a mi-prFE scheme supporting polynomial arity. In Section 6.6, we define and construct indistinguishability obfuscation for pseudorandom functionalities with polynomial size domain (pPRIO) using a mi-prFE scheme supporting constant arity. In Section 6.7 and 6.8 we show how to instantiate the random oracle using a prIO scheme via various applications.

6.2 PRELIMINARIES

In this section we introduce some additional preliminaries for this chapter. We refer to Chapter 2 for preliminaries on lattices, Section 4.2.7 for properties of GSW Homomorphic encryption and evaluation and Section 4.2.8 for properties of BGG^+ homomorphic evaluation procedures.

6.2.1 Hardness Assumptions

The following variant of the evasive LWE assumption will be used in Section 6.3. This strengthens Assumption 2.6 in that it considers non-uniform samplers and replaces the negligible term in Equation (2.3) with negligible function in another parameter κ , which can be much larger than λ . The reason why we need this strengthened version of the assumption is that we need prFE to satisfy stronger security notion than prCT security that we call non-uniform κ -prCT security for the application to mi-prFE. We refer to Remark 35 for the detailed discussion on why we need the stronger security definition for prFE for the application to mi-prFE.

Assumption 6.14 (Non-Uniform κ -Evasive LWE). Let $n, m, t, m', q, \lambda \in \mathbb{N}$ be parameters defined as in Assumption 2.6 and $\text{Samp} = \{\text{Samp}_\lambda\}_\lambda$ be a non-uniform sampler that takes as input 1^λ and outputs $\mathbf{S}, \mathbf{P}, \text{aux}$ as in Assumption 2.6. For non-uniform adversaries $\mathcal{A}_0 = \{\mathcal{A}_{0,\lambda}\}_\lambda$ and $\mathcal{A}_1 = \{\mathcal{A}_{1,\lambda}\}_\lambda$, we define the advantage functions $\mathcal{A}_{\text{Adv}_0}^{\text{PRE}}(\lambda)$ and $\mathcal{A}_{\text{Adv}_1}^{\text{POST}}(\lambda)$ as in Equation (2.1) and Equation (2.2), respectively. For a function $\kappa := \kappa(\lambda)$ of the security parameter λ , we say that the non-uniform κ -evasive LWE assumption with respect to the sampler class \mathcal{SC} holds if for every non-uniform sampler $\text{Samp} \in \mathcal{SC}$ and a non-uniform adversary Adv_1 such that $\text{Size}(\text{Samp}) \leq \text{poly}(\lambda')$ and $\text{Size}(\text{Adv}_1) \leq \text{poly}(\kappa)$ for $\lambda'(\lambda) \leq \kappa(\lambda)$, there exists another non-uniform adversary Adv_0 and a polynomial $Q(\cdot)$ such that

$$\mathcal{A}_{\text{Adv}_0}^{\text{PRE}}(\lambda) \geq \mathcal{A}_{\text{Adv}_1}^{\text{POST}}(\lambda)/Q(\lambda') - \text{negl}(\kappa) \quad \text{and} \quad \text{Size}(\mathcal{A}_0) \leq Q(\lambda') \cdot \text{Size}(\mathcal{A}_1). \quad (6.13)$$

Note that in the case λ' is superpolynomial in λ , $\text{Samp}(1^\lambda)$ outputs \mathbf{S}, \mathbf{P} , and aux

whose sizes are polynomial in λ' and thus superpolynomial in λ . We require the above assumption with $\kappa = 2^{\text{poly}(\lambda)}$ in our construction.

The following lemma is an adaptation of Theorem 2.7 for the stronger version of evasive LWE assumption defined in Assumption 6.14. The proof is almost the same as that for Theorem 2.7. We provide it here for completeness.

Lemma 6.15. *Let $n, m, t, m', q, \lambda \in \mathbb{N}$ be parameters defined as in Assumption 2.6 and $\text{Samp} = \{\text{Samp}_\lambda\}_\lambda$ be a non-uniform sampler that takes as input 1^λ and outputs $\mathbf{S}, \text{aux} = (\text{aux}_1, \text{aux}_2)$, and \mathbf{P} as in Theorem 2.7. For a non-uniform adversaries \mathcal{A} , we define the advantage functions $\mathcal{A}_{\text{Adv}}^{\text{PRE}'}(\lambda)$ and $\mathcal{A}_{\text{Adv}}^{\text{POST}'}(\lambda)$ as in Equation (2.4) and Equation (2.5), respectively. Then, for a function $\kappa := \kappa(\lambda)$ of the security parameter λ , under the non-uniform κ -evasive LWE assumption (Assumption 6.14), if $\text{Size}(\text{Samp}) \leq \text{poly}(\lambda')$ and $\text{Size}(\text{Adv}_1) \leq \text{poly}(\kappa)$ for $\lambda'(\lambda) \leq \kappa(\lambda)$, there exists another non-uniform adversary Adv_0 and a polynomial $Q(\cdot)$ such that*

$$\mathcal{A}_{\text{Adv}_0}^{\text{PRE}'}(\lambda) \geq \mathcal{A}_{\text{Adv}_1}^{\text{POST}'}(\lambda)/Q(\lambda') - \text{negl}(\kappa) \quad \text{and} \quad \text{Size}(\mathcal{A}_0) \leq Q(\lambda') \cdot \text{Size}(\mathcal{A}_1).$$

Proof. Let us consider an adversary \mathcal{A}_1 and a sampler Samp with size being polynomial in κ and λ' respectively and $\epsilon = \mathcal{A}_{\mathcal{A}_1}^{\text{POST}'}$. Then, the same adversary is able to distinguish either (1) $(\mathbf{B}, \mathbf{SB} + \mathbf{E}, \mathbf{K}, \text{aux}_1, \text{aux}_2)$ from $(\mathbf{B}, \mathbf{C}_0, \mathbf{K}, \text{aux}_1, \text{aux}_2)$ with advantage at least $\epsilon/2$ or (2) $(\mathbf{B}, \mathbf{C}_0, \mathbf{K}, \text{aux}_1, \text{aux}_2)$ from $(\mathbf{B}, \mathbf{C}_0, \mathbf{K}, \mathbf{c}, \text{aux}_2)$ with advantage at least $\epsilon/2$. If the latter is the case, then we can obtain an adversary \mathcal{A}_0 that distinguishes $(\mathbf{B}, \mathbf{SB} + \mathbf{E}, \mathbf{SP} + \mathbf{E}', \text{aux}_1, \text{aux}_2)$ from $(\mathbf{B}, \mathbf{C}_0, \mathbf{C}', \mathbf{c}, \text{aux}_2)$ with advantage $\epsilon/2$. This can be seen by observing that \mathcal{A}_1 can be turned into an adversary that distinguishes $(\text{aux}_1, \text{aux}_2)$ from $(\mathbf{c}, \text{aux}_2)$ and then turned into an adversary that distinguishes $(\mathbf{B}, \mathbf{C}_0, \mathbf{K}, \text{aux}_1, \text{aux}_2)$ from $(\mathbf{B}, \mathbf{C}_0, \mathbf{K}, \mathbf{c}, \text{aux}_2)$ by sampling $(\mathbf{B}, \mathbf{C}_0, \mathbf{K})$ by itself, where we sample \mathbf{B} with the corresponding trapdoor and then sample \mathbf{K} using it. We therefore assume that the former is the case. Then, by invoking the non-uniform κ -evasive LWE with respect to the sampler Samp , we obtain another adversary \mathcal{A}_0 whose size is bounded by

$Q(\lambda') \cdot \text{Size}(\mathcal{A}_1)$ and distinguishing advantage against $(\mathbf{B}, \mathbf{SB} + \mathbf{E}, \mathbf{SP} + \mathbf{E}', \text{aux}_1, \text{aux}_2)$ and $(\mathbf{B}, \mathbf{C}_0, \mathbf{C}', \text{aux}_1, \text{aux}_2)$ is at least $\epsilon/2Q(\lambda')$. Then, \mathcal{A}_0 is able to distinguish either (1) $(\mathbf{B}, \mathbf{SB} + \mathbf{E}, \mathbf{SP} + \mathbf{E}', \text{aux}_1, \text{aux}_2)$ from $(\mathbf{B}, \mathbf{C}_0, \mathbf{C}', \mathbf{c}, \text{aux}_2)$ with advantage at least $\epsilon/4Q(\lambda')$ or (2) $(\mathbf{B}, \mathbf{C}_0, \mathbf{C}', \mathbf{c}, \text{aux}_2)$ from $(\mathbf{B}, \mathbf{C}_0, \mathbf{C}', \text{aux}_1, \text{aux}_2)$ with advantage at least $\epsilon/4Q(\lambda')$. If the former is the case, we are done. If the latter is the case, we are still able to convert it into a distinguisher against $(\mathbf{B}, \mathbf{SB} + \mathbf{E}, \mathbf{SP} + \mathbf{E}', \text{aux}_1, \text{aux}_2)$ and $(\mathbf{B}, \mathbf{C}_0, \mathbf{C}', \mathbf{c}, \text{aux}_2)$ by the similar argument to the above. ■

Assumption 6.16 (Decisional Bilinear Diffie-Hellman). Let $e : \mathbb{G} \times \mathbb{G} \rightarrow \mathbb{G}_T$ be a bilinear map, where $e(g^a, g^b) = e(g, g)^{ab}$ and $a, b \in \mathbb{Z}_p$, for a cyclic group \mathbb{G} of order p with a generator g . The Decisional Bilinear Diffie-Hellman assumption (DBDH) states that for any PPT adversary

$$\left(g^\alpha, g^\beta, g^\gamma, e(g, g)^{\alpha\beta\gamma} \mid \alpha, \beta, \gamma \leftarrow \mathbb{Z}_p \right) \approx_c \left(g^\alpha, g^\beta, g^\gamma, T \mid \alpha, \beta, \gamma \leftarrow \mathbb{Z}_p, T \leftarrow \mathbb{G}_T \right)$$

6.2.2 Puncturable Pseudorandom Functions

Syntax. A puncturable pseudorandom function (PPRF) $F : \mathcal{K} \times \mathcal{X} \rightarrow \mathcal{Y}$ with key space \mathcal{K} , input space \mathcal{X} and output space \mathcal{Y} has the following syntax.

Setup $(1^\lambda) \rightarrow K$. The setup algorithm takes as input the security parameter λ and outputs a key $K \in \mathcal{K}$.

puncture $(K, x) \rightarrow K_x$. The puncture algorithm takes as input a PRF key $K \in \mathcal{K}$ and an input $x \in \mathcal{X}$, and outputs a punctured key K_x .

Eval $(K_x, x') \rightarrow y$. The evaluation algorithm takes as input a punctured key K_x an input $x' \in \mathcal{X}$, such that $x \neq x'$ and outputs $y \in \mathcal{Y}$. It outputs \perp if $x = x'$.

Definition 6.1. (Correctness) A PPRF scheme is said to be correct if for any $K \in \mathcal{K}$, $x, x' \in \mathcal{X}$ such that $x \neq x'$, we have

$$\Pr[\text{Eval}(K_x, x') = F(K, x') \mid K_x \leftarrow \text{puncture}(K, x)] = 1.$$

Definition 6.2. (Security) A PPRF scheme is said to be (selectively) secure if the advantage of a PPT adversary \mathcal{A} in the following experiment is negligible.

1. \mathcal{A} on input 1^λ outputs the challenge input x^\star .
2. The challenger samples a random key $K \leftarrow \mathcal{K}$ and a bit $\beta \leftarrow \{0, 1\}$. Then, it computes $y = F(K, x)$ if $\beta = 0$, else it sample $y \leftarrow \mathcal{Y}$ uniformly at random. It also computes $K_{x^\star} \leftarrow \text{puncture}(K, x^\star)$ and sends K_{x^\star}, y to \mathcal{A} .
3. \mathcal{A} outputs a guess bit β' .

\mathcal{A} wins if $\beta = \beta'$.

We know that PPRF with the security defined above from one-way functions [100, 47, 51, 134].

6.2.3 Pseudorandom Functional Encryption

In this section we define an additional security notion for functional encryption for pseudorandom functionalities. We refer to Section 5.3 for definitions of a prFE scheme.

Definition 6.3 (Non-uniform κ -prCT Security). For a prFE scheme for function family $\{\mathcal{F}_{\text{prm}} = \{f : \mathcal{X}_{\text{prm}} \rightarrow \mathcal{Y}_{\text{prm}}\}\}_{\text{prm}}$, parameter $\text{prm} = \text{prm}(\lambda)$, and function $\kappa \stackrel{\text{def}}{=} \kappa(\lambda)$ of λ , let $\text{Samp} = \{\text{Samp}_\lambda\}_\lambda$ be a non-uniform polynomial-time algorithm that on input 1^λ , outputs

$$(f_1, \dots, f_{Q_{\text{key}}}, x_1, \dots, x_{Q_{\text{msg}}}, \text{aux} \in \{0, 1\}^*)$$

where Q_{key} is the number of key queries, Q_{msg} is the number of message queries, and $f_i \in \mathcal{F}_{\text{prm}}, x_j \in \mathcal{X}_{\text{prm}}$ for all $i \in [Q_{\text{key}}], j \in [Q_{\text{msg}}]$.

For non-uniform adversaries $\mathcal{A}_0 := \{\mathcal{A}_{0,\lambda}\}_\lambda$ and $\mathcal{A}_1 := \{\mathcal{A}_{1,\lambda}\}_\lambda$, we define the following advantage functions:

$$\begin{aligned} \text{Adv}_{\mathcal{A}_0}^{\text{PRE}}(\lambda) &\stackrel{\text{def}}{=} \Pr[\mathcal{A}_0(\text{aux}, \{f_i, f_i(x_j)\}_{i \in [Q_{\text{key}}], j \in [Q_{\text{msg}}]}) = 1] \\ &\quad - \Pr[\mathcal{A}_0(\text{aux}, \{f_i, \Delta_{i,j} \leftarrow \mathcal{Y}_{\text{prm}}\}_{i \in [Q_{\text{key}}], j \in [Q_{\text{msg}}]}) = 1] \end{aligned}$$

$$\begin{aligned} \text{Adv}_{\mathcal{A}_1}^{\text{POST}}(\lambda) &\stackrel{\text{def}}{=} \Pr[\mathcal{A}_1(\text{mpk}, \text{aux}, \{f_i, \text{ct}_j \leftarrow \text{Enc}(\text{mpk}, x_j), \text{sk}_{f_i}\}_{i \in [Q_{\text{key}}], j \in [Q_{\text{msg}}]}) = 1] \\ &\quad - \Pr[\mathcal{A}_1(\text{mpk}, \text{aux}, \{f_i, \delta_j \leftarrow \text{CT}, \text{sk}_{f_i}\}_{i \in [Q_{\text{key}}], j \in [Q_{\text{msg}}]}) = 1] \end{aligned}$$

where $(f_1, \dots, f_{Q_{\text{key}}}, x_1, \dots, x_{Q_{\text{msg}}}, \text{aux} \in \{0, 1\}^*) \leftarrow \text{Samp}(1^\lambda)$, $(\text{mpk}, \text{msk}) \leftarrow \text{Setup}(1^\lambda, \text{prm})$ and \mathcal{CT} is the ciphertext space. We say that a prFE scheme for function family \mathcal{F}_{prm} is secure in the non-uniform κ setting with respect to the sampler class \mathcal{SC} if for every sampler $\text{Samp} \in \mathcal{SC}$ and an adversary \mathcal{A}_1 such that $\text{Size}(\text{Samp}) \leq \text{poly}(\lambda')$ and $\text{Size}(\mathcal{A}_1) \leq \text{poly}(\kappa)$ for $\lambda' \leq \kappa$, there exists another adversary \mathcal{A}_0 such that

$$\mathcal{A}_{\mathcal{A}_0}^{\text{PRE}}(\lambda) \geq \mathcal{A}_{\mathcal{A}_1}^{\text{POST}}(\lambda)/Q(\lambda') - \text{negl}(\kappa) \quad (6.14)$$

and $\text{Size}(\mathcal{A}_0) \leq \text{Size}(\mathcal{A}_1) \cdot Q(\lambda')$ for some polynomial $Q(\cdot)$.

In Section 6.A.1, we prove the following theorem.

Theorem 6.17. *Let $\kappa = 2^{\lambda^c}$ for some constant c . Assuming non-uniform κ -evasive LWE (Assumption 6.14), subexponentially secure PRF against non-uniform adversary, and non-uniform sub-exponential LWE (Assumption 2.5), there exists a prFE scheme for function class $\mathcal{F}_{L(\lambda), \ell(\lambda), \text{dep}(\lambda)} = \{f : \{0, 1\}^L \rightarrow \{0, 1\}^\ell\}$ satisfying κ -prCT security as per Definition 6.3, with respect to a specific sampler class, with efficiency*

$$|\text{mpk}| = L \cdot \text{poly}(\text{dep}, \lambda), \quad |\text{sk}_f| = \ell \cdot \text{poly}(\text{dep}, \lambda), \quad |\text{ct}| = L \cdot \text{poly}(\text{dep}, \lambda).$$

where $\text{dep} = \text{poly}(\lambda)$ is the depth bound on the functions supported by the scheme.

Remark 31 (Comparison between Definition 6.3 and Definition 5.13). We remark that Definition 6.3 strengthens Definition 5.13 in two aspects. First of all, it considers non-uniform adversaries instead of uniform adversaries. Secondly, it is parameterized by κ and the additive term $\text{negl}(\lambda)$ is replaced by $\text{negl}(\kappa)$ in Equation (6.14). By taking κ asymptotically larger than λ (e.g., $\kappa := \lambda^\lambda$), we can make the additive term $\text{negl}(\kappa)$ much smaller than $\text{negl}(\lambda)$. We note that these changes are introduced to prove the security of our mi-prFE in Section 6.3. We refer to Remark 35 for the discussion on why these changes are necessary for the security proof there.

Theorem 6.18 ([12]). *Assuming LWE and evasive LWE assumptions, there exists a secure (Definition 5.13) prFE scheme, with respect to a specific sampler class, for*

function class $\mathcal{F}_{L(\lambda), \ell(\lambda), \text{dep}(\lambda)} = \{f : \{0, 1\}^L \rightarrow \{0, 1\}^\ell\}$ satisfying

$$|\text{mpk}| = L \cdot \text{poly}(\text{dep}, \lambda), \quad |\text{sk}_f| = \ell \cdot \text{poly}(\text{dep}, \lambda), \quad |\text{ct}| = L \cdot \text{poly}(\text{dep}, \lambda).$$

where $\text{dep} = \text{poly}(\lambda)$ is the depth bound on the functions supported by the scheme.

6.2.4 Predicate Encryption

Consider a function family $\{\mathcal{F}_{\text{prm}} = \{f : \mathcal{X}_{\text{prm}} \rightarrow \{0, 1\}\}\}_{\text{prm}}$, for a parameter $\text{prm} = \text{prm}(\lambda)$.

Syntax. A PE scheme PE for function family \mathcal{F}_{prm} consists of polynomial time algorithms (Setup, KeyGen, Enc, Dec) defined as follows.

Setup($1^\lambda, \text{prm}$) \rightarrow (mpk, msk). The setup algorithm takes as input the security parameter λ and a parameter prm , and outputs a master public key mpk and master secret key msk⁷.

KeyGen(msk, f) \rightarrow sk_f . The key generation algorithm takes as input the master secret key msk and a function $f \in \mathcal{F}_{\text{prm}}$ and it outputs a functional secret key sk_f .

Enc(mpk, x, μ) \rightarrow ct. The encryption algorithm takes as input a master secret key msk, an attribute $x \in \mathcal{X}_{\text{prm}}$, and message $\mu \in \{0, 1\}$, and outputs a ciphertext ct.

Dec(mpk, $\text{sk}_f, f, \text{ct}$) \rightarrow $\{0, 1\} \cup \perp$. The decryption algorithm takes as input the master public key mpk, secret key sk_f , function f and ciphertext ct, and outputs a string $\mu' \in \{0, 1\} \cup \perp$.

Definition 6.4 (Correctness.). For every $\lambda \in \mathbb{N}$, $\mu \in \{0, 1\}$, $x \in \mathcal{X}_{\text{prm}}$, $f \in \mathcal{F}_{\text{prm}}$, if

⁷We assume w.l.o.g that msk includes mpk.

$f(x) = 1$, then

$$\Pr \left[\text{Dec}(\text{mpk}, \text{KeyGen}(\text{msk}, f), f, \text{Enc}(\text{mpk}, x, \mu)) = \mu \right] = 1 - \text{negl}(\lambda)$$

where $(\text{mpk}, \text{msk}) \leftarrow \text{Setup}(1^\lambda, \text{prm})$.

Definition 6.5 (Selective IND_r Security). A PE scheme is said to satisfy selective IND_r security if there exists a negligible function $\text{negl}(\cdot)$ such that for all $\lambda \in \mathbb{N}$, we have

$$\Pr \left[\begin{array}{l} (\mathbf{x}, \text{prm}) \leftarrow \mathcal{A}(1^\lambda); \\ (\text{mpk}, \text{msk}) \leftarrow \text{Setup}(1^\lambda, \text{prm}); \\ \beta' = \beta : (\mu, \text{st}) \leftarrow \mathcal{A}^{\text{KeyGen}(\text{msk}, \cdot)}(\text{mpk}); \\ \text{ct}_0 \leftarrow \text{Enc}(\text{mpk}, \mathbf{x}, \mu), \text{ct}_1 \leftarrow C\mathcal{T}; \\ \beta \leftarrow \{0, 1\}, \beta' \leftarrow \mathcal{A}^{\text{KeyGen}(\text{msk}, \cdot)}(\text{st}, \text{ct}_\beta) \end{array} \right] - \frac{1}{2} \leq \text{negl}(\lambda), \quad (6.15)$$

where $C\mathcal{T}$ is the ciphertext space of the scheme and the adversary \mathcal{A} is admissible in the sense that for all key query f made by \mathcal{A} , it holds that $f(\mathbf{x}) = 0$.

Predicate encryption schemes for (bounded depth) circuits satisfying the above security notion are known from LWE [108, 111, 172].

6.2.5 Multi-Input Predicate Encryption

In this section we define multi-input Predicate Encryption (mi-PE), adapting the syntax from [25]. Consider a function family $\{\mathcal{F}_{\text{prm}} = \{f : (\mathcal{X}_{\text{prm}})^n \rightarrow \{0, 1\}\}\}_{\text{prm}}$, for a parameter $\text{prm} = \text{prm}(\lambda)$, where each \mathcal{F}_{prm} is a finite collection of n -ary functions. Each function $f \in \mathcal{F}_{\text{prm}}$ takes as input strings x_1, \dots, x_n , where each $x_i \in \mathcal{X}_{\text{prm}}$ and outputs $f(x_1, \dots, x_n) \in \{0, 1\}$.

Syntax. A mi-PE scheme miPE_n for n -ary function family \mathcal{F}_{prm} consists of polynomial time algorithms $(\text{Setup}, \text{KeyGen}, \text{Enc}_1, \dots, \text{Enc}_n, \text{Dec})$ defined as follows.

$\text{Setup}(1^\lambda, 1^n, \text{prm}) \rightarrow (\text{mpk}, \text{msk})$. The setup algorithm takes as input the security parameter λ , the function arity n and a parameter prm and outputs a master public

key mpk and master secret key msk .

$\text{KeyGen}(\text{msk}, f) \rightarrow \text{sk}_f$. The key generation algorithm takes as input the master secret key msk and a function $f \in \mathcal{F}_{\text{prm}}$ and it outputs a functional secret key sk_f .

$\text{Enc}_i(\text{msk}, x_i, \mu_i) \rightarrow \text{ct}_i$ for $i \in [n]$. The encryption algorithm for the i^{th} slot takes as input a master secret key msk , an attribute $x_i \in \mathcal{X}_{\text{prm}}$, and message $\mu_i \in \{0, 1\}$, and outputs a ciphertext ct_i .

$\text{Dec}(\text{mpk}, \text{sk}_f, f, \text{ct}_1, \text{ct}_2, \dots, \text{ct}_n) \rightarrow \{0, 1\}^n \cup \perp$. The decryption algorithm takes as input the master public key mpk , secret key sk_f , function f and n ciphertexts $\text{ct}_1, \dots, \text{ct}_n$ and outputs a string $\mu' \in \{0, 1\}^n \cup \perp$.
Next, we define correctness and security.

Correctness: For every $\lambda \in \mathbb{N}$, $\mu_1, \dots, \mu_n \in \{0, 1\}$, $x_1, \dots, x_n \in \mathcal{X}_{\text{prm}}$, $f \in \mathcal{F}_{\text{prm}}$, it holds that if $f(x_1, \dots, x_n) = 1$, then

$$\Pr \left[\text{Dec} \left(\begin{array}{c} \text{mpk}, \text{KeyGen}(\text{msk}, f), f, \\ \text{Enc}_1(\text{msk}, x_1, \mu_1), \dots, \text{Enc}_n(\text{msk}, x_n, \mu_n) \end{array} \right) = (\mu_1, \dots, \mu_n) \right] = 1 - \text{negl}(\lambda)$$

where the probability is over the choice of $(\text{mpk}, \text{msk}) \leftarrow \text{Setup}(1^\lambda, 1^n, \text{prm})$ and over the internal randomness of KeyGen and $\text{Enc}_1, \dots, \text{Enc}_n$.

Definition 6.6 (Sim-Security.). For a miPE scheme for function family $\{\mathcal{F}_{\text{prm}} = \{f : (\mathcal{X}_{\text{prm}})^n \rightarrow \{0, 1\}\}\}_{\text{prm}}$, parameter $\text{prm} = \text{prm}(\lambda)$, a stateful adversary \mathcal{A} , and a simulator algorithm $\{\text{Sim}_i\}_{i \in [n]}$, we define the Sim-security game, $\text{Exp}_{\text{miPE}, \mathcal{A}}$, as follows.

1. **Query phase:** On input $1^\lambda, 1^n, \text{prm}$, \mathcal{A} outputs the following in an arbitrary order.
 - a) **Key Queries:** \mathcal{A} issues polynomial number of key queries, say $q_0 = q_0(\lambda)$. For each key query $k \in [q_0]$, \mathcal{A} chooses a function $f_k \in \mathcal{F}_{\text{prm}}$.

- b) **Ciphertext Queries:** \mathcal{A} issues polynomial number of ciphertext queries for each slot, say $q_i = q_i(\lambda)$ for the i^{th} slot. We use $(x_i^{j_i}, \mu_i^{j_i})$ to denote the j_i -th ciphertext query corresponding to the i -th slot, where $j_i \in [q_i]$ and $i \in [n]$.
2. **Setup phase:** On input $1^\lambda, 1^n, \text{prm}, \{f_k\}_{k \in [q_0]}$, the challenger samples $(\text{mpk}, \text{msk}) \leftarrow \text{Setup}(1^\lambda, 1^n, \text{prm})$, a bit $\beta \leftarrow \{0, 1\}$ and does the following.
- a) It computes $\text{sk}_{f_k} \leftarrow \text{KeyGen}(\text{msk}, f_k)$.
- b) If $\beta = 0$, it computes $\text{ct}_i^{j_i} \leftarrow \text{Enc}_i(\text{msk}, x_i^{j_i}, \mu_i^{j_i})$, else if $\beta = 1$, it computes $\text{ct}_i^{j_i} \leftarrow \text{Sim}_i(\text{msk})$ for $i \in [n], j_i \in [q_i]$.
- It returns $(\text{mpk}, \{\text{sk}_{f_k}\}_{k \in [q_0]}, \{\text{ct}_i^{j_i}\}_{i \in [n], j_i \in [q_i]})$ to \mathcal{A} .

3. **Output phase:** \mathcal{A} outputs a guess bit β' as the output of the experiment.

For the adversary to be *admissible*, we require that it holds that $f_k(x_1^{j_1}, \dots, x_n^{j_n}) = 0$ for every $i \in [n], j_i \in [q_i]$, and $k \in [q_0]$. We define the advantage $\text{Adv}_{\text{miPE}, \mathcal{A}}^{\text{Sim}}$ of \mathcal{A} in the security game as

$$\text{Adv}_{\text{miPE}, \mathcal{A}}^{\text{Sim}}(1^\lambda) := |\Pr[\text{Exp}_{\text{miPE}, \mathcal{A}}(1^\lambda) = 1 | \beta = 0] - \Pr[\text{Exp}_{\text{miPE}, \mathcal{A}}(1^\lambda) = 1 | \beta = 1]|.$$

The miPE scheme is said to satisfy Sim-security if for any stateful PPT adversary \mathcal{A} , there exists a PPT simulator algorithm $\{\text{Sim}_i\}_{i \in [n]}$ such that $\text{Adv}_{\text{miPE}, \mathcal{A}}^{\text{Sim}}(1^\lambda) = \text{negl}(\lambda)$.

6.2.6 ID-Based Non-Interactive Key Exchange

In this section, we give the definitions for an identity-based non-interactive key exchange scheme, for two parties, adapted from [88].

Syntax. An identity-based non-interactive key exchange (ID-NIKE) scheme for identity space \mathcal{ID} has the following syntax.

Setup $(1^\lambda) \rightarrow (\text{mpk}, \text{msk})$ The setup algorithm takes as input the security parameter 1^λ and outputs a master public key mpk and a master secret key msk .

Extract $(\text{mpk}, \text{msk}, \text{id}) \rightarrow \text{usk}_{\text{id}}$. The key extraction algorithm takes as input the master public key mpk , the master secret key msk and an identity $\text{id} \in \mathcal{ID}$. It outputs a

user secret key usk_{id} for id .

$\text{Share}(\text{mpk}, \text{usk}_{\text{id}_1}, \text{id}_2) \rightarrow K$. The share algorithm takes as input the master public key mpk , a user secret key usk_{id_1} for an identity $\text{id}_1 \in \mathcal{ID}$ and an identity $\text{id}_2 \in \mathcal{ID}$. It outputs a shared key K .

Definition 6.7 (Correctness). An ID-NIKE scheme for an identity space \mathcal{ID} is correct if for any $\lambda \in \mathbb{N}$, $(\text{mpk}, \text{msk}) \leftarrow \text{Setup}(1^\lambda)$, any $\text{id}_1, \text{id}_2 \in \mathcal{ID}$ we have

$$\text{Share}(\text{mpk}, \text{usk}_{\text{id}_1}, \text{id}_2) = \text{Share}(\text{mpk}, \text{usk}_{\text{id}_2}, \text{id}_1)$$

where $\text{usk}_{\text{id}_1} \leftarrow \text{Extract}(\text{mpk}, \text{msk}, \text{id}_1)$ and $\text{usk}_{\text{id}_2} \leftarrow \text{Extract}(\text{mpk}, \text{msk}, \text{id}_2)$.

Definition 6.8 (Security). We say that an ID-NIKE scheme for an identity space \mathcal{ID} is secure if the advantage function

$$\text{Adv}_{\text{ID-NIKE}, \mathcal{A}}(\lambda) = \Pr [\text{Exp}_{\text{ID-NIKE}, \mathcal{A}}(\lambda) = 1] - \frac{1}{2}$$

is negligible for all PPT adversaries \mathcal{A} . Here, experiment $\text{Exp}_{\text{ID-NIKE}, \mathcal{A}}$ is defined as follows:

1. **Setup Phase.** The experiment samples $(\text{mpk}, \text{msk}) \leftarrow \text{Setup}(1^\lambda)$, and gives \mathcal{A} the mpk .
2. **Query Phase.** The adversary can make the following queries in an arbitrary order.
 - Extraction Query: \mathcal{A} sends an extraction query for an identity $\text{id} \in \mathcal{ID}$. The experiment returns usk_{id} to the adversary, where $\text{usk}_{\text{id}} \leftarrow \text{Extract}(\text{mpk}, \text{msk}, \text{id})$.
 - Challenge Query : \mathcal{A} sends a pair $(\text{id}_1^*, \text{id}_2^*) \in \mathcal{ID} \times \mathcal{ID}$ as the challenge query. The experiment samples a bit $\beta \leftarrow \{0, 1\}$ and returns $K^* \leftarrow \text{Share}(\text{mpk}, \text{usk}_{\text{id}_1^*}, \text{id}_2^*)$, where $\text{usk}_{\text{id}_1^*} \leftarrow \text{Extract}(\text{mpk}, \text{msk}, \text{id}_1^*)$, if $\beta = 0$ and returns $K^* \leftarrow \mathbb{G}_T$ if $\beta = 1$.
3. **Output Phase.** \mathcal{A} outputs a guess bit β' , and the experiment outputs 1 if $\beta = \beta'$.

We only quantify over \mathcal{A} that guarantees that does not make extraction queries for id_1^* and id_2^* .

6.3 MULTI-INPUT FE FOR PSEUDORANDOM FUNCTIONALITIES

In this section, we construct our main tool – multi-input functional encryption for pseudorandom functionalities.

6.3.1 Definition

In this section we give the definitions for multi-input functional encryption for pseudorandom functionalities (mi-prFE).

Consider a function family $\{\mathcal{F}_{\text{prm}} = \{f : (\mathcal{X}_{\text{prm}})^n \rightarrow \mathcal{Y}_{\text{prm}}\}\}_{\text{prm}}$, for a parameter $\text{prm} = \text{prm}(\lambda)$, where each \mathcal{F}_{prm} is a finite collection of n -ary functions. Each function $f \in \mathcal{F}_{\text{prm}}$ takes as input strings x_1, \dots, x_n , where each $x_i \in \mathcal{X}_{\text{prm}}$ and outputs $f(x_1, \dots, x_n) \in \mathcal{Y}_{\text{prm}}$.

Syntax. A miprfe scheme mi-prFE_n for n -ary function family \mathcal{F}_{prm} consists of polynomial time algorithms $(\text{Setup}, \text{KeyGen}, \text{Enc}_1, \dots, \text{Enc}_n, \text{Dec})$ defined as follows.

$\text{Setup}(1^\lambda, 1^n, \text{prm}) \rightarrow (\text{mpk}, \text{msk})$. The setup algorithm takes as input the security parameter λ , the function arity n and a parameter prm and outputs a master public key mpk and master secret key msk ⁸.

$\text{KeyGen}(\text{msk}, f) \rightarrow \text{sk}_f$. The key generation algorithm takes as input the master secret key msk and a function $f \in \mathcal{F}_{\text{prm}}$ and it outputs a functional secret key sk_f .

$\text{Enc}_i(\text{msk}, x) \rightarrow \text{ct}$. The encryption algorithm for the i -th slot takes as input the master secret key msk and an input $x \in \mathcal{X}_{\text{prm}}$ and outputs a ciphertext $\text{ct}_i \in \mathcal{CT}$, where \mathcal{CT} is the ciphertext space.

$\text{Dec}(\text{mpk}, \text{sk}_f, f, \text{ct}_1, \dots, \text{ct}_n) \rightarrow y$. The decryption algorithm takes as input the

⁸We assume w.l.o.g that msk includes mpk .

master public key mpk , secret key sk_f , function f and n ciphertexts $\text{ct}_1, \dots, \text{ct}_n$, and outputs $y \in \mathcal{Y}_{\text{prm}}$.

Definition 6.9 (Correctness). A mi-prFE scheme is said to be correct if for every prm , n -ary function $f \in \mathcal{F}_{\text{prm}}$ and input tuple $(x_1, \dots, x_n) \in \mathcal{X}_{\text{prm}}^n$ we have

$$\Pr \left[\begin{array}{l} (\text{mpk}, \text{msk}) \leftarrow \text{Setup}(1^\lambda, 1^n, \text{prm}), \text{sk}_f \leftarrow \text{KeyGen}(\text{msk}, f), \\ \text{Dec}(\text{mpk}, \text{sk}_f, f, \text{Enc}_1(\text{msk}, x_1), \dots, \text{Enc}_n(\text{msk}, x_n)) = f(x_1, \dots, x_n) \end{array} \right] \geq 1 - \text{negl}(\lambda). \quad (6.16)$$

Definition 6.10 (κ -Security). Let $\kappa = \kappa(\lambda)$ be a function in λ . For a mi-prFE scheme for function family $\{\mathcal{F}_{\text{prm}} = \{f : (\mathcal{X}_{\text{prm}})^n \rightarrow \mathcal{Y}_{\text{prm}}\}_{\text{prm}}, \text{parameter } \text{prm} = \text{prm}(\lambda), \text{ let Samp be a PPT algorithm that on input } 1^\lambda, \text{ outputs}$

$$\left(\{f_k\}_{k \in [q_0]}, \{x_1^{j_1}\}_{j_1 \in [q_1]}, \dots, \{x_n^{j_n}\}_{j_n \in [q_n]}, \text{aux} \in \{0, 1\}^* \right)$$

where q_0 is the number of key queries, q_i is the number of encryption queries for the i -th slot, $f_1, \dots, f_{q_0} \in \mathcal{F}_{\text{prm}}$ and $x_i^{j_i} \in \mathcal{X}_{\text{prm}}$ for all $i \in [n], j_i \in [q_i]$. We say that the mi-prFE scheme satisfies κ -security with respect to the sampler class \mathcal{SC} if for every PPT sampler $\text{Samp} \in \mathcal{SC}$ there exists a PPT simulator algorithm $\{\text{Sim}_i\}_{i \in [n]}$ such that

$$\begin{aligned} & \text{If } \left(1^\kappa, \{f_k, f_k(x_1^{j_1}, \dots, x_n^{j_n})\}_{k \in [q_0], j_1 \in [q_1], \dots, j_n \in [q_n]}, \text{aux} \right) \\ & \approx_c \left(1^\kappa, \{f_k, \Delta_{k, j_1, \dots, j_n}\}_{k \in [q_0], j_1 \in [q_1], \dots, j_n \in [q_n]}, \text{aux} \right) \end{aligned} \quad (6.17)$$

$$\begin{aligned} & \text{then } \left(\text{mpk}, \{f_k, \text{sk}_{f_k}\}_{k \in [q_0]}, \{\text{ct}_i^{j_i}\}_{i \in [n], j_i \in [q_i]}, \text{aux} \right) \\ & \approx_c \left(\text{mpk}, \{f_k, \text{sk}_{f_k}\}_{k \in [q_0]}, \{\delta_i^{j_i}\}_{i \in [n], j_i \in [q_i]}, \text{aux} \right), \end{aligned} \quad (6.18)$$

where $\kappa \geq \lambda^n$, $(\text{mpk}, \text{msk}) \leftarrow \text{Setup}(1^\lambda, 1^n, \text{prm}), \text{sk}_{f_k} \leftarrow \text{KeyGen}(\text{msk}, f_k), \text{ct}_i^{j_i} \leftarrow \text{Enc}(\text{msk}, x_i^{j_i}), \delta_i^{j_i} \leftarrow \text{Sim}_i(\text{msk})$, and $\Delta_{k, j_1, \dots, j_n} \leftarrow \mathcal{Y}_{\text{prm}}$ for $i \in [n], j_i \in [q_i]$, and $k \in [q_0]$.

Remark 32. Note that 1^κ in Equation (6.17) is introduced for the purpose of padding, allowing the distinguisher for the distributions to run in time polynomial in κ and requiring

the distinguishing advantage to be negligible in κ .⁹ The reason why we require $\kappa \geq \lambda^n$ is that the input length to the distinguisher is polynomial in λ^n anyway and in order for the padding to make sense, κ should satisfy this condition. If we need κ to be larger, this doubly strengthens the requirement for the precondition, as it means we want the distributions in Equation (6.17) to be indistinguishable against an adversary with a longer running time and smaller advantage. Ideally, we want κ to be as small as λ^n to make the requirement weaker. However, the security proof for our construction in Section 6.3.2 for general n requires large κ as an artifact of the proof technique. In the special case of n being constant, we can achieve $\kappa = \lambda^n$.

The following variant of security will be necessary in Section 6.6.

Definition 6.11 (Pseudorandomness of the Last Slot Ciphertext). We say that a mi-prFE scheme satisfies κ -pseudorandomness of the last slot ciphertext property if it satisfies κ -security as per defined in Definition 6.10 where the simulator Sim_n corresponding to the last slot ciphertext outputs a random string of the same length as $\text{ct}_n^{j_n}$.

Remark 33. It is shown in [19, 65] that there is no mi-prFE that satisfies the above style security for all general samplers. Therefore, when we use the security of mi-prFE, we invoke the security *with respect to a specific sampler class* that is induced by the respective applications. For simplicity, we sometimes will treat as if there was mi-prFE that is secure for all the samplers.

6.3.2 Construction for n-input prFE

In this section we provide our construction of a multi-input functional encryption scheme for pseudorandom functionalities for function family $\mathcal{F}_{nL(\lambda),d(\lambda)} = \{f : \{\{0, 1\}^L\}^n \rightarrow \{0, 1\}\}$, where the depth of a function $f \in \mathcal{F}$ is at most $d(\lambda) = \text{poly}(\lambda)$. Each function $f \in \mathcal{F}$ takes as input strings $x_1, \dots, x_n \in \{0, 1\}^L$ and outputs $f(x_1, \dots, x_n) \in \{0, 1\}$. We consider the case of arity n being constant and the general case of n being arbitrary polynomial in λ . While we provide separate security proofs for these cases, we have

⁹This is due to our convention, where the running time of the distinguisher should be polynomial in its input length and the distinguishing advantage should be negligible in its input length. Please refer to Section 6.2 for the details.

unified description of the construction. The reason why we consider the proofs separately is that we can base the security of the scheme on a weaker assumption when n is constant than the general case.

Building Blocks. Our construction uses the following building blocks.

1. A secret key encryption scheme $\text{SKE} = (\text{SKE.Setup}, \text{SKE.Enc}, \text{SKE.Dec})$. When we set up the scheme SKE , we run it on a scaled version of the security parameter Λ , instead of the usual security parameter λ . We will explain how to set Λ in Remark 34. We denote the ciphertext space of the scheme by $\mathcal{CT}_{\text{SKE}, \Lambda}$ and the key space of the scheme by $\mathcal{K}_{\text{SKE}, \Lambda}$. In the following, we drop Λ and denote them as $\mathcal{CT}_{\text{SKE}}$ and \mathcal{K}_{SKE} respectively.
2. n single-input FE scheme for pseudorandom functionality $\text{prFE}_1, \dots, \text{prFE}_n$. For $i \in [n]$, $\text{prFE}_i = (\text{prFE}_i.\text{Setup}, \text{prFE}_i.\text{KeyGen}, \text{prFE}_i.\text{Enc}, \text{prFE}_i.\text{Dec})$ for circuit class $\mathcal{C}_{\text{inp}_i(\lambda), \text{dep}_i(\lambda), \text{out}_i(\lambda)}$ consisting of circuits with input length $\text{inp}_i(\lambda)$, maximum depth $\text{dep}_i(\lambda)$ and output length $\text{out}_i(\lambda)$. We denote the ciphertext space of the prFE_i scheme by $\mathcal{CT}_{\text{prFE}_i}$. For our construction, we set the following parameters

- $\text{inp}_1 = n \cdot L$, $\text{dep}_1 = d$, and $\text{out}_1 = 1$.
- $\text{inp}_i = |\text{SKE.key}| + (n-i)L + n\Lambda$, $\text{dep}_i = \text{poly}(d, \lambda)$, and $\text{out}_i = |\text{prFE}_{i-1}.\text{ct}|$ for $i \in [2, n]$, where $\text{SKE.key} \in \mathcal{CT}_{\text{SKE}}$ and $\text{prFE}_{i-1}.\text{ct} \in \mathcal{CT}_{\text{prFE}_{i-1}}$.

3. We also use $n - 1$ pseudorandom functions $\text{PRF}_1, \dots, \text{PRF}_{n-1}$. Similarly to the case of SKE , we use Λ to setup these instances of PRF . We specify the domain and codomain of the functions as $\text{PRF}_i : \{0, 1\}^\Lambda \times \{0, 1\}^\Lambda \rightarrow \{0, 1\}^{\text{len}_i}$ where len_i is the length of randomness used in $\text{prFE}_i.\text{Enc}$ for $i \in [n - 1]$.

We describe our construction of $\text{mi-prFE} = (\text{Setup}, \text{KeyGen}, \text{Enc}_1, \dots, \text{Enc}_n, \text{Dec})$ in the following.

$\text{Setup}(1^\lambda, 1^n, \text{prm}) \rightarrow (\text{mpk}, \text{msk})$. The setup algorithm does the following.

- Generate $(\text{prFE}_i.\text{mpk}, \text{prFE}_i.\text{msk}) \leftarrow \text{prFE}_i.\text{Setup}(1^\lambda, 1^{\text{prm}_i})$ for all $i \in [n]$,
- Generate $\text{SKE.sk} \leftarrow \text{SKE.Setup}(1^\Lambda)$.
- Output $\text{mpk} := (\{\text{prFE}_i.\text{mpk}\}_{i \in [n]})$ and $\text{msk} := (\text{SKE.sk}, \{\text{prFE}_i.\text{msk}, \text{prFE}_i.\text{mpk}\}_{i \in [n]})$.

$\text{KeyGen}(\text{msk}, f) \rightarrow \text{sk}_f$. The key generation algorithm does the following.

- Parse $\text{msk} = (\text{SKE.sk}, \{\text{prFE}_i.\text{msk}, \text{prFE}_i.\text{mpk}\}_{i \in [n]})$.
- Compute $\text{prFE}_1.\text{sk}_f \leftarrow \text{prFE}_1.\text{KeyGen}(\text{prFE}_1.\text{msk}, f)$.
- Output $\text{sk}_f := \text{prFE}_1.\text{sk}_f$.

$\text{Enc}_i(\text{msk}, \mathbf{x}_i) \rightarrow \text{ct}_i$. For $i \in [n - 1]$, the Enc_i algorithm outputs a function secret key corresponding to prFE_{i+1} -th instance in the following way.

- Parse $\text{msk} = (\text{SKE.sk}, \{\text{prFE}_i.\text{msk}, \text{prFE}_i.\text{mpk}\}_{i \in [n]})$.
- Sample $\mathbf{r}_i \leftarrow \{0, 1\}^\Lambda$.
- Compute $\text{SKE.ct}_i \leftarrow \text{SKE.Enc}(\text{SKE.sk}, \mathbf{x}_i)$.
- Define $F_i := F_i[\text{SKE.ct}_i, \mathbf{r}_i, \text{prFE}_i.\text{mpk}]$ as in Figure 6.1.¹⁰
- Compute $\text{prFE}_{i+1}.\text{sk} \leftarrow \text{prFE}_{i+1}.\text{KeyGen}(\text{prFE}_{i+1}.\text{msk}, F_i)$.
- Output $\text{ct}_i := \text{prFE}_{i+1}.\text{sk}$.

$\text{Enc}_n(\text{msk}, \mathbf{x}_n) \rightarrow \text{ct}_n$. The Enc_n algorithm does the following.

- Parse $\text{msk} = (\text{SKE.sk}, \{\text{prFE}_i.\text{msk}, \text{prFE}_i.\text{mpk}\}_{i \in [n]})$.
- For $i \in [n - 1]$, sample $K_i \leftarrow \{0, 1\}^\Lambda$.
- Compute $\text{prFE}_n.\text{ct} \leftarrow \text{prFE}_n.\text{Enc}(\text{prFE}_n.\text{mpk}, (\text{SKE.sk}, \mathbf{x}_n, K_1, \dots, K_{n-1}))$.
- Output $\text{ct}_n := \text{prFE}_n.\text{ct}$.

$\text{Dec}(\text{mpk}, \text{sk}_f, f, \text{ct}_1, \dots, \text{ct}_n) \rightarrow y \in \{0, 1\}$. The decryption algorithm does the following.

- Parse $\text{mpk} = (\{\text{prFE}_i.\text{mpk}\}_{i \in [n]})$, $\text{sk}_f = \text{prFE}_1.\text{sk}_f$, $\text{ct}_i = \text{prFE}_{i+1}.\text{sk}$ for $i \in [n - 1]$, and $\text{ct}_n = \text{prFE}_n.\text{ct}$.

¹⁰The hardwired values are not hidden, even if we don't output them explicitly.

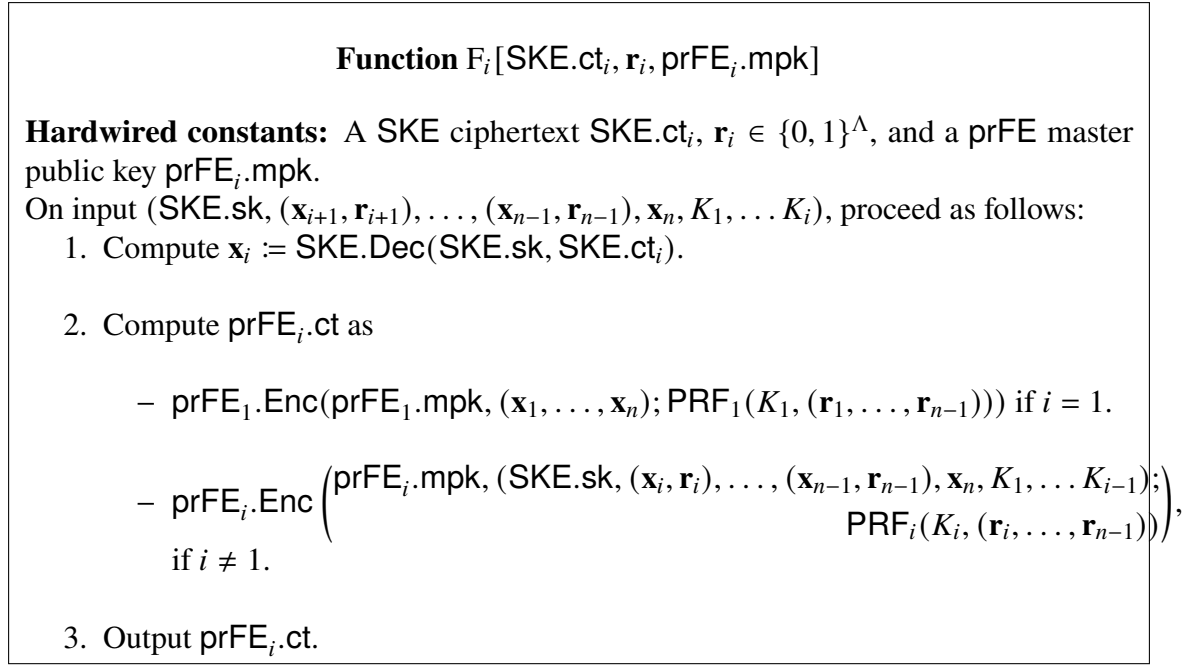


Figure 6.1: Function F_i

- For $i = n, \dots, 2$ and do the following.
 1. Compute $\text{prFE}_{i-1}.\text{ct} := \text{prFE}_i.\text{Dec}(\text{prFE}_i.\text{mpk}, \text{prFE}_i.\text{sk}, F_{i-1}, \text{prFE}_i.\text{ct})$.
 2. If $i = 2$ output $\text{prFE}_1.\text{ct}$, else set $i := i - 1$ and go to Step 1.
- Output $y := \text{prFE}_1.\text{Dec}(\text{prFE}_1.\text{mpk}, \text{prFE}_1.\text{sk}_f, f, \text{prFE}_1.\text{ct})$.

Remark 34. We consider two cases of parameter settings for the construction. One is the case of n being constant. In this case, we simply set $\Lambda = \lambda$. In the general case of $n = \text{poly}(\lambda)$, we do something more complex. In this case, we assume that PRF and SKE have subexponential security. This means that there exists $0 < \delta < 1$ such that there is no adversary with size 2^{λ^δ} and distinguishing advantage $2^{-\lambda^\delta}$ against SKE and PRF for all sufficiently large λ . In the security proof, we require PRF and SKE to be secure even against an adversary that takes 1^κ as an input and thus runs in polynomial time in κ . To satisfy this requirement, we run SKE and PRF with respect to a larger security parameter Λ that satisfies $2^{\Lambda^\delta} \geq \kappa^{\omega(1)}$. An example choice would be to take $\Lambda := (n^2\lambda)^{1/\delta}$.

Efficiency. The scheme has the following parameters : $|\text{mpk}| = \text{poly}(n, L, d, \Lambda, \lambda)$, $|\text{sk}_f| = \text{poly}(d, \lambda)$, $|\text{ct}_1| = nL\text{poly}(\text{dep}, \lambda)$, $|\text{ct}_i| = \text{poly}(n, L, d, \Lambda, \lambda)$ for $i \in [2, n]$.

Correctness. We prove the correctness of our scheme via the following theorem.

Theorem 6.19. *Suppose prFE_i for $i \in [n]$ and SKE are correct, then the above construction of mi-prFE satisfies correctness as defined in Definition 6.9.*

Proof. To prove the theorem, we first prove the following statement.

Claim 6.20. For $i = n, \dots, 2$, we have

$$\Pr[\text{prFE}_i.\text{Dec}(\text{prFE}_i.\text{mpk}, \text{prFE}_i.\text{sk}, F_{i-1}, \text{prFE}_i.\text{ct}) = \text{prFE}_{i-1}.\text{ct}] = 1 \quad (6.19)$$

where

$$\text{prFE}_{i-1}.\text{ct} = \begin{cases} \text{prFE}_{i-1}.\text{Enc} \left(\text{prFE}_{i-1}.\text{mpk}, (\text{SKE}.\text{sk}, (\mathbf{x}_{i-1}, \mathbf{r}_{i-1}), \dots, (\mathbf{x}_{n-1}, \mathbf{r}_{n-1}), \mathbf{x}_n, K_1, \dots, K_{i-2}) \right. \\ \left. ; \text{PRF}_{i-1}(K_{i-1}, (\mathbf{r}_{i-1}, \dots, \mathbf{r}_{n-1})) \right) & \text{if } i \neq 2 \\ \text{prFE}_1.\text{Enc}(\text{prFE}_1.\text{mpk}, (\mathbf{x}_1, \dots, \mathbf{x}_n); \text{PRF}_1(K_1, (\mathbf{r}_1, \dots, \mathbf{r}_{n-1}))) & \text{if } i = 2. \end{cases}$$

Proof. We prove this by induction.

Base Case: For $i = n$, we show that

$$\Pr[\text{prFE}_n.\text{Dec}(\text{prFE}_n.\text{mpk}, \text{prFE}_n.\text{sk}, F_{n-1}, \text{prFE}_n.\text{ct}) = \text{prFE}_{n-1}.\text{ct}] = 1.$$

From the correctness of prFE_n scheme, we have with probability 1

$$\begin{aligned} & \text{prFE}_n.\text{Dec}(\text{prFE}_n.\text{mpk}, \text{prFE}_n.\text{sk}, F_{n-1}, \text{prFE}_n.\text{ct}) \\ &= F_{n-1}[\text{SKE}.\text{ct}_{n-1}, \mathbf{r}_{n-1}, \text{prFE}_{n-1}.\text{mpk}](\text{SKE}.\text{sk}, \mathbf{x}_n, K_1, \dots, K_{n-1}). \end{aligned}$$

Next, by the definition of F_{n-1} and the correctness of the SKE scheme, we have

$$\text{prFE}_n.\text{Dec}(\text{prFE}_n.\text{mpk}, \text{prFE}_n.\text{sk}, F_{n-1}, \text{prFE}_n.\text{ct})$$

$$= \text{prFE}_{n-1}.\text{Enc} \left(\text{prFE}_{n-1}.\text{mpk}, (\text{SKE}.\text{sk}, (\mathbf{x}_{n-1}, \mathbf{r}_{n-1}), \mathbf{x}_n, K_1, \dots, K_{n-2}); \right. \\ \left. \text{PRF}_{n-1}(K_{n-1}, \mathbf{r}_{n-1}) \right) \quad (6.20)$$

which proves the base case.

Inductive Step: For the inductive step, suppose Equation (6.19) holds for some $i \in [3, n]$ then we prove the same statement for $i - 1$. Consider

$$\begin{aligned} & \text{prFE}_{i-1}.\text{Dec}(\text{prFE}_{i-1}.\text{mpk}, \text{prFE}_{i-1}.\text{sk}, F_{i-2}, \text{prFE}_{i-1}.\text{ct}) \\ &= F_{i-2}[\text{SKE}.\text{ct}_{i-2}, \mathbf{r}_{i-2}, \text{prFE}_{i-2}.\text{mpk}](\text{SKE}.\text{sk}, (\mathbf{x}_{i-1}, \mathbf{r}_{i-1}), \dots, (\mathbf{x}_{n-1}, \mathbf{r}_{n-1}), \mathbf{x}_n, K_1, \dots, K_{i-2}) \\ &= \begin{cases} \text{prFE}_{i-2}.\text{Enc} \left(\text{prFE}_{i-2}.\text{mpk}, (\text{SKE}.\text{sk}, (\mathbf{x}_{i-2}, \mathbf{r}_{i-2}), \dots, \mathbf{x}_n, K_1, \dots, K_{i-3}); \right. \\ \quad \left. \text{PRF}_{i-2}(K_{i-2}, (\mathbf{r}_{i-2}, \dots, \mathbf{r}_{n-1})) \right) & \text{if } i \neq 3 \\ \text{prFE}_1.\text{Enc}(\text{prFE}_1.\text{mpk}, (\mathbf{x}_1, \dots, \mathbf{x}_n); \text{PRF}_1(K_1, (\mathbf{r}_1, \dots, \mathbf{r}_{n-1}))) & \text{if } i = 3. \end{cases} \end{aligned} \quad (6.21)$$

where in the first equality we use $\text{prFE}_{i-1}.\text{sk} = \text{prFE}_{i-1}.\text{KeyGen}(\text{prFE}_{i-1}.\text{msk}, F_{i-2})$ and $\text{prFE}_{i-1}.\text{ct} = \text{prFE}_{i-1}.\text{Enc}(\text{prFE}_{i-1}.\text{mpk}, (\text{SKE}.\text{sk}, (\mathbf{x}_{i-1}, \mathbf{r}_{i-1}), \dots, (\mathbf{x}_{n-1}, \mathbf{r}_{n-1}), \mathbf{x}_n, K_1, \dots, K_{i-2}); \text{PRF}_{i-1}(K_{i-1}, (\mathbf{r}_{i-1}, \dots, \mathbf{r}_{n-1})))$ which follows from the assumption for i . The second equality follows from the definition of F_{i-2} and the correctness of the SKE scheme.

This completes the proof of the inductive step. ■

Using the above claim we get $\text{prFE}.\text{ct}_1 = \text{prFE}_1.\text{Enc}(\text{prFE}_1.\text{mpk}, (\mathbf{x}_1, \dots, \mathbf{x}_n); \text{PRF}_1(K_1, (\mathbf{r}_1, \dots, \mathbf{r}_{n-1})))$ from Step 6.3.2 of the decryption algorithm with probability 1. From the correctness of prFE_1 scheme, the decryption Step 6.3.2 outputs

$$\begin{aligned} y &= \text{prFE}_1.\text{Dec}(\text{prFE}_1.\text{mpk}, \text{prFE}_1.\text{sk}_f, f, \text{prFE}_1.\text{ct}) \\ &= \text{prFE}_1.\text{Dec} \left(\text{prFE}_1.\text{mpk}, \text{prFE}_1.\text{sk}_f, f, \text{prFE}_1.\text{Enc}(\text{prFE}_1.\text{mpk}, (\mathbf{x}_1, \dots, \mathbf{x}_n); \right. \\ & \quad \left. \text{PRF}_1(K_1, (\mathbf{r}_1, \dots, \mathbf{r}_{n-1}))) \right) \end{aligned}$$

$$=f(\mathbf{x}_1, \dots, \mathbf{x}_n)$$

with probability 1. ■

6.3.3 Security Proof for General n

Theorem 6.21. *Let $\mathcal{SC}_{\text{mi-prFE}}$ be a sampler class for mi-prFE. Suppose prFE_i scheme satisfies non-uniform κ -prCT security as per Definition 6.3 for $\kappa = \lambda^{n^2 \log \lambda}$ with respect to the sampler class that contains all $\text{Samp}_{\text{prFE}}(1^\lambda)$, induced by $\text{Samp}_{\text{mi-prFE}} \in \mathcal{SC}_{\text{mi-prFE}}$, as in Section 6.3.3, SKE satisfies sub-exponential INDr security and PRF_i is sub-exponentially secure, then mi-prFE constructed above satisfies security for $\kappa = \lambda^{n^2 \log \lambda}$ as in Definition 6.11. Note that this in particular implies the κ -security defined in Definition 6.10.*

Proof. Consider a sampler $\text{Samp}_{\text{mi-prFE}}$ that generates the following:

1. **Key Queries.** It issues q_0 number of functions f_1, \dots, f_{q_0} for key queries.
2. **Ciphertext Queries.** It issues q_i number of messages for ciphertext queries for slot i . We use $\mathbf{x}_i^{j_i}$ to denote the j_i -th ciphertext query corresponding to the i -th slot, where $j_i \in [q_i]$ and $i \in [n]$.
3. **Auxiliary Information.** It outputs the auxiliary information $\text{aux}_{\mathcal{A}}$.

To prove the security of mi-prFE as per Definition 6.11, we first define $\{\text{Sim}_i\}_{i \in [n]}$ as follows. Observe that Sim_n outputs random string as is required by Definition 6.11.

$\text{Sim}_i(\text{msk}) \rightarrow \text{ct}_i$ for $i \in [n - 1]$.

- Parse $\text{msk} = (\text{SKE.sk}, \{\text{prFE}_i.\text{msk}, \text{prFE}_i.\text{mpk}\}_{i \in [n]})$.
- Sample $\mathbf{r}_i \leftarrow \{0, 1\}^\Lambda$ and $\gamma_i \leftarrow \mathcal{CT}_{\text{SKE}}$.
- Compute $\text{prFE}_{i+1}.\text{sk} \leftarrow \text{prFE}_{i+1}.\text{KeyGen}(\text{prFE}_{i+1}.\text{msk}, F_i[\gamma_i, \mathbf{r}_i, \text{prFE}_i.\text{mpk}])$.
- Output $\text{ct}_i := \text{prFE}_{i+1}.\text{sk}$.

$\text{Sim}_n(\text{msk}) \rightarrow \text{ct}_n$. Sample $\delta_n \leftarrow C\mathcal{T}_{\text{prFE}_n}$ and output $\text{ct}_n := \delta_n$.
Then, it suffices to show

$$\begin{aligned} & \left(\begin{array}{c} \text{aux}_{\mathcal{A}}, \{\text{prFE}_i.\text{mpk}\}_{i \in [n]}, \{\text{prFE}_n.\text{ct}^{j_n}\}_{j_n \in [q_n]}, \\ \{f_k, \text{sk}_{f_k} = \text{prFE}_1.\text{sk}_{f_k}\}_{k \in [q_0]}, \{\text{SKE.ct}_i^{j_i}, \mathbf{r}_i^{j_i}, \text{prFE}_{i+1}.\text{sk}^{j_i}\}_{i \in [n-1], j_i \in [q_i]}, \end{array} \right) \\ & \approx_c \left(\begin{array}{c} \text{aux}_{\mathcal{A}}, \{\text{prFE}_i.\text{mpk}\}_{i \in [n]}, \{\delta_n^{j_n}\}_{j_n \in [q_n]} \\ \{f_k, \text{sk}_{f_k} = \text{prFE}_1.\text{sk}_{f_k}\}_{k \in [q_0]}, \{\gamma_i^{j_i}, \mathbf{r}_i^{j_i}, \text{prFE}_{i+1}.\text{sk}^{j_i}\}_{i \in [n-1], j_i \in [q_i]}, \end{array} \right) \quad (6.22) \end{aligned}$$

where $(\text{aux}_{\mathcal{A}}, \{f_k\}_k, \{\mathbf{x}_i^{j_i}\}_{i, j_i}) \leftarrow \text{Samp}_{\text{mi-prFE}}(1^\lambda)$,

$(\text{mpk} = \{\text{prFE}_i.\text{mpk}\}_i, \text{msk} = (\text{SKE.sk}, \{\text{prFE}_i.\text{msk}, \text{prFE}_i.\text{mpk}\}_i) \leftarrow \text{Setup}(1^\lambda, 1^n, \text{prm})$,

$\text{prFE}_1.\text{sk}_{f_k} \leftarrow \text{prFE}_1.\text{KeyGen}(\text{prFE}_1.\text{msk}, f_k) \quad \text{for } k \in [q_0]$,

$\text{SKE.ct}_i^{j_i} \leftarrow \text{SKE.Enc}(\text{SKE.sk}, \mathbf{x}_i^{j_i}), \quad \gamma_i^{j_i} \leftarrow C\mathcal{T}_{\text{SKE}}, \quad \mathbf{r}_i^{j_i} \leftarrow \{0, 1\}^\Lambda$,

$\text{prFE}_{i+1}.\text{sk}^{j_i} \leftarrow \begin{cases} \text{prFE}_{i+1}.\text{KeyGen}(\text{prFE}_{i+1}.\text{msk}, F_i[\text{SKE.ct}_i^{j_i}, \mathbf{r}_i^{j_i}, \text{prFE}_i.\text{mpk}]) & \text{in LHS of Eq. (6.22)} \\ \text{prFE}_{i+1}.\text{KeyGen}(\text{prFE}_{i+1}.\text{msk}, F_i[\gamma_i^{j_i}, \mathbf{r}_i^{j_i}, \text{prFE}_i.\text{mpk}]) & \text{in RHS of Eq. (6.22)} \end{cases}$

$\text{prFE}_n.\text{ct}^{j_n} \leftarrow \text{prFE}_n.\text{Enc}(\text{prFE}_n.\text{mpk}, (\text{SKE.sk}, \mathbf{x}_n^{j_n}, K_1^{j_n}, \dots, K_{n-1}^{j_n})), \quad \delta_n^{j_n} \leftarrow C\mathcal{T}_{\text{prFE}_n}$,

$K_i^{j_n} \leftarrow \{0, 1\}^\Lambda, \quad \text{for } i \in [n-1], j_i \in [q_i], \text{ and } j_n \in [q_n]$

assuming we have

$$\begin{aligned} & \left(1^\kappa, \text{aux}_{\mathcal{A}}, \left\{ f_k, f_k(\mathbf{x}_1^{j_1}, \dots, \mathbf{x}_n^{j_n}) \right\}_{k \in [q_0], j_1 \in [q_1], \dots, j_n \in [q_n]} \right) \\ & \approx_c \left(1^\kappa, \text{aux}_{\mathcal{A}}, \left\{ f_k, \Delta_k^{j_1, \dots, j_n} \right\}_{k \in [q_0], j_1 \in [q_1], \dots, j_n \in [q_n]} \right) \quad (6.23) \end{aligned}$$

where $(\text{aux}_{\mathcal{A}}, \{f_k\}_k, \{\mathbf{x}_i^{j_i}\}_{i, j_i}) \leftarrow \text{Samp}_{\text{mi-prFE}}(1^\lambda)$, and $\Delta_k^{j_1, \dots, j_n} \leftarrow \{0, 1\}$. We prove this in the following two steps.

- **Step 1.** We first show that Equation (6.23) implies

$$\left(\begin{array}{c} 1^\kappa, \quad \text{aux}_{\mathcal{A}}, \quad \text{prFE.mpk}_1 \\ \{\text{SKE.ct}_i^{j_i}\}_{i \in [n-1], j_i \in [q_i]} \quad \{\text{prFE}_1.\text{ct}^j\}_{j \in [q_1] \times \dots \times [q_n]} \\ \{f_k, \text{sk}_{f_k} = \text{prFE}_1.\text{sk}_{f_k}\}_{k \in [q_0]} \end{array} \right)$$

$$\approx_c \left(\begin{array}{c} 1^\kappa, \text{aux}_{\mathcal{A}}, \text{prFE.mpk}_1 \\ \{\text{SKE.ct}_i^{j_i}\}_{i \in [n-1], j_i \in [q_{\text{msg}}]}, \{\Delta^{\mathbf{j}}\}_{\mathbf{j} \in [q_1] \times \dots \times [q_n]} \\ \{f_k, \text{sk}_{f_k} = \text{prFE}_1.\text{sk}_{f_k}\}_{k \in [q_0]} \end{array} \right) \quad (6.24)$$

where $\mathbf{j} = (j_1, \dots, j_n) \in [q_1] \times \dots \times [q_n]$, $\Delta^{\mathbf{j}} \leftarrow \mathcal{CT}_{\text{prFE}_1}$, and

$$\text{prFE}_1.\text{ct}^{\mathbf{j}} \leftarrow \text{prFE}_1.\text{Enc}(\text{prFE}_1.\text{mpk}, (\mathbf{x}_1^{j_1}, \dots, \mathbf{x}_n^{j_n})).$$

- **Step 2.** We prove that Equation (6.24) implies Equation (6.22).

Step 1. We show the following lemma.

Lemma 6.22. *If SKE satisfies subexponential IND_r security, Equation (6.23) implies Equation (6.24).*

Proof. We first prove the following:

$$\begin{aligned} & \left(1^\kappa, \text{aux}_{\mathcal{A}}, \left\{ f_k, f_k(\mathbf{x}_1^{j_1}, \dots, \mathbf{x}_n^{j_n}) \right\}_{k \in [q_0], j_1 \in [q_1], \dots, j_n \in [q_n]}, \left\{ \text{SKE.ct}_i^{j_i} \leftarrow \text{SKE.Enc}(\text{SKE.sk}, \mathbf{x}_i^{j_i}) \right\}_{i \in [n-1], j_i \in [q_i]} \right) \\ & \approx_c \left(1^\kappa, \text{aux}_{\mathcal{A}}, \left\{ f_k, \Delta_k^{\mathbf{j}_1, \dots, \mathbf{j}_n} \right\}_{k \in [q_0], j_1 \in [q_1], \dots, j_n \in [q_n]}, \left\{ \text{SKE.ct}_i^{j_i} \leftarrow \text{SKE.Enc}(\text{SKE.sk}, \mathbf{x}_i^{j_i}) \right\}_{i \in [n-1], j_i \in [q_i]} \right), \end{aligned} \quad (6.25)$$

where $(\text{aux}_{\mathcal{A}}, \{f_k\}_k, \{\mathbf{x}_i^{j_i}\}_{i,j_i}) \leftarrow \text{Samp}(1^\lambda)$, $\text{SKE.sk} \leftarrow \text{SKE.Setup}(1^\Lambda)$, and $\Delta_k^{\mathbf{j}_1, \dots, \mathbf{j}_n} \leftarrow \{0, 1\}$. To prove this, we observe

$$\begin{aligned} & \left(1^\kappa, \text{aux}_{\mathcal{A}}, \left\{ f_k, f_k(\mathbf{x}_1^{j_1}, \dots, \mathbf{x}_n^{j_n}) \right\}_{k, j_1, \dots, j_n}, \left\{ \text{SKE.ct}_i^{j_i} \leftarrow \text{SKE.Enc}(\text{SKE.sk}, \mathbf{x}_i^{j_i}) \right\}_{i, j_i} \right) \\ & \approx_c \left(1^\kappa, \text{aux}_{\mathcal{A}}, \left\{ f_k, f_k(\mathbf{x}_1^{j_1}, \dots, \mathbf{x}_n^{j_n}) \right\}_{k, j_1, \dots, j_n}, \left\{ \gamma_i^{j_i} \leftarrow \mathcal{CT}_{\text{SKE}} \right\}_{i, j_i} \right) \end{aligned} \quad (6.26)$$

$$\approx_c \left(1^\kappa, \text{aux}_{\mathcal{A}}, \left\{ f_k, \Delta_k^{\mathbf{j}_1, \dots, \mathbf{j}_n} \right\}_{k, j_1, \dots, j_n}, \left\{ \gamma_i^{j_i} \leftarrow \mathcal{CT}_{\text{SKE}} \right\}_{i, j_i} \right) \quad (6.27)$$

$$\approx_c \left(1^\kappa, \text{aux}_{\mathcal{A}}, \left\{ f_k, \Delta_k^{\mathbf{j}_1, \dots, \mathbf{j}_n} \right\}_{k, j_1, \dots, j_n}, \left\{ \text{SKE.ct}_i^{j_i} \leftarrow \text{SKE.Enc}(\text{SKE.sk}, \mathbf{x}_i^{j_i}) \right\}_{i, j_i} \right). \quad (6.28)$$

Here, we justify each step of the equations above. We can see that Equation (6.26) follows from subexponential IND_r security of SKE, since SKE.sk is used only for computing $\{\text{SKE.ct}_i^{j_i}\}_{i, j_i}$ and not used anywhere else. Note that by our choice of the parameter $2^{\Lambda^\delta} > \kappa^{\omega(1)}$, we can use the security of SKE even for an adversary who runs in time

polynomial in κ . We can see that Equation (6.27) follows from Equation (6.23) by noting that adding random strings does not make the task of disinguishing the two distributions any easier. Finally, Equation (6.28) follows from IND_r security of SKE again.

We then consider a sampler Samp_1 that on input 1^κ outputs

$$\left(f_1, \dots, f_{q_0}, \{(\mathbf{x}_1^{j_1}, \dots, \mathbf{x}_n^{j_n})\}_{j_1 \in [q_1], \dots, j_n \in [q_n]}, \text{aux}_1 \stackrel{\text{def}}{=} \left(1^\kappa, \text{aux}_{\mathcal{A}}, \{\text{SKE.ct}_i^{j_i}\}_{i \in [n-1], j_i \in [q_i]} \right) \right).$$

By the security guarantee of prFE_1 with sampler Samp_1 and Equation (6.25), we obtain Equation (6.24). \blacksquare

Step 2. To prove that Equation (6.24) implies Equation (6.22), we prove the following statement.

Lemma 6.23. *For $h \in [n]$ and an adversary \mathcal{A} , let us consider the following distinguishing advantage:*

$$\begin{aligned} \text{Adv}_{\mathcal{A}}^h(\lambda) \stackrel{\text{def}}{=} & \left| \mathcal{A} \left(\begin{array}{l} \text{aux}_{\mathcal{A}}, \{\text{prFE.mpk}_i\}_{i \in [h]}, \{\text{SKE.ct}_i^{j_i}, \mathbf{r}_i^{j_i}\}_{i \in [n-1], j_i \in [q_i]} \\ \{f_k, \text{sk}_{f_k} = \text{prFE}_1.\text{sk}_{f_k}\}_{k \in [q_0]}, \\ \{\text{prFE}_{i+1}.\text{sk}^{j_i}\}_{i \in [h-1], j_i \in [q_i]}, \{\text{prFE}_h.\text{ct}^{\mathbf{j}}\}_{\mathbf{j} \in [q_h] \times \dots \times [q_n]}, \end{array} \right) \right. \\ & \left. - \mathcal{A} \left(\begin{array}{l} \text{aux}_{\mathcal{A}}, \{\text{prFE.mpk}_i\}_{i \in [h]}, \{\text{SKE.ct}_i^{j_i}, \mathbf{r}_i^{j_i}\}_{i \in [n-1], j_i \in [q_i]} \\ \{f_k, \text{sk}_{f_k} = \text{prFE}_1.\text{sk}_{f_k}\}_{k \in [q_0]}, \\ \{\text{prFE}_{i+1}.\text{sk}^{j_i}\}_{i \in [h-1], j_i \in [q_i]}, \{\Delta^{\mathbf{j}}\}_{\mathbf{j} \in [q_h] \times \dots \times [q_n]} \end{array} \right) \right| \quad (6.29) \end{aligned}$$

where $\mathbf{j} = (j_h, \dots, j_n)$, $\Delta^{\mathbf{j}} \leftarrow \text{CT}_{\text{prFE}_h}$,

$\text{prFE}_{i+1}.\text{sk}^{j_i} \leftarrow \text{prFE}_{i+1}.\text{KeyGen}(\text{prFE}_{i+1}.\text{msk}, F_i[\text{SKE.ct}_i^{j_i}, \mathbf{r}_i, \text{prFE}_i.\text{mpk}])$, and

$$\text{prFE}_h.\text{ct}^{\mathbf{j}} \leftarrow \text{prFE}_h.\text{Enc} \left(\text{prFE}_h.\text{mpk}, \left(\text{SKE.sk}, (\mathbf{x}_h^{j_h}, \mathbf{r}_h^{j_h}), \dots, (\mathbf{x}_{n-1}^{j_{n-1}}, \mathbf{r}_{n-1}^{j_{n-1}}), (\mathbf{x}_n^{j_n}, K_1^{j_n}, \dots, K_{h-1}^{j_{n-1}}) \right) \right).$$

Then, for every $h^* := \{h_\lambda^* \in [2, n(\lambda)]\}_\lambda$ and every non-uniform adversary $\mathcal{A} = \{\mathcal{A}_\lambda\}_\lambda$ such that $\text{Size}(\mathcal{A}) < \text{poly}(\kappa)$,¹¹ there exists another non-uniform adversary $\mathcal{B} = \{\mathcal{B}_\lambda\}_\lambda$

¹¹Here, we deviate from our convention that the adversary runs in polynomial time in its input length. Note that κ here may be super-polynomial in the input length to \mathcal{A} .

and a polynomial Q such that

$$\text{Adv}_{\mathcal{B}}^{h^*-1}(\lambda) \geq \text{Adv}_{\mathcal{A}}^{h^*}(\lambda)/Q(\lambda') - \text{negl}(\kappa) \quad \text{and} \quad \text{Size}(\mathcal{B}) \leq Q(\lambda') \cdot \text{Size}(\mathcal{A})$$

assuming the security of prFE as per Definition 6.3 with respect to κ and the subexponential security of PRF , where $\lambda' := \lambda^n$.

Proof. We invoke the security of prFE_{h^*} with non-uniform sampler Samp_{h^*} that takes as input the security parameter 1^λ and outputs

$$\left(\begin{array}{ll} \text{Functions:} & \left\{ F_{h^*-1}^{j_{h^*-1}} = F_{h^*-1}^{j_{h^*-1}}[\text{SKE.ct}_{h^*-1}^{j_{h^*-1}}, \mathbf{r}_{h^*-1}^{j_{h^*-1}}, \text{prFE}_{h^*-1}.\text{mpk}] \right\}_{j_{h^*-1} \in [q_{h^*-1}]}, \\ \text{Inputs:} & \left\{ \mathbf{x}^{j_{h^*}, \dots, j_n} \stackrel{\text{def}}{=} \left(\text{SKE.sk}, (\mathbf{x}_{h^*}^{j_{h^*}}, \mathbf{r}_{h^*}^{j_{h^*}}), \dots, (\mathbf{x}_n^{j_n}, K_1^{j_n}, \dots, K_{h^*-1}^{j_n}) \right) \right\}_{j_{h^*} \in [q_{h^*}], \dots, j_n \in [q_n]}, \\ \text{Auxiliary Information:} & \text{aux}_{h^*} \stackrel{\text{def}}{=} \left(\text{aux}_{\mathcal{A}}, \{\text{prFE}_i.\text{mpk}\}_{i \in [h^*-1]}, \{f_k, \text{sk}_{f_k} = \text{prFE}_1.\text{sk}_{f_k}\}_{k \in [q_0]}, \right. \\ & \left. \{\text{SKE.ct}_i^{j_i}, \mathbf{r}_i^{j_i}\}_{i \in [n-1], j_i \in [q_i]}, \{\text{prFE}_{i+1}.\text{sk}^{j_i}\}_{i \in [h^*-2], j_i \in [q_i]} \right) \end{array} \right).$$

We can see that the size of Samp is $\text{poly}(\lambda^n) = \text{poly}(\lambda')$, since $q_0 q_1 \dots q_n = \text{poly}(\lambda^n)$.

We also consider the following distributions:

$$\begin{aligned} & \left(\text{prFE}_{h^*}.\text{mpk}, \left\{ F_{h^*-1}^{j_{h^*-1}}, \text{prFE}_{h^*}.\text{sk}^{j_{h^*-1}} \right\}_{j_{h^*-1}}, \left\{ \text{prFE}_{h^*}.\text{Enc}(\mathbf{x}^{j_{h^*}, \dots, j_n}) \right\}_{j_{h^*}, \dots, j_n}, \text{aux}_{h^*} \right) \\ \text{and} \quad & \left(\text{prFE}_{h^*}.\text{mpk}, \left\{ F_{h^*-1}^{j_{h^*-1}}, \text{prFE}_{h^*}.\text{sk}^{j_{h^*-1}} \right\}_{j_{h^*-1}}, \left\{ \Delta^{j_{h^*}, \dots, j_n} \leftarrow \mathcal{CT}_{\text{prFE}_{h^*}} \right\}_{j_{h^*}, \dots, j_n}, \text{aux}_{h^*} \right). \end{aligned} \tag{6.30}$$

By the security guarantee of prFE_{h^*} with respect to Samp_{h^*} , we can see that an adversary \mathcal{A} that can distinguish the distributions in Equation (6.30) with advantage more than ϵ can be converted into another adversary \mathcal{B} that can distinguish the following distributions with advantage more than $\epsilon' \stackrel{\text{def}}{=} \epsilon/Q(\lambda') - \text{negl}(\kappa)$ satisfying $\text{Size}(\mathcal{B}) \leq Q(\lambda')\text{Size}(\mathcal{A})$ for some polynomial Q :

$$\left(\left\{ F_{h^*-1}^{j_{h^*-1}} \right\}_{j_{h^*-1}}, \left\{ F_{h^*-1}^{j_{h^*-1}}(\mathbf{x}^{j_{h^*}, \dots, j_n}) \right\}_{j_{h^*-1}, \dots, j_n}, \text{aux}_{h^*} \right)$$

$$\text{and } \left(\left\{ F_{h^*-1}^{j_{h^*-1}} \right\}_{j_{h^*-1}}, \left\{ \Delta_{h^*-1, \dots, j_n}^{j_{h^*-1}, \dots, j_n} \leftarrow \mathcal{CT}_{\text{prFE}_{h^*}} \right\}_{j_{h^*-1}, \dots, j_n}, \text{aux}_{h^*} \right). \quad (6.31)$$

By inspection, one can see that the distributions in Equation (6.30) are equivalent to those in Equation (6.29) with $h = h^*$.¹² Therefore, to complete the proof, it suffices to show that \mathcal{B} can be used as a distinguisher against the distributions in Equation (6.29) with $h = h^* - 1$ whose advantage is at least $\epsilon' - \text{negl}(\kappa)$. To show this, we first observe that the second distribution in Equation (6.31) is equivalent to that in Equation (6.29) with $h = h^* - 1$. Therefore, it suffices to prove that \mathcal{A} cannot distinguish the first distribution in Equation (6.31) from the first distribution in Equation (6.29) with $h = h^* - 1$ with more than negligible advantage in κ . To show this, we consider the following sequence of hybrids.

Hyb₁. This is the first distribution of Equation (6.31). Recall that we have

$$F_{h^*-1}^{j_{h^*-1}}(\mathbf{x}^{j_{h^*-1}, \dots, j_n}) = \text{prFE}_{h^*-1} \cdot \text{Enc} \left(\text{prFE}_{h^*-1} \cdot \text{mpk}, \left(\text{SKE.sk}, (\mathbf{x}_{h^*-1}^{j_{h^*-1}}, \mathbf{r}_{h^*-1}^{j_{h^*-1}}), \dots, (\mathbf{x}_n^{j_n}, \mathbf{r}_1^{j_n}, \dots, \mathbf{r}_{h^*-2}^{j_n}) \right); \right. \\ \left. \text{PRF}_{h^*-1}(K_{h^*-1}^{j_n}, (\mathbf{r}_{h^*-1}^{j_{h^*-1}}, \dots, \mathbf{r}_{n-1}^{j_{n-1}})) \right).$$

Hyb₂. This hybrid is the same as the previous one except that we replace $\text{PRF}_{h^*-1}(K_{h^*-1}^{j_n}, \cdot)$ with the real random function $\mathbf{R}^{j_n}(\cdot)$ for each $j_n \in [q_{\text{msg}}]$. Since $K_{h^*-1}^{j_n}$ is not used anywhere else, we can use the subexponential security of PRF to conclude that this hybrid is computationally indistinguishable from the previous one. In particular, by our choice of the parameter $2^{\Lambda^\delta} > \kappa^{\omega(1)}$, we can conclude that the adversary cannot distinguish this hybrid from the previous one with advantage more than $\text{negl}(\kappa)$.

Hyb₃. This hybrid is the same as the previous one except that we output a failure symbol when there exist $(j_{h^*-1}, \dots, j_{n-1}) \neq (j'_{h^*-1}, \dots, j'_{n-1})$ such that

¹²Equation (6.30) includes additional terms $\{F_{h^*-1}^{j_{h^*-1}}\}_{j_{h^*-1}}$ while Equation (6.29) does not. We ignore this difference, since these terms can be efficiently computed from aux_{h^*} and does not affect the indistinguishability.

$(\mathbf{r}_{h^*-1}^{j_{h^*-1}}, \dots, \mathbf{r}_{n-1}^{j_{n-1}}) = (\mathbf{r}_{h^*-1}^{j'_{h^*-1}}, \dots, \mathbf{r}_{n-1}^{j'_{n-1}})$. We show that the probability of this happening is negligible in κ . To prove this, it suffices to show that there are no $i \in [h^* - 1, n - 1]$, $j, j' \in [q_i]$ satisfying $j \neq j'$ and $\mathbf{r}_i^j = \mathbf{r}_i^{j'}$. The probability of this happening can be bounded by $(q_1^2 + \dots + q_{n-1}^2)/2^{2\Lambda}$ by taking the union bound with respect to all the combinations of i, j, j' . By our choice of Λ , this is bounded by $\text{negl}(\kappa)$.

Hyb₄. In this hybrid, we replace $F_{h^*-1}^{j_{h^*-1}}(\mathbf{x}^{j_{h^*}, \dots, j_n})$ with

$$\text{prFE}_{h^*-1}.\text{ct}^{j_{h^*-1}, \dots, j_n} = \text{prFE}_{h^*-1}.\text{Enc} \left(\begin{array}{c} \text{prFE}_{h^*-1}.\text{mpk}, \\ (\text{SKE.sk}, (\mathbf{x}_{h^*-1}^{j_{h^*-1}}, \mathbf{r}_{h^*-1}^{j_{h^*-1}}), \dots, (\mathbf{x}_n^{j_n}, K_1^{j_n}, \dots, K_{h^*-2}^{j_n})) \end{array} \right).$$

Namely, we use fresh randomness for each encryption instead of deriving the randomness by $R^{j_n}(\mathbf{r}_{h^*-1}^{j_{h^*-1}}, \dots, \mathbf{r}_{n-1}^{j_{n-1}})$. We claim that this change is only conceptual. To see this, we observe that unless the failure condition introduced in **Hyb₃** is satisfied, every invocation of the function R^{j_n} is with respect to a fresh input and thus the output can be replaced with a fresh randomness.

Hyb₅. In this hybrid, we remove the failure event. Namely, we always outputs $\text{prFE}_{h^*-1}.\text{ct}^{j_{h^*-1}, \dots, j_n}$ regardless of whether the failure event happens or not. Since the failure event happens with probability only $\text{negl}(\kappa)$ probability, we conclude that the adversary is not able to distinguish this hybrid from the previous one with more than $\text{negl}(\kappa)$ probability.

Noting that the final hybrid is equivalent to the first distribution in Equation (6.29) with $h = h^* - 1$, we complete the proof. \blacksquare

Lemma 6.24. Assuming that $\text{Adv}_{\mathcal{A}}^1(\lambda)$ (defined in Equation (6.29)) is negligible in κ for all non-uniform adversary \mathcal{A} such that $\text{Size}(\mathcal{A}) = \text{poly}(\kappa)$, we have $\text{Adv}_{\mathcal{B}}^n(\lambda) = \text{negl}(\lambda)$ for all non-uniform adversary \mathcal{B} such that $\text{Size}(\mathcal{B}) = \text{poly}(\lambda)$.

Proof. For the sake of contradiction, suppose that there exists an adversary $\mathcal{B} = \{\mathcal{B}_\lambda\}_\lambda$

such that $\epsilon_n(\lambda) \stackrel{\text{def}}{=} \text{Adv}_{\mathcal{B}}^n(\lambda)$ is non-negligible and $t_n(\lambda) \stackrel{\text{def}}{=} \text{Size}(\mathcal{B})$ is polynomial in λ . In particular, this implies that there exists an infinite set $\mathcal{L} \subseteq \mathbb{N}$ and some polynomial p such that $\epsilon_n(\lambda) \geq 1/p(\lambda)$ and $t_n(\lambda) \leq p(\lambda)$ for all $\lambda \in \mathcal{L}$. We then define $t_i(\lambda) \stackrel{\text{def}}{=} p(\lambda)\lambda^{n(n-i)\log \lambda}$ and $\epsilon_i(\lambda) \stackrel{\text{def}}{=} 1/p(\lambda)\lambda^{n(n-i)\log \lambda}$ for $i = 1, \dots, n-1$. We can also see that $t_1(\lambda) < \kappa$ and $\epsilon_1(\lambda) > 1/\kappa$ for sufficiently large λ . This implies that there exists $\lambda_0 \in \mathbb{N}$ such that for all $\lambda > \lambda_0$, there is no adversary \mathcal{A}_λ such that $\text{Size}(\mathcal{A}_\lambda) \leq t_1(\lambda)$ and $\text{Adv}_{\mathcal{A}_\lambda}^1(\lambda) \geq \epsilon_1(\lambda)$. We then consider the following statement that is parameterized by λ and $h \in [n]$:

Statement $_{\lambda,h}$: There exists an adversary \mathcal{A}_λ such that $\text{Size}(\mathcal{A}_\lambda) \leq t_h(\lambda)$ and $\text{Adv}_{\mathcal{A}_\lambda}^h(\lambda) \geq \epsilon_h(\lambda)$.

For each $\lambda \in \mathcal{L} \cap \mathbb{N}_{>\lambda_0}$, there exists $h_\lambda^* \in [2, n]$ such that **Statement $_{\lambda,h_\lambda^*-1}$** is false and **Statement $_{\lambda,h_\lambda^*}$** is true, since **Statement $_{\lambda,1}$** is false and **Statement $_{\lambda,n}$** is true. However, by applying Theorem 6.23 to the adversary guaranteed by **Statement $_{\lambda,h_\lambda^*}$** being true for the sequence $\{h_\lambda^*\}_\lambda$, we obtain another adversary $\mathcal{A}' = \{\mathcal{A}'_\lambda\}_\lambda$ such that $\text{Size}(\mathcal{A}'_\lambda) \leq t_{h^*}Q(\lambda^n)$ and $\text{Adv}_{\mathcal{A}'_\lambda}^{h^*-1}(\lambda) \geq \epsilon_{h^*}/Q(\lambda^n) - \text{negl}(\kappa) \geq \epsilon_{h^*}/2Q(\lambda^n)$ for some polynomial Q . In particular, $\text{Size}(\mathcal{A}'_\lambda) \leq t_{h^*-1}(\lambda)$ and $\text{Adv}_{\mathcal{A}'_\lambda}^{h^*-1}(\lambda) \geq \epsilon_{h^*-1}(\lambda)$ for all sufficiently large λ , since we have $\lambda^{n \log \lambda} > Q(\lambda^n)$ for all sufficiently large λ . However, this contradicts the above assertion that **Statement $_{\lambda,h_\lambda^*-1}$** is false. This concludes the proof. \blacksquare

The following lemma completes Step 2 of the proof of Theorem 6.21.

Lemma 6.25. *If SKE is IND_r secure, Equation (6.24) implies Equation (6.22).*

Proof. We first observe that Equation (6.24) is equivalent to saying $\text{Adv}_{\mathcal{A}}^1(\lambda) = \text{negl}(\kappa)$ for all \mathcal{A} with $\text{Size}(\mathcal{A}) = \text{poly}(\kappa)$. By Theorem 6.24, this implies that $\text{Adv}_{\mathcal{A}}^n(\lambda) = \text{negl}(\lambda)$ for all \mathcal{A} with $\text{Size}(\mathcal{A}) = \text{poly}(\lambda)$. Namely, we have

$$\left(\text{aux}_{\mathcal{A}}, \{\text{prFE}_i.\text{mpk}\}_{i \in [n]}, \{f_k, \text{sk}_{f_k} = \text{prFE}_1.\text{sk}_{f_k}\}_{k \in [q_0]} \right) \\ \left(\{\text{SKE.ct}_i^{j_i}, \mathbf{r}_i^{j_i}, \text{prFE}_{i+1}.\text{sk}^{j_i}\}_{i \in [n-1], j_i \in [q_i]}, \{\text{prFE}_n.\text{ct}^{j_n}\}_{j_n \in [q_n]} \right)$$

$$\approx_c \left(\begin{array}{c} \text{aux}_{\mathcal{A}}, \{\text{prFE}_i.\text{mpk}\}_{i \in [n]}, \{f_k, \text{sk}_{f_k} = \text{prFE}_1.\text{sk}_{f_k}\}_{k \in [q_0]} \\ \left\{ \text{SKE.ct}_i^{j_i}, \mathbf{r}_i^{j_i}, \text{prFE}_{i+1}.\text{sk}^{j_i} \right\}_{\substack{i \in [n-1], \\ j_i \in [q_i]}}, \left\{ \delta_n^{j_n} \right\}_{j_n \in [q_n]} \end{array} \right).$$

By IND_r security of SKE, we have

$$\begin{aligned} & \left(\begin{array}{c} \text{aux}_{\mathcal{A}}, \{\text{prFE}_i.\text{mpk}\}_{i \in [n]}, \{f_k, \text{sk}_{f_k} = \text{prFE}_1.\text{sk}_{f_k}\}_{k \in [q_0]} \\ \left\{ \text{SKE.ct}_i^{j_i}, \mathbf{r}_i^{j_i}, \text{prFE}_{i+1}.\text{sk}^{j_i} \right\}_{\substack{i \in [n-1], \\ j_i \in [q_i]}}, \left\{ \delta_n^{j_n} \right\}_{j_n \in [q_n]} \end{array} \right) \\ & \approx_c \left(\begin{array}{c} \text{aux}_{\mathcal{A}}, \{\text{prFE}_i.\text{mpk}\}_{i \in [n]}, \{f_k, \text{sk}_{f_k} = \text{prFE}_1.\text{sk}_{f_k}\}_{k \in [q_0]} \\ \left\{ \gamma_i^{j_i}, \mathbf{r}_i^{j_i}, \text{prFE}_{i+1}.\text{sk}^{j_i} \right\}_{\substack{i \in [n-1], \\ j_i \in [q_i]}}, \left\{ \delta_n^{j_n} \right\}_{j_n \in [q_n]} \end{array} \right). \end{aligned}$$

Combining the above equations, Equation (6.22) readily follows. ■

We conclude the proof of Theorem 6.21. ■

Remark 35 (Comparison with [164]). We note that in high level, overall structure of our security proof above is similar to that of witness encryption in [164]. In both proofs, the main step considers parameterized distributions $\{\mathcal{D}_{h,b}\}_{h \in [n], b \in \{0,1\}}$ and shows that $\mathcal{D}_{1,0} \approx_c \mathcal{D}_{1,1}$ holds if $\mathcal{D}_{n,0}$ and $\mathcal{D}_{n,1}$ are indistinguishable even with subexponentially small advantage against subexponential time adversary. To show this claim, [164] uses the evasive LWE assumption, while we use the security of prFE, which in turn is reduced to the evasive LWE assumption. While this difference stems simply from the fact that we introduce the intermediate primitive of prFE to construct mi-prFE instead of directly constructing it from evasive LWE, there are more fundamental differences as well. In particular, we identify certain subtle issues in the proof by [164] and fix these by strengthening the assumptions. We elaborate on this in the following.

- **On the multiplicative invocation of evasive LWE.** To prove $\mathcal{D}_{1,0} \approx_c \mathcal{D}_{1,1}$, [164] assumes that there exists an adversary \mathcal{A}_1 that distinguishes them with non-negligible advantage ϵ and polynomial time t for the sake of contradiction. They then invoke the evasive LWE assumption with respect to an appropriately defined sampler Samp_1 to conclude that there exists a distinguishing adversary \mathcal{A}_2 against $\mathcal{D}_{2,0}$ and $\mathcal{D}_{2,1}$. This process continues multiple times, where they invoke evasive LWE with respect to the security parameter $\lambda_j := 2^j \lambda$ and an adversary \mathcal{A}_j for

the j -th invocation to obtain another adversary \mathcal{A}_{j+1} , where \mathcal{A}_k is a distinguisher against $\mathcal{D}_{k,0}$ and $\mathcal{D}_{k,1}$. Denoting the distinguishing advantage against $\mathcal{D}_{j,0}$ and $\mathcal{D}_{j,1}$ of \mathcal{A}_j by ϵ_j , we have $\epsilon_{j+1} \geq \epsilon_j / \text{poly}_j(\lambda_j)$, where poly_j is a polynomial that is determined by the sampler Samp_j . Finally, they obtain a distinguisher \mathcal{A}_n against $\mathcal{D}_{n,0}$ and $\mathcal{D}_{n,1}$, where $\epsilon_n = \epsilon / \text{poly}_1(\lambda_1) \text{poly}_2(\lambda_2) \cdots \text{poly}_n(\lambda_n)$ and the running time being $\text{poly}_1(\lambda_1) \text{poly}_2(\lambda_2) \cdots \text{poly}_n(\lambda_n)$. They derive the conclusion by saying

$$\text{poly}_1(\lambda_1) \text{poly}_2(\lambda_2) \cdots \text{poly}_n(\lambda_n) = \text{poly}(\lambda_1 \cdots \lambda_n) = \text{poly}(2^{n^2} \lambda^n) \quad (6.32)$$

and setting the parameter so that there is no adversary of this running time and distinguishing advantage against $\mathcal{D}_{n,0}$ and $\mathcal{D}_{n,1}$. However, a subtlety is that $\text{poly}_1(\lambda_1) \text{poly}_2(\lambda_2) \cdots \text{poly}_n(\lambda_n) = \text{poly}(\lambda_1 \cdots \lambda_n)$ is not necessarily true. For example, one can consider the setting where we have $\text{poly}_j(\lambda) = \lambda^{2^j}$. This example may look a bit artificial, but it does not contradict the evasive LWE assumption, since j is treated as a constant asymptotically. In words, the issue arises from the fact that even if each polynomial has a constant exponent, the maximum of the exponents can be arbitrarily large function in λ , when we consider non-constant number of polynomials. In this setting, \mathcal{A}_n 's distinguishing advantage is too small to derive the contradiction.

The above issue occurs due to the invocations of evasive LWE super-constant times. To resolve the problem, we consider non-uniform sampler $\{\text{Samp}_{h^*}\}_{h^*}$ that hardwires the "best" index h^* and invoke the evasive LWE only with respect to this sampler (See Theorem 6.23 and 6.24). We avoid the above problem, since we invoke the evasive LWE only once in the proof. However, this solution entails the strengthening of the assumption where we consider non-uniform samplers. We believe that the same strengthening of the assumption is required for the proof in [164] as well.

- **On the additive term of evasive LWE.** Here, we also discuss the other subtlety that arises in the proof by [164]. To focus on the issue, we ignore the first issue discussed above and assume $\text{poly}_j(\lambda) = \lambda^c$ holds for some fixed $c \in \mathbb{N}$ that does not depend on j , which makes Equation (6.32) correct. In the proof of [164] (and in our explanation above), we implicitly ignore the negligible additive term when applying evasive LWE. Namely, when we apply the assumption with respect to \mathcal{A}_j , the lower bound for the advantage ϵ_{j+1} of \mathcal{A}_{j+1} should be $\epsilon_{j+1} \geq \epsilon_j / \text{poly}(\lambda_j) - \text{negl}(\lambda_j)$ rather than $\epsilon_{j+1} \geq \epsilon_j / \text{poly}(\lambda_j)$. This does not cause any difference when we consider the setting where ϵ_j is non-negligible in λ_j . However, for larger j , ϵ_j is negligible function in the security parameter λ_j . Concretely, the lower bound on ϵ_{n-1} obtained by ignoring the additive term is $\epsilon / \text{poly}(2^{(n-1)^2} \lambda^{n-1})$.¹³ If we apply evasive LWE once more with respect to $\lambda_n = 2^n \lambda^n$ to complete the proof, we have $\epsilon_n \geq \epsilon_{n-1} / \text{poly}(\lambda_n) - \text{negl}(\lambda_n)$. The RHS of the inequality may be negative, since $\epsilon_{n-1} / \text{poly}(\lambda_n)$ is some specific negligible function in λ_n and this may be smaller than the second term $\text{negl}(\lambda_n)$. Therefore, what we can derive here is the trivial

¹³Namely, the actual value of ϵ_{n-1} may be even smaller.

bound $\epsilon_n \geq 0$, which is not enough for our purpose. To fix this issue, we introduce additional parameter κ and then modify the assumption so that the additive term is negligibly small in κ , which is set much larger than λ . Again, we believe that the same strengthening of the assumption is required for the proof in [164] as well. Finally, we note that the above problem does not occur when n is constant. This is because in that case, we have ϵ_n is non-negligible and thus the problem of ϵ_n may be smaller than negl_n does not occur.

6.3.4 Security Proof for Constant n (with Weaker Assumption)

Here, we prove the security of our construction in the case of n being a constant. The reason why we consider the security proof separately for this special case is that we can give a proof from better assumptions than the general case. In more detail, the security is proven assuming the standard security notion for prFE, rather than the non-uniform and κ version of it. As a result, the security of the mi-prFE is reduced to the (plain) evasive LWE instead of non-uniform κ -evasive LWE. The reason why we can achieve this is that in the case of n being constant, we can avoid all the subtleties that arise in the general case. We refer to Remark 35 for more discussions.

Theorem 6.26. *Let $SC_{\text{mi-prFE}}$ be a samples class for mi-prFE. Suppose prFE _{i} scheme satisfies prCT security as per Definition 5.13 with respect to the sampler class that contains all $\text{Samp}_{\text{prFE}}(1^\lambda)$, induced by $\text{Samp}_{\text{mi-prFE}} \in SC_{\text{mi-prFE}}$, as in Section 6.3.3, SKE satisfies IND_r security and PRF _{i} is secure, then mi-prFE constructed in Section 6.3.2 for constant n satisfies security for $\kappa = \lambda^n$ as in Definition 6.11. Note that this in particular implies the κ -security defined in Definition 6.10.*

Proof. The proof of this theorem largely follows that of Theorem 6.21. The crucial difference is that to get the equivalent of Theorem 6.24, we simply invoke the (non- κ , uniform) security of prFE n -times. We sketch the proof below while highlighting the difference. We consider the same simulator as in the proof of Theorem 6.21 and divide the proof steps into Step 1 and Step 2 in the same manner.

- We start with Step 1, which consists of proving that Equation (6.23) implies Equation (6.24). This is proven in the same manner as Theorem 6.22. However here, since we set $\kappa = \lambda^n = \text{poly}(\lambda)$, we do not need subexponential IND_r security

and only (polynomial) INDr security suffices for SKE.

- We then move to prove Step 2. The goal here is to prove Equation (6.24) implies Equation (6.22).
 - We first observe that the uniform version of Theorem 6.23 holds by the same proof, where we only consider constant h^* rather than arbitrary sequence $h^* = \{h_\lambda^*\}_\lambda$ and uniform PPT adversaries. Notice that then the sampler is now uniform, since it no longer has to hardwire the sequence $\{h_\lambda^*\}_\lambda$. In this setting, non-uniform κ -prCT security collapses to the (plain) prCT security, since $\kappa = \lambda^n = \text{poly}(\lambda)$ and the sampler is uniform. Therefore, plain prCT security is sufficient for the proof. Furthermore, since we set $\kappa = \lambda^n = \text{poly}(\lambda)$, we do not need subexponential security for PRF and standard security suffices.
 - We then consider an analogue of Theorem 6.24, which asserts that if $\text{Adv}_{\mathcal{A}}^1(\lambda)$ (defined in Equation (6.29)) is negligible for all (uniform) PPT adversary that runs in polynomial time in λ , then so is $\text{Adv}_{\mathcal{A}}^n(\lambda)$. This is proven by observing that the indistinguishability of the distributions in Equation (6.29) for $h = h^* - 1$ implies that for h^* by the analogue of Theorem 6.23 explained in the previous item. By applying this n -times, we obtain the conclusion.
 - We finally conclude the proof by the same argument as Theorem 6.25.

This completes the proof of Theorem 6.26. ■

We encapsulate the results of this section using the following theorems.

Theorem 6.27 (mi-prFE for poly arity). *Let $\kappa = \lambda^{n^2 \log \lambda}$. Assume non-uniform κ -evasive LWE (Assumption 6.14), subexponentially secure PRF against non-uniform adversary, and non-uniform sub-exponential LWE (Assumption 2.5). Then there exists a mi-prFE scheme for arity $n = \text{poly}(\lambda)$, supporting functions with input length L and bounded polynomial depth $d = d(\lambda)$, and satisfying κ -security (Definition 6.10) with efficiency $|\text{mpk}| = \text{poly}(n, L, d, \Lambda, \lambda)$, $|\text{sk}_f| = \text{poly}(d, \lambda)$, $|\text{ct}_1| = nL\text{poly}(\text{dep}, \lambda)$, $|\text{ct}_i| = \text{poly}(n, L, d, \Lambda, \lambda)$ for $i \in [2, n]$ where $\Lambda := (n^2\lambda)^{1/\delta}$.*

Theorem 6.28 (mi-prFE for constant arity). *Let $\kappa = \lambda^n$. Assume evasive LWE (Assumption 2.6) and LWE (Assumption 2.5). Then there exists a mi-prFE scheme for arity $n = O(1)$, supporting functions with input length L and bounded polynomial depth $d = d(\lambda)$, and satisfying κ -security (Definition 6.10) with efficiency $|\text{mpk}| = \text{poly}(n, L, d, \lambda)$, $|\text{sk}_f| = \text{poly}(d, \lambda)$, $|\text{ct}_1| = nL\text{poly}(\text{dep}, \lambda)$, $|\text{ct}_i| =$*

$\text{poly}(n, L, d, \lambda)$ for $i \in [2, n]$.

6.4 MULTI-INPUT PREDICATE ENCRYPTION FOR POLYNOMIAL ARITY FOR P

In this section, we provide our construction of multi-input predicate encryption (miPE) for all circuits as an application of mi-prFE. Similarly to Section 6.3, we consider two settings where the arity is either constant or arbitrary polynomial. The former setting leads to a construction with weaker assumption where we only require (plain) evasive LWE and LWE. This improves the construction by [22], which additionally requires non-standard tensor LWE assumption. The latter setting leads to a construction with polynomial arity and for all circuits under the stronger non-uniform κ -evasive LWE assumption. This resolves the open problem posed by [22].

6.4.1 Construction

In this section, we give a construction of a miPE scheme using mi-prFE and PE. The construction will support functions with arity $n = n(\lambda)$ where input string for each arity is in $\{0, 1\}^L$ and the output is $\{0, 1\}$. We further restrict the depth of the circuits that implement the function by a parameter dep . We denote this function class by \mathcal{F}_{prm} , where prm is the set of the parameters n, L, dep . Namely, \mathcal{F}_{prm} consists of n -ary functions that takes as input strings x_1, \dots, x_n , where each $x_i \in \{0, 1\}^L$ and outputs $f(x_1, \dots, x_n) \in \{0, 1\}$. The message space for each slot is $\{0, 1\}$. Namely, we have $\mu_1, \dots, \mu_n \in \{0, 1\}$ in the following.

Building Blocks. Below, we list the building blocks required for our construction.

1. A single input predicate encryption scheme $\text{PE} = \text{PE}(\text{Setup}, \text{KeyGen}, \text{Enc}, \text{Dec})$ for function family supporting functions $f : \{0, 1\}^{nL} \rightarrow \{0, 1\}$ that can be represented as circuits with depth at most dep . We denote by this function class by $\mathcal{F}_{\text{prm}_{\text{PE}}}$, where $\text{prm}_{\text{PE}} = (1^{nL}, 1^{\text{dep}})$ is the parameter that specifies the circuit class. We also assume that the scheme has message space $\{0, 1\}^n$ and satisfies IND ρ security (Definition 6.5). We can

instantiate such a PE by using the construction by [109] or by the combination of lockable obfuscation [172] (a.k.a compute-and-compare obfuscation [172]) and ABE for circuits by [43], for example.

We use \mathcal{CT}_{PE} to denote the ciphertext space, $\ell_{\text{ct}}^{\text{PE}}$ to denote the ciphertext length and $d_{\text{Enc}}^{\text{PE}}$ to denote the depth of the circuit required to compute the PE.Enc algorithm.

2. A n -input FE for pseudorandom functionalities $\text{mi-prFE} = \text{mi-prFE}(\text{Setup}, \text{KeyGen}, \text{Enc}_1, \dots, \text{Enc}_n, \text{Dec})$ for function family $\mathcal{F}_{L'(\lambda), d_{\text{Enc}}^{\text{PE}}}$ consisting of circuits with input space $\{0, 1\}^{L'}$ and output space $\{0, 1\}$ where we set $L' = L + \lambda + 1$ for our construction and with maximum depth $d_{\text{Enc}}^{\text{PE}}$. We denote the parameters that specify $\mathcal{F}_{L'(\lambda), d_{\text{Enc}}^{\text{PE}}}$ by $\text{prm}_{\text{mi-prFE}}$.
3. A pseudorandom function $\text{PRF} : \{0, 1\}^\Lambda \times \{0, 1\}^{(n-1)\Lambda} \rightarrow \{0, 1\}^{R_{\text{len}}}$, where $\{0, 1\}^\Lambda$ and $\{0, 1\}^{(n-1)\Lambda}$ are the key space and input space respectively and R_{len} is the length of randomness used in the PE.Enc algorithm. We will discuss how to set Λ in Remark 36.

Next, we describe our construction.

$\text{Setup}(1^\lambda, 1^n, \text{prm}) \rightarrow (\text{mpk}, \text{msk})$. The setup algorithm does the following.

- Generate $(\text{mi-prFE.mpk}, \text{mi-prFE.msk}) \leftarrow \text{mi-prFE.Setup}(1^\lambda, 1^n, \text{prm}_{\text{mi-prFE}})$.
- Generate $(\text{PE.mpk}, \text{PE.msk}) \leftarrow \text{PE.Setup}(1^\Lambda, \text{prm}_{\text{PE}})$.
- Output $\text{mpk} = (\text{mi-prFE.mpk}, \text{PE.mpk})$ and $\text{msk} = (\text{mi-prFE.msk}, \text{PE.msk})$.

$\text{KeyGen}(\text{mpk}, \text{msk}, f) \rightarrow \text{sk}_f$. The key generation algorithm does the following.

- Parse $\text{mpk} = (\text{mi-prFE.mpk}, \text{PE.mpk})$ and $\text{msk} = (\text{mi-prFE.msk}, \text{PE.msk})$.
- Compute $\text{PE.sk}_f \leftarrow \text{PE.KeyGen}(\text{PE.msk}, f)$.
- Compute $\text{mi-prFE.sk}_F \leftarrow \text{mi-prFE.KeyGen}(\text{mi-prFE.msk}, F[\text{PE.mpk}])$ where $F[\text{PE.mpk}]$ is defined as follows:

$$\begin{aligned} F[\text{PE.mpk}] & ((\mathbf{x}_1, \mu_1, \mathbf{r}_1), \dots, (\mathbf{x}_{(n-1)}, \mu_{(n-1)}, \mathbf{r}_{(n-1)}), (\mathbf{x}_n, \mu_n, \mathbf{r}_n)) \\ &= \text{PE.Enc}(\text{PE.mpk}, (\mathbf{x}_1, \dots, \mathbf{x}_n), (\mu_1, \dots, \mu_n); \text{PRF}(\mathbf{r}_n, (\mathbf{r}_1, \dots, \mathbf{r}_{(n-1)}))) \end{aligned}$$

- Output $\text{sk}_f = (\text{PE.sk}_f, \text{mi-prFE.sk}_F)$ ¹⁴.

$\text{Enc}_i(\text{msk}, \mathbf{x}_i, \mu_i) \rightarrow \text{ct}_i$ for $1 \leq i \leq n$. The slot i encryption algorithm does the following.

- Parse $\text{msk} = (\text{mi-prFE.msk}, \text{PE.msk})$.
- For $1 \leq i \leq n$, sample $\mathbf{r}_i \leftarrow \{0, 1\}^\Lambda$ and compute $\text{mi-prFE.ct}_i \leftarrow \text{mi-prFE.Enc}_i(\text{mi-prFE.msk}, (\mathbf{x}_i, \mu_i, \mathbf{r}_i))$.
- Output $\text{ct}_i := \text{mi-prFE.ct}_i$.

$\text{Dec}(\text{mpk}, \text{sk}_f, f, \text{ct}_1, \dots, \text{ct}_n) \rightarrow y \in \{0, 1\}^n \cup \{\perp\}$. The decryption algorithm does the following.

- Parse $\text{mpk} = (\text{mi-prFE.mpk}, \text{PE.mpk})$, $\text{sk}_f = (\text{PE.sk}_f, \text{mi-prFE.sk}_F)$ and $\{\text{ct}_i = \text{mi-prFE.ct}_i\}_{i \in [n]}$.
- Compute $\text{PE.ct} = \text{mi-prFE.Dec}(\text{mi-prFE.mpk}, \text{mi-prFE.sk}_F, f, \text{mi-prFE.ct}_1, \dots, \text{mi-prFE.ct}_n)$.
- Compute $y = \text{PE.Dec}(\text{PE.mpk}, \text{PE.sk}_f, f, \text{PE.ct})$.
- Output y .

Remark 36. Here, we discuss how we set Λ . Similarly to the case of mi-prFE in Section 6.3, we consider two cases of parameter settings for the construction. One is the case of n being constant. In this case, we simply set $\Lambda = \lambda$. In the general case of $n = \text{poly}(\lambda)$, we assume that PRF and SKE are subexponentially secure. This means that there exists $0 < \delta < 1$ such that there is no adversary with size 2^{λ^δ} and distinguishing advantage $2^{-\lambda^\delta}$. Similarly to the case of mi-prFE (See Remark 34 for further discussion), we set Λ so that it satisfies $2^{\Lambda^\delta} \geq \kappa^{\omega(1)}$. An example choice would be to take $\Lambda := (n^2\lambda)^{1/\delta}$.

¹⁴Note that one can compute mi-prFE.sk_F in the Setup algorithm and output it as a part of public key rather than the secret key as it does not require the knowledge of f . We output this here for notational convenience.

Correctness. We prove the correctness of our scheme using the following theorem.

Theorem 6.29. *Assume PE and mi-prFE schemes are correct, and PRF is secure. Then the above construction of miPE scheme is correct.*

Proof. From the correctness of the mi-prFE scheme and definition of function F , we have

$$\text{mi-prFE.Dec}(\text{mi-prFE.mpk}, \text{mi-prFE.sk}_F, F, \text{mi-prFE.ct}_1, \dots, \text{mi-prFE.ct}_n) = \text{PE.ct}$$

with probability 1, where

$$\text{PE.ct} = \text{PE.Enc}(\text{PE.mpk}, (\mathbf{x}_1, \dots, \mathbf{x}_n), (\mu_1, \dots, \mu_n); \text{PRF}(\mathbf{r}_n, (\mathbf{r}_1, \dots, \mathbf{r}_{(n-1)}))) .$$

Next, using the security of PRF, with all but negl advantage, $\text{PRF}(\mathbf{r}_n, (\mathbf{r}_1, \dots, \mathbf{r}_{(n-1)}))$ is indistinguishable from $R \leftarrow \{0, 1\}^{R_{\text{len}}}$. Since

$$\text{PE.Dec}(\text{PE.mpk}, \text{PE.sk}_f, f, \text{PE.Enc}(\text{PE.mpk}, (\mathbf{x}_1, \dots, \mathbf{x}_n), (\mu_1, \dots, \mu_n); R)) = (\mu_1, \dots, \mu_n)$$

holds for randomly chosen R with all but negl probability if $f(\mathbf{x}_1, \dots, \mathbf{x}_n) = 1$ by the correctness of PE scheme, the above holds even for $R = \text{PRF}(\mathbf{r}_n, (\mathbf{r}_1, \dots, \mathbf{r}_{(n-1)}))$ with all but negl probability. Hence, the correctness. \blacksquare

6.4.2 Security

Here, we prove the security of our scheme. The following theorem asserts the security of the scheme for the case of n being arbitrary polynomial.

Theorem 6.30. *Assume mi-prFE scheme is secure (as per Definition 6.10) with respect to $\kappa = \lambda^{n^2 \log \lambda}$ and the sampler class containing the sampler $\text{Samp}_{\text{mi-prFE}}$ as defined in Eq. 6.34, PE scheme is sub-exponentially secure (Definition 6.5) and PRF is sub-exponentially secure. Then the construction of miPE is secure as per Definition 6.6.*

Proof. Suppose the adversary \mathcal{A} outputs the following:

1. **Key Queries.** It issues q_0 number of functions f_1, \dots, f_{q_0} for key queries.
2. **Ciphertext Queries.** It issues q_i number of messages for ciphertext queries for slot i . We use $(\mathbf{x}_i^{j_i}, \mu_i^{j_i})$ to denote the j_i -th ciphertext query corresponding to the i -th slot, where $j_i \in [q_i]$ and $i \in [n]$.
3. **Auxiliary Information.** It outputs the auxiliary information $\text{aux}_{\mathcal{A}}$.

To prove the security as per Definition 6.6 we first define $\{\text{Sim}_i\}_{i \in [n]}$ as follows.

$\text{Sim}_i(\text{msk}) \rightarrow \text{ct}_i$ for $i \in [n]$. Parse $\text{msk} = (\text{mi-prFE.msk}, \text{PE.msk})$. Set $\text{Sim}_i(\text{msk}) = \text{mi-prFE.Sim}_i(\text{mi-prFE.msk})$, where $\{\text{mi-prFE.Sim}_i\}_{i \in [n]}$ is the simulator whose existence is guaranteed by the security of mi-prFE (See Definition 6.10).

Then, it suffices to show

$$\begin{aligned} & \left(\begin{array}{c} \text{aux}_{\mathcal{A}}, \text{mpk} = (\text{mi-prFE.mpk}, \text{PE.mpk}), \\ \{F[\text{PE.mpk}], f_k, \text{sk}_{f_k} = (\text{PE.sk}_{f_k}, \text{mi-prFE.sk}_F)\}_{k \in [q_0]}, \\ \{\text{ct}_i^{j_i} = \text{mi-prFE.ct}_i^{j_i}\}_{i \in [n], j_i \in [q_i]} \end{array} \right) \\ & \approx_c \left(\begin{array}{c} \text{aux}_{\mathcal{A}}, \text{mpk} = (\text{mi-prFE.mpk}, \text{PE.mpk}), \\ \{F[\text{PE.mpk}], f_k, \text{sk}_{f_k} = (\text{PE.sk}_{f_k}, \text{mi-prFE.sk}_F)\}_{k \in [q_0]}, \\ \{\delta_i^{j_i} \leftarrow \text{mi-prFE.Sim}_i(\text{msk})\}_{i \in [n], j_i \in [q_i]} \end{array} \right) \end{aligned} \quad (6.33)$$

where

$$(\text{aux}_{\mathcal{A}}, \{f_k\}_k, \{(\mathbf{x}_i^{j_i}, \mu_i^{j_i})\}_{i, j_i}) \leftarrow \mathcal{A}(1^\lambda),$$

$$(\text{mpk} = (\text{mi-prFE.mpk}, \text{PE.mpk}), \text{msk} = (\text{mi-prFE.msk}, \text{PE.msk})) \leftarrow \text{Setup}(1^\lambda, 1^n, \text{prm}),$$

$$\text{mi-prFE.sk}_F \leftarrow \text{mi-prFE.KeyGen}(\text{mi-prFE.msk}, F[\text{PE.mpk}]),$$

$$\text{PE.sk}_{f_k} \leftarrow \text{PE.KeyGen}(\text{PE.msk}, f_k) \text{ for } k \in [q_0],$$

$$\text{mi-prFE.ct}_i^{j_i} \leftarrow \text{mi-prFE.Enc}_i(\text{mi-prFE.msk}, (\mathbf{x}_i^{j_i}, \mu_i^{j_i}, \mathbf{r}_i^{j_i})), \mathbf{r}_i^{j_i} \leftarrow \{0, 1\}^\Lambda \text{ for } i \in [n], j_i \in [q_i].$$

We invoke the security of mi-prFE with sampler $\text{Samp}_{\text{mi-prFE}}$ that outputs

$$\left(\begin{array}{ll} \text{Function:} & F[\text{PE.mpk}], \\ \text{Inputs:} & \left\{ \left((\mathbf{x}_1^{j_1}, \mu_1^{j_1}, \mathbf{r}_1^{j_1}), \dots, (\mathbf{x}_n^{j_n}, \mu_n^{j_n}, \mathbf{r}_n^{j_n}) \right) \right\}_{i \in [n], j_i \in [q_i]}, \\ \text{Auxiliary Information:} & \text{aux} = (\text{aux}_{\mathcal{A}}, \text{PE.mpk}, \{f_k, \text{PE.sk}_{f_k}\}_{k \in [q_0]}) \end{array} \right) \quad (6.34)$$

From the security guarantee of the mi-prFE scheme with sampler $\text{Samp}_{\text{mi-prFE}}$ and simulator $\{\text{mi-prFE.Sim}_i\}_{i \in [n]}$, we have that

$$\begin{aligned} & \left(\begin{array}{l} \text{aux, mi-prFE.mpk, } F[\text{PE.mpk}], \text{mi-prFE.sk}_F, \\ \{\text{mi-prFE.ct}_i^{j_i} \leftarrow \text{mi-prFE.Enc}_i(\text{mi-prFE.msk}, (\mathbf{x}_i^{j_i}, \mu_i^{j_i}, \mathbf{r}_i^{j_i}))\}_{i \in [n], j_i \in [q_i]} \end{array} \right) \\ & \approx_c \left(\begin{array}{l} \text{aux, mi-prFE.mpk, } F[\text{PE.mpk}], \text{mi-prFE.sk}_F, \\ \{\text{mi-prFE.ct}_i^{j_i} \leftarrow \text{mi-prFE.Sim}_i(\text{mi-prFE.msk})\}_{i \in [n], j_i \in [q_i]} \end{array} \right) \\ & \quad \text{if } \left(1^\kappa, \text{aux}, \{F[\text{PE.mpk}], \text{PE.ct}_{j_1, \dots, j_n}\}_{i \in [n], j_i \in [q_i]} \right) \\ & \quad \approx_c \left(1^\kappa, \text{aux}, \{F[\text{PE.mpk}], \text{PE.}\Delta_{j_1, \dots, j_n}\}_{i \in [n], j_i \in [q_i]} \right) \end{aligned}$$

where $\text{PE.}\Delta_{j_1, \dots, j_n} \leftarrow \mathcal{CT}_{\text{PE}}$ for $i \in [n], j_i \in [q_i]$ and

$$\begin{aligned} \text{PE.ct}_{j_1, \dots, j_n} &= F[\text{PE.mpk}] \left(\left(\mathbf{x}_1^{j_1}, \mu_1^{j_1}, \mathbf{r}_1^{j_1} \right), \dots, \left(\mathbf{x}_{(n-1)}^{j_{(n-1)}}, \mu_{(n-1)}^{j_{(n-1)}}, \mathbf{r}_{(n-1)}^{j_{(n-1)}} \right), \left(\mathbf{x}_n^{j_n}, \mu_n^{j_n}, \mathbf{r}_n^{j_n} \right) \right) \\ &= \text{PE.Enc} \left(\text{PE.mpk}, (\mathbf{x}_1^{j_1}, \dots, \mathbf{x}_n^{j_n}), (\mu_1^{j_1}, \dots, \mu_n^{j_n}); \text{PRF} \left(\mathbf{r}_n^{j_n}, (\mathbf{r}_1^{j_1}, \dots, \mathbf{r}_{(n-1)}^{j_{(n-1)}}) \right) \right) \end{aligned}$$

and $\text{PE.}\Delta_{j_1, \dots, j_n} \leftarrow \mathcal{CT}_{\text{PE}}$ for $i \in [n], j_i \in [q_i]$.

Thus to prove Equation (6.33), it suffices to show

$$\left(\begin{array}{l} 1^\kappa, \text{aux}_{\mathcal{A}}, F[\text{PE.mpk}], \text{PE.mpk}, \{f_k, \text{PE.sk}_{f_k}\}_{k \in [q_0]}, \\ \left\{ \text{PE.ct}_{j_1, \dots, j_n} = \text{PE.Enc} \left(\text{PE.mpk}, (\mathbf{x}_1^{j_1}, \dots, \mathbf{x}_n^{j_n}), (\mu_1^{j_1}, \dots, \mu_n^{j_n}); \right. \right. \\ \quad \left. \left. \text{PRF}(\mathbf{r}_n^{j_n}, (\mathbf{r}_1^{j_1}, \dots, \mathbf{r}_{(n-1)}^{j_{(n-1)}})) \right) \right\}_{i \in [n], j_i \in [q_i]} \end{array} \right) \quad (6.35)$$

$$\approx_c \left(1^\kappa, \text{aux}_{\mathcal{A}}, F[\text{PE.mpk}], \text{PE.mpk}, \{f_k, \text{PE.sk}_{f_k}\}_{k \in [q_0]}, \right. \\ \left. \{\text{PE}.\Delta_{j_1, \dots, j_n} \leftarrow \mathcal{CT}_{\text{PE}}\}_{i \in [n], j_i \in [q_i]} \right)$$

We prove the above via the following sequence of hybrids.

Hyb₁. This is Equation (6.35).

$$\left(1^\kappa, \text{aux}_{\mathcal{A}}, F[\text{PE.mpk}], \text{PE.mpk}, \{f_k, \text{PE.sk}_{f_k}\}_{k \in [q_0]}, \right. \\ \left. \left\{ \text{PE.ct}_{j_1, \dots, j_n} = \text{PE.Enc} \left(\text{PE.mpk}, (\mathbf{x}_1^{j_1}, \dots, \mathbf{x}_n^{j_n}), (\mu_1^{j_1}, \dots, \mu_n^{j_n}); \text{PRF}(\mathbf{r}_n^{j_n}, (\mathbf{r}_1^{j_1}, \dots, \mathbf{r}_{(n-1)}^{j_{(n-1)}})) \right) \right\}_{i \in [n], j_i \in [q_i]} \right)$$

Hyb₂. This hybrid is the same as the previous one except that we replace $\text{PRF}(\mathbf{r}_n^{j_n}, \cdot)$ with the real random function $\mathbf{R}^{j_n}(\cdot)$ for each $j_n \in [q_n]$. Since $\mathbf{r}_n^{j_n}$ is not used anywhere else, we can use the subexponential security of PRF to conclude that this hybrid is computationally indistinguishable from the previous one. In particular, by our choice of the parameter $2^{\Lambda^\delta} > \kappa^{\omega(1)}$, we can conclude that the adversary cannot distinguish this hybrid from the previous one with advantage more than $\text{negl}(\kappa)$.

Hyb₄. This hybrid is the same as the previous one except that we output a failure symbol when there exist $(j_1, \dots, j_{n-1}) \neq (j'_1, \dots, j'_{n-1})$ such that $(\mathbf{r}_1^{j_1}, \dots, \mathbf{r}_{n-1}^{j_{n-1}}) = (\mathbf{r}_1^{j'_1}, \dots, \mathbf{r}_{n-1}^{j'_{n-1}})$. We show that the probability of this happening is negligible in κ . To prove this, it suffices to show that there are no $i \in [n]$, $j, j' \in [q_i]$ satisfying $j \neq j'$ and $\mathbf{r}_i^j = \mathbf{r}_i^{j'}$. The probability of this happening can be bounded by $(q_1^2 + \dots + q_{n-1}^2)/2^{2\Lambda}$ by taking the union bound with respect to all the combinations of i, j, j' . By our choice that $2^{\Lambda^\delta} > \kappa^{\omega(1)}$, this probability is bounded by $\text{negl}(\kappa)$.

Hyb₅. This hybrid is the same as the previous one except that we compute

$$\text{PE.ct}_{j_1, \dots, j_n} \leftarrow \text{PE.Enc} \left(\text{PE.mpk}, (\mathbf{x}_1^{j_1}, \dots, \mathbf{x}_n^{j_n}), (\mu_1^{j_1}, \dots, \mu_n^{j_n}) \right).$$

Namely, we use fresh randomness for each encryption instead of deriving the

randomness from $R^{j_n}(\cdot)$. This change is only conceptual.

In this hybrid, the view of the adversary is

$$\left(\begin{array}{c} 1^\kappa, \text{aux}_{\mathcal{A}}, F[\text{PE.mpk}], \text{PE.mpk}, \{f_k, \text{PE.sk}_{f_k}\}_{k \in [q_0]}, \\ \left\{ \text{PE.ct}_{j_1, \dots, j_n} \leftarrow \text{PE.Enc} \left(\text{PE.mpk}, (\mathbf{x}_1^{j_1}, \dots, \mathbf{x}_n^{j_n}), (\mu_1^{j_1}, \dots, \mu_n^{j_n}) \right) \right\}_{i \in [n], j_i \in [q_i]} \end{array} \right)$$

Hyb₆. This hybrid is the same as the previous one except that we use sub-exponential security of underlying PE scheme to replace $\text{PE.Enc}(\text{PE.mpk}, (\mathbf{x}_1^{j_1}, \dots, \mathbf{x}_n^{j_n}), (\mu_1^{j_1}, \dots, \mu_n^{j_n}))$ with $\text{PE}.\delta_{j_1, \dots, j_n} \leftarrow \mathcal{CT}_{\text{PE}}$ for all $i \in [n]$ and $j_i \in [q_i]$. By admissibility of \mathcal{A} in the miPE security, we have $f(\mathbf{x}_1^{j_1}, \dots, \mathbf{x}_n^{j_n}) = 0$ which satisfies the admissibility of the single input PE security game. We therefore replace each of $\text{PE.ct}_{j_1, \dots, j_n}$ with a random string by the PE security. By our choice of the parameter $2^{\Lambda^\delta} > \kappa^{\omega(1)}$ and subexponential security of PE, the distinguishing advantage of the adversary when we replace single ciphertext is $\text{negl}(\kappa)$. Since there are $q_1 \cdots q_n = \text{poly}(\kappa)$ ciphertexts, we can conclude that the distinguishing advantage between this hybrid from the previous one is at most $\text{poly}(\kappa)\text{negl}(\kappa) = \text{negl}(\kappa)$.

In this hybrid, the view of the adversary is

$$\left(\begin{array}{c} 1^\kappa, \text{aux}_{\mathcal{A}}, F[\text{PE.mpk}], \text{PE.mpk}, \{f_k, \text{PE.sk}_{f_k}\}_{k \in [q_0]}, \\ \left\{ \text{PE}.\delta_{j_1, \dots, j_n} \leftarrow \mathcal{CT}_{\text{PE}} \right\}_{i \in [n], j_i \in [q_i]} \end{array} \right)$$

which is the distribution on RHS in Equation (6.35).

Hence, the proof. ■

Constant arity case. In the special case of n being a constant, we can base the security of the scheme on weaker security requirements for the underlying ingredients. Concretely, we have the following theorem.

Theorem 6.31. *Assume mi-prFE scheme is secure (as per Definition 6.10) with respect to $\kappa = \lambda^n$ and the sampler class containing the sampler $\text{Samp}_{\text{mi-prFE}}$ as defined in*

Eq. 6.34, PE scheme is secure (Definition 6.5) and PRF is secure. Then the construction of miPE is secure as per Definition 6.6.

The proof of the above theorem is exactly the same as Theorem 6.30 except that here $\kappa = \lambda^n$. Since $\kappa = \text{poly}(\lambda)$, we do not need subexponential security for PRF and PE. In addition, since mi-prFE for constant arity with $\kappa = \lambda^n$ can be constructed from (plain) evasive LWE, which is weaker assumption than non-uniform κ -version of it that is necessary for the general case.

We encapsulate the results of this section using the following theorems.

Theorem 6.32 (miPE for poly arity). *Let $\kappa = \lambda^{n^2 \log \lambda}$. Assuming non-uniform κ -evasive LWE (Assumption 6.14), subexponentially secure PRF against non-uniform adversary, and non-uniform sub-exponential LWE (Assumption 2.5), there exists a miPE scheme for arity $n = \text{poly}(\lambda)$, supporting functions of bounded polynomial depth $\text{dep} = \text{dep}(\lambda)$, and satisfying security as per Definition 6.6.*

Theorem 6.33 (miPE for constant arity). *Assuming evasive LWE (Assumption 2.6) and LWE (Assumption 2.5), there exists a mi-prFE scheme for arity $n = O(1)$, supporting functions of bounded polynomial depth $d = d(\lambda)$, and satisfying security as per Definition 6.6.*

6.5 INDISTINGUISHABILITY OBFUSCATION FOR PSEUDORANDOM FUNCTIONALITIES

6.5.1 Definition

In this section we give the definitions for indistinguishability obfuscation for pseudorandom functionalities (prIO) for circuits.

Syntax. An indistinguishability obfuscator for pseudorandom functionalities consists of the following algorithms.

$iO(1^\lambda, C) \rightarrow \tilde{C}$. The obfuscation algorithm takes as input the security parameter λ and a circuit $C : \{0, 1\}^n \rightarrow \{0, 1\}^m$ with arbitrary n and m . It outputs an obfuscated circuit \tilde{C} .

$\text{Eval}(\tilde{C}, x) \rightarrow y$. The evaluation algorithm takes as input an obfuscated circuit \tilde{C} and an input $x \in \{0, 1\}^n$. It outputs y .

A uniform PPT machine iO is an indistinguishability obfuscator for pseudorandom functionalities w.r.t parameter $\kappa = \kappa(\lambda)$ if it satisfies the following properties.

Definition 6.12 (Polynomial Slowdown). For all security parameters $\lambda \in \mathbb{N}$, for any circuit C and every input x , the evaluation time of $iO(1^\lambda, C)$ on x is at most polynomially slower than the run time of the circuit C on x .

Definition 6.13 (Correctness). For all security parameters $\lambda \in \mathbb{N}$, for all integers n, m , all circuits $C : \{0, 1\}^n \rightarrow \{0, 1\}^m$, and all input $x \in \{0, 1\}^n$, we have that:

$$\Pr[C' \leftarrow iO(1^\lambda, C) : C'(x) = C(x)] = 1$$

where the probability is taken over the coin-tosses of the obfuscator iO .

Definition 6.14 (Indistinguishability for Pseudorandom Functionality). For the security parameter $\lambda = \lambda(\lambda)$, let Samp be a PPT algorithm that on input 1^λ , outputs

$$(C_0, C_1, \text{aux} \in \{0, 1\}^*)$$

where $C_0 : \{0, 1\}^n \rightarrow \{0, 1\}^m$ and $C_1 : \{0, 1\}^n \rightarrow \{0, 1\}^m$ have the same description size. We say that a prIO scheme is secure with respect to the sampler class \mathcal{SC} if for every PPT sampler $\text{Samp} \in \mathcal{SC}$ the following holds.

$$\text{If } (1^\kappa, \{C_0(x)\}_{x \in \{0, 1\}^n}, \text{aux}) \approx_c (1^\kappa, \{\Delta_x\}_{x \in \{0, 1\}^n}, \text{aux}) \approx_c (1^\kappa, \{C_1(x)\}_{x \in \{0, 1\}^n}, \text{aux}) \quad (6.36)$$

$$\text{then } (iO(1^\lambda, C_0), \text{aux}) \approx_c (iO(1^\lambda, C_1), \text{aux}) \quad (6.37)$$

where $\Delta_x \leftarrow \{0, 1\}^m$ and $\kappa \geq 2^n$.

Remark 37. Note that 1^κ in the precondition is introduced for the purpose of padding, allowing the distinguisher for the distributions to run in time polynomial in κ . The reason why we require $\kappa \geq 2^n$ is that the input length to the distinguisher is polynomial in 2^n anyway and in order for the padding to make sense, κ should satisfy this condition.

Remark 38. It is shown in [19, 65] that there is no prIO scheme satisfying the above security for all general samplers. Therefore, when we use the security of prIO, we invoke the security *with respect to a specific sampler class* that is induced by the respective applications.

6.5.2 Construction

Our construction follows the blueprint of the multi-input FE to iO conversion by Ananth and Jain [29]. To obfuscate a circuit with input domain $\{0, 1\}^n$, we generate a mi-prFE instance for arity $n + 1$. We then let C be encrypted in position $n + 1$, and the two inputs 0 and 1 be encrypted in position i for $i \in [n]$. The $2n + 1$ ciphertexts together with a secret key for the universal circuit and the public parameters would form the iO .

Building Blocks. We use a $(n + 1)$ -input prFE scheme $\text{mi-prFE} = \text{mi-prFE}(\text{Setup}, \text{KeyGen}, \{\text{Enc}_i\}_{i \in [n+1]}, \text{Dec})$ for the circuit class with fixed input length, bounded depth, and binary output. We require mi-prFE to satisfy κ -security defined as per Definition 6.10. We can instantiate the scheme by our construction in Section 6.3.2 with $\kappa = \lambda^{n^2 \log \lambda}$.

Next, we describe the construction of prIO for all circuits.

$iO(1^\lambda, C)$. Given as input the security parameter 1^λ and a circuit $C : \{0, 1\}^n \rightarrow \{0, 1\}^m$ for arbitrary input length n , output length m , and description size L , do the following:

- Run $(\text{mpk}, \text{msk}) \leftarrow \text{mi-prFE.Setup}(1^\lambda, 1^{n+1}, \text{prm})$, where prm specifies message length L and the maximum depth d of the circuits supported by the mi-prFE instance. We set d to be the depth of the universal circuit U that, upon input an n -ary circuit C and vector $\mathbf{x} \in \{0, 1\}^n$, outputs $U(C, \mathbf{x}) = C(\mathbf{x})$.

- Compute $\text{ct}_{n+1} \leftarrow \text{mi-prFE.Enc}_{n+1}(\text{msk}, C)$.
- For $i \in [n]$ and $b \in \{0, 1\}$, compute $\text{ct}_{i,b} = \text{mi-prFE.Enc}_i(\text{msk}, b)$.
- Compute $\text{sk}_U \leftarrow \text{mi-prFE.KeyGen}(\text{msk}, U)$, where U is defined in the first item above.
- Output $\tilde{C} = (\{ \text{ct}_{i,b} \}_{i \in [n], b \in \{0,1\}}, \text{ct}_{n+1}, \text{sk}_U, \text{mpk})$

Eval(\tilde{C}, \mathbf{x}). Given as input an obfuscated circuit \tilde{C} and an input $\mathbf{x} \in \{0, 1\}^n$, do the following:

1. Parse $\tilde{C} = (\{ \text{ct}_{i,b} \}_{i \in [n], b \in \{0,1\}}, \text{ct}_{n+1}, \text{sk}_U, \text{mpk})$.
2. Output $\text{mi-prFE.Dec}(\text{mpk}, \text{sk}_U, U, \text{ct}_{x_1}, \dots, \text{ct}_{x_n}, \text{ct}_{n+1})$.

Remark 39. We note that the input size of mi-prFE scheme varies for slot 0, where we encrypt C , and slot i , where we encrypt a bit b , for $i \in [n]$. To make the input size consistent throughout the slots, we can pad the bit b with $\mathbf{0}$ (say) such that $|b\mathbf{0}| = |C|$ and give a mi-prFE key for a circuit U which on input $(C, b_1\mathbf{0}, \dots, b_n\mathbf{0})$, where $b_i \in \{0, 1\}$, simply discards the padding and outputs $C(b_1, \dots, b_n)$.

Correctness. The correctness of the scheme follows in a straightforward manner from the correctness of the underlying mi-prFE scheme and the definition of the universal circuit U .

6.5.3 Security

Theorem 6.34. Let $\mathcal{SC}_{\text{prIO}}$ be a sampler class for prIO. Suppose mi-prFE scheme is secure (Definition 6.10) for $\kappa = \kappa(\lambda)$ w.r.t sampler class that contains all $\text{Samp}_{\text{mi-prFE},0}$ and $\text{Samp}_{\text{mi-prFE},1}$, induced by $\text{Samp}_{\text{prIO}} \in \mathcal{SC}_{\text{prIO}}$, as defined in Equation (6.41). Then the prIO scheme satisfies security as defined in Definition 6.14 with $\kappa = \kappa(\lambda)$.

Proof. Consider a sampler $\text{Samp}_{\text{prIO}}$ that generates the following:

1. **Obfuscation Query.** It issues $C_0, C_1 : \{0, 1\}^n \rightarrow \{0, 1\}^m$ with the same size L as an obfuscation query.

2. **Auxiliary Information.** It outputs the auxiliary information $\text{aux}_{\mathcal{A}}$.

To prove the security as per Definition 6.14, we show that

$$\left(\{\text{ct}_{i,b}\}_{i \in [n], b \in \{0,1\}}, \text{ct}_{n+1}^0, \text{sk}_U, \text{mpk}, \text{aux}_{\mathcal{A}} \right) \approx_c \left(\{\text{ct}_{i,b}\}_{i \in [n], b \in \{0,1\}}, \text{ct}_{n+1}^1, \text{sk}_U, \text{mpk}, \text{aux}_{\mathcal{A}} \right) \quad (6.38)$$

$$\text{if } (1^\kappa, \{C_0(x_i)\}_{\forall x_i \in \mathcal{X}_\lambda}, \text{aux}_{\mathcal{A}}) \approx_c (1^\kappa, \{\Delta_i \leftarrow \mathcal{Y}_\lambda\}_i, \text{aux}_{\mathcal{A}}) \approx_c (1^\kappa, \{C_1(x_i)\}_{\forall x_i \in \mathcal{X}_\lambda}, \text{aux}_{\mathcal{A}}) \quad (6.39)$$

where

$$\begin{aligned} (C_0, C_1, \text{aux}_{\mathcal{A}}) &\leftarrow \text{Samp}_{\text{prIO}}(1^\lambda), \\ (\text{msk}, \text{mpk}) &\leftarrow \text{mi-prFE.Setup}(1^\lambda, 1^{n+1}, \text{prm}), \\ \text{sk}_U &\leftarrow \text{mi-prFE.KeyGen}(\text{msk}, U), \\ \text{ct}_{n+1}^0 &\leftarrow \text{mi-prFE.Enc}_{n+1}(\text{msk}, C_0), \text{ct}_{n+1}^1 \leftarrow \text{mi-prFE.Enc}_{n+1}(\text{msk}, C_1), \\ \text{ct}_{i,b} &= \text{mi-prFE.Enc}_i(\text{msk}, b), \text{ for } i \in [n], b \in \{0, 1\}. \end{aligned}$$

From the right hand indistinguishability of Equation (6.39), we have that

$$\begin{aligned} &\left(\{\text{ct}_{i,b}\}_{i \in [n], b \in \{0,1\}}, \text{ct}_{n+1}^0, \text{sk}_U, \text{mpk}, \text{aux}_{\mathcal{A}} \right) \\ &\approx_c \left(\{\gamma_{i,b} \leftarrow \text{mi-prFE.Sim}_i(\text{msk})\}_{i \in [n], b \in \{0,1\}}, \gamma_{n+1} \leftarrow \text{mi-prFE.Sim}_{n+1}(\text{msk}), \text{sk}_U, \text{mpk}, \text{aux}_{\mathcal{A}} \right) \end{aligned} \quad (6.40)$$

using the mi-prFE security with simulator $\{\text{mi-prFE.Sim}_i\}_{i \in [n+1]}$ and sampler

$\text{Samp}_{\text{mi-prFE},0}$ that outputs

$$\left(\begin{array}{ll} \text{Function:} & \text{Universal Circuit } U \\ \text{Inputs:} & \{x_1^{j_1} = j_1, \dots, x_n^{j_n} = j_n\}_{j_1 \in \{0,1\}, \dots, j_n \in \{0,1\}}, x_{n+1} = C_0 \\ \text{Auxiliary Information:} & \text{aux}_{\mathcal{A}} \end{array} \right) \quad (6.41)$$

To see this we note that the pseudorandomness of $U(x_1 \in \{0, 1\} \dots, x_n \in \{0, 1\}, x_{n+1} = C_{n+1} \in \{0, 1\}^L) = C_0(x_1, \dots, x_n)$ is implied from Equation (6.39).

Similarly, from the right hand indistinguishability of Equation (6.39), we have

$$\begin{aligned} & \left(\{\text{ct}_{i,b}\}_{i \in [n], b \in \{0,1\}}, \text{ct}_{n+1}^1, \text{sk}_U, \text{mpk}, \text{aux}_{\mathcal{A}} \right) \\ & \approx_c \left(\{\gamma_{i,b} \leftarrow \text{mi-prFE.Sim}_i(\text{msk})\}_{i \in [n], b \in \{0,1\}}, \gamma_{n+1} \leftarrow \text{mi-prFE.Sim}_{n+1}(\text{msk}), \text{sk}_U, \text{mpk}, \text{aux}_{\mathcal{A}} \right) \end{aligned} \quad (6.42)$$

using the mi-prFE security with simulator $\{\text{mi-prFE.Sim}_i\}_{i \in [0,n]}$ and sampler $\text{Samp}_{\text{mi-prFE},1}$ whose output is same as that of Samp_0 except that $x_{n+1} = C_1$.

From Equation (6.40) and Equation (6.42), we have

$$\begin{aligned} & \left(\{\text{ct}_{i,b}\}_{i \in [n], b \in \{0,1\}}, \text{ct}_{n+1}^0, \text{sk}_U, \text{mpk}, \text{aux}_{\mathcal{A}} \right) \\ & \approx_c \left(\{\text{ct}_{i,b}\}_{i \in [n], b \in \{0,1\}}, \text{ct}_{n+1}^1, \text{sk}_U, \text{mpk}, \text{aux}_{\mathcal{A}} \right) \end{aligned}$$

hence the proof. ■

The following theorem holds for the specific class of samplers we described in the proof.

Theorem 6.35. *Let $\kappa = \lambda^{n^2 \log \lambda}$. Assuming non-uniform κ -evasive LWE (Assumption 6.14), subexponentially secure PRF against non-uniform adversary, and non-uniform sub-exponential LWE (Assumption 2.5), there exists a prIO scheme for circuits with input domain $\{0, 1\}^n$ satisfying security as per Definition 6.14.*

6.6 POLYNOMIAL DOMAIN IO FOR PSEUDORANDOM FUNCTIONALITIES

In this section, we define and construct indistinguishability obfuscation for pseudorandom functionalities with polynomial size domain (pPRIO). The advantage of considering this restricted variant is that we can base the security of the construction on (plain) evasive LWE, rather than non-uniform κ -variant of it. Here, we introduce online-offline property and reusable security and provide a construction satisfying these properties. The reason why we introduce these properties is that they seem to be useful for some applications. Indeed, in our companion paper [12], we use pPRIO with these properties to construct ABE with optimized parameter size. We consider two variants of pPRIO: pPRIO with fixed input domain and pPRIO with flexible domain. As the name suggests, the latter is more flexible and desirable. To obtain the latter, we first construct the former from mi-prFE in Section 6.6.2 and then convert it into the latter in Section 6.6.3.

6.6.1 Definition

Definition 6.15 (Syntax). A pPRIO scheme consists of the following algorithms.

$\text{Obf}(1^\lambda, C) \rightarrow \text{Obf}$. The obfuscation algorithm takes as input the security parameter λ and a circuit $C : [N] \rightarrow [M]$ with $\text{size}(C) \leq L$ for some arbitrary polynomial $L = L(\lambda)$. It outputs an obfuscation of the circuit Obf .

We consider a definition where the Obf algorithm can be decomposed into the following two phases.

$\text{ObfOff}(1^\lambda, 1^L) \rightarrow (\text{obf}_{\text{off}}, \text{st})$. The offline obfuscation algorithm takes as input the security parameter λ and the circuit size bound L . It outputs obf_{off} and st .

$\text{ObfOn}(\text{st}, C) \rightarrow \text{obf}_{\text{on}}$. The online obfuscation algorithm takes as input the security parameter st and the circuit C and outputs obf_{on} .

The final output of Obf is $\text{Obf} = (\text{obf}_{\text{off}}, \text{obf}_{\text{on}})$.

$\text{Eval}(\text{Obf}, x) \rightarrow y$. The evaluation algorithm takes as input an obfuscated circuit Obf and an input $x \in [N]$. It outputs $y \in [M]$.

We consider the following syntax variants of a pPRIO scheme. We consider fixed input domain pPRIO and flexible input domain pPRIO. In the former, we need that ObfOff algorithm should know the input domain $[N]$ of C that is going to be input to ObfOn while in the latter, we do not.

Definition 6.16 (Fixed Input Domain pPRIO). A *fixed input domain* pPRIO scheme has syntax as in Definition 6.15 except that ObfOff takes N (in binary) as an additional input.

Definition 6.17 (Flexible Input Domain pPRIO). A *flexible input domain* pPRIO scheme has syntax exactly as in Definition 6.15, i.e., a variant without inputting N into ObfOff .

Next, we define the properties of a pPRIO scheme.

Definition 6.18 (Correctness for Flexible Input Domain Case). For all security parameters $\lambda \in \mathbb{N}$, for any $C : [N] \rightarrow [M]$, $L = L(\lambda)$ such that $\text{size}(C) \leq L$ and every input $x \in [N]$, we have that:

$$\Pr \left[\text{Eval}(\text{Obf}, x) = C(x) \mid \begin{array}{l} \text{Obf} = (\text{obf}_{\text{off}}, \text{obf}_{\text{on}}), (\text{obf}_{\text{off}}, \text{st}) \leftarrow \text{ObfOff}(1^\lambda, 1^L), \\ \text{obf}_{\text{on}} \leftarrow \text{ObfOn}(\text{st}, C) \end{array} \right] = 1$$

where the probability is taken over the coin-tosses of the obfuscator Obf . The correctness for the fixed input domain case is defined as above, except that $\text{ObfOff}(1^\lambda, 1^L)$ is replaced with $\text{ObfOff}(1^\lambda, 1^L, N)$.

We introduce a bit unusual efficiency requirement for pPRIO. Namely, we require that pPRIO can obfuscate circuits from exponentially large input domain efficiently. One may think that with this property, pPRIO is equivalent to prIO. However, for pPRIO, we require that it can securely obfuscate a circuit only when it has polynomial size domain and do not require any security guarantee when it has superpolynomial domain size. This relaxation leads to the construction with weaker assumption (i.e., uniform and non- κ variant of evasive LWE).

The following efficiency requirement is the consequence of inputting N into ObfOff in *binary* form and the requirement that ObfOff is a PPT algorithm, which implies that it runs in polynomial time in the input length. We explicitly require this for emphasizing the property.

Definition 6.19 (Efficiency for Fixed Input Domain). We require that the running time of $\text{ObfOff}(1^\lambda, 1^L, N)$ is bounded by $\text{poly}(\lambda, L, \log N)$.

We introduce the security definition for pPRIO .

Definition 6.20 (Security for Flexible Input Domain). Let Samp be a PPT algorithm that on input 1^λ , outputs

$$(1^{N_1+N_2+\dots+N_Q}, 1^L, \text{aux}, C^1, \dots, C^Q), \quad \text{where } C^i : [N_i] \rightarrow [M_i], \text{ size}(C^i) \leq L$$

where we enforce Samp to output $1^{N_1+N_2+\dots+N_Q}$ to make sure that all N_i are bounded by $\text{poly}(\lambda)$. We say that a flexible input domain pPRIO scheme is secure with respect to the sampler class \mathcal{SC} if for every PPT sampler $\text{Samp} \in \mathcal{SC}$ the following holds.

$$\text{If } (\text{aux}, 1^L, \{C^1(i)\}_{i \in [N_1]}, \dots, \{C^Q(i)\}_{i \in [N_Q]}) \approx_c (\text{aux}, 1^L, \{\Delta_i^1\}_{i \in [N_1]}, \dots, \{\Delta_i^Q\}_{i \in [N_Q]}) \quad (6.43)$$

$$\text{then } (\text{aux}, \text{obf}_{\text{off}}, \text{obf}_{\text{on}}^1, \dots, \text{obf}_{\text{on}}^Q) \approx_c (\text{aux}, \text{obf}_{\text{off}}, \delta^1 \leftarrow C\mathcal{T}^1, \dots, \delta^Q \leftarrow C\mathcal{T}^Q), \quad (6.44)$$

where $\Delta_i^j \leftarrow [M_j]$ for $j \in [Q], i \in [N_j]$, $(\text{obf}_{\text{off}}, \text{st}) \leftarrow \text{ObfOff}(1^\lambda, 1^L)$, $\text{obf}_{\text{on}}^j \leftarrow \text{ObfOn}(\text{st}, C^j)$ for $j \in [Q]$, and $C\mathcal{T}^j$ denotes the set of binary strings of the same length as the output of $\text{obf}_{\text{on}}(\text{st}, C^j)$ algorithm.

Definition 6.21 (Security for Fixed Input Domain). The security for flexible input domain pPRIO is defined exactly as in Definition 6.20, except that we have $N_1 = N_2 = \dots = N_Q$.

Remark 40 (Comparison with Definition 6.14). Here, we compare the above security definitions for pPRIO with that for prIO (Definition 6.14). On the one hand, the above security definition is weaker than Definition 6.14 in that the adversary is not allowed to submit a circuit with exponential size input domain. This restriction is captured by the

above definitions where we enforce the adversary to output the size of the input domain for each circuit in *unary*.

On the other hand, the above security definitions are stronger than Definition 6.14 in two folds. First, it requires the online part of the obfuscation obf_{on} to be *pseudorandom*, whereas Definition 6.14 only requires the obfuscation of the circuits are *computationally indistinguishable* (when their truth tables are pseudorandom). Secondly, the above security definition considers a security game where the same internal state st is reused for obfuscating *multiple* circuits, whereas only a single circuit is obfuscated in Definition 6.14. The above security definition may look odd, but pPRIO with this security notion has proven useful in our companion paper [12] to obtain optimal parameters for ABE schemes.

Remark 41. Similar to Remark 38, there is no pPRIO scheme satisfying the above security for all general samplers. Therefore, when we use the security of pPRIO, we invoke the security *with respect to a specific sampler class* that is induced by the respective applications.

6.6.2 Construction for Fixed Input Domain

In this section, we provide a construction of fixed input domain pPRIO scheme $\text{pPRIO} = (\text{Obf} = (\text{ObfOff}, \text{ObfOn}), \text{Eval})$ for circuit class $C = \{C : [\lambda^c] \rightarrow \{0, 1\}\}$, for some constant $c \in \mathbb{N}$, from mi-prFE for $c + 1$ arity. The construction is extended to be able to handle flexible input domain in Section 6.6.3.

Building Blocks. We use the following ingredient for our construction.

1. A mi-prFE scheme $\text{mi-prFE} = \text{mi-prFE}(\text{Setup}, \text{KeyGen}, \{\text{Enc}_i\}_{i \in [c+1]}, \text{Dec})$ for the circuit class with fixed input length $L = |C|$, bounded depth $d = d(\lambda)$, and binary output. We require mi-prFE to satisfy the pseudorandomness of the last slot ciphertext as per Definition 6.11. We can instantiate the scheme by our construction in Section 6.3.2.

Next, we describe our construction.

$\text{ObfOff}(1^\lambda, 1^L, N = \lambda^c)$. The offline phase of obfuscation algorithm does the following.

- Run $(\text{mi-prFE.mpk}, \text{mi-prFE.msk}) \leftarrow \text{mi-prFE.Setup}(1^\lambda, 1^{c+1}, \text{prm})$, where prm specifies the maximum size L of the circuit that is going to input to ObfOn and the maximum depth d of the circuit class supported by the mi-prFE instance. We set d to be the depth of the universal circuit U that, upon input a c -ary circuit C and vector $\mathbf{x} \in [\lambda]^c$, outputs $U(C, \mathbf{x}) = C(\mathbf{x})$, where \mathbf{x} is interpreted as an integer in $[\lambda^c]$ by some efficient bijective mapping between $[\lambda^c]$ and $[\lambda]^c$.
- For $i \in [c]$ and $j \in [\lambda]$, compute $\text{mi-prFE.ct}_{i,j} = \text{mi-prFE.Enc}_i(\text{mi-prFE.msk}, j)$.¹⁵
- Compute $\text{mi-prFE.sk}_U = \text{mi-prFE.KeyGen}(\text{mi-prFE.msk}, U)$, where U is supported by the mi-prFE instance because of our choice of prm .
- Output $\text{obf}_{\text{off}} := (\text{mi-prFE.mpk}, \{\text{mi-prFE.ct}_{i,j}\}_{i \in [c], j \in [\lambda]}, \text{mi-prFE.sk}_U)$ and $\text{st} = \text{mi-prFE.msk}$.

$\text{ObfOn}(\text{st}, C)$. The online phase of obfuscation algorithm does the following.

- Run $\text{mi-prFE.ct}_{c+1} \leftarrow \text{mi-prFE.Enc}_{c+1}(\text{mi-prFE.msk}, C)$.
- Output $\text{obf}_{\text{on}} := \text{mi-prFE.ct}_{c+1}$.

$\text{Eval}(\text{Obf}, \mathbf{x})$. The evaluation algorithm does the following.

1. Parse $\text{Obf} = (\text{obf}_{\text{off}}, \text{obf}_{\text{on}})$ where $\text{obf}_{\text{off}} = (\text{mi-prFE.mpk}, \{\text{mi-prFE.ct}_{i,j}\}_{i \in [c], j \in [\lambda]}, \text{mi-prFE.sk}_U)$ and $\text{obf}_{\text{on}} = \text{mi-prFE.ct}_{c+1}$.
2. Output $\text{mi-prFE.Dec}(\text{mi-prFE.mpk}, \text{mi-prFE.sk}_U, U, \text{mi-prFE.ct}_{1,x_1}, \dots, \text{mi-prFE.ct}_{c,x_c}, \text{mi-prFE.ct}_{c+1})$, where $\mathbf{x} \in [\lambda]^c$ is mapped to $(x_1, \dots, x_c) \in [\lambda]^c$ by the bijective mapping between $[\lambda^c]$ and $[\lambda]^c$.

Correctness. The correctness of the scheme follows in a straightforward manner from the correctness of the underlying mi-prFE scheme and the definition of the universal circuit U .

¹⁵Using Remark 39, it is safe to use this instance of mi-prFE to encrypt messages with smaller length than L .

Efficiency. The scheme satisfies

$$|\text{Obf}_{\text{off}}| = \text{poly}(c, L, \lambda) = \text{poly}(\lambda, L), |\text{Obf}_{\text{on}}| = \text{poly}(c, L, \lambda) = \text{poly}(\lambda, L),$$

where $(\text{Obf}_{\text{off}}, \text{Obf}_{\text{on}}) \leftarrow \text{Obf}(1^\lambda, C)$ for circuit $C : [N] \rightarrow [M]$ whose size is bounded by $L = L(\lambda)$.

Security. We prove the security of the fixed input pPRIO using the following theorem.

Theorem 6.36. *Let $\mathcal{SC}_{\text{pPRIO}}$ be a sampler class for pPRIO. Suppose mi-prFE scheme is secure (Definition 6.10) with $\kappa = \lambda^{c+1}$ and sampler class that contains all $\text{Samp}_{\text{mi-prFE}}$ induced by $\text{Samp}_{\text{prIO}} \in \mathcal{SC}_{\text{prIO}}$, as defined in Equation (6.47). Then the pPRIO scheme satisfies security as defined in Definition 6.21.*

Proof. Consider a sampler Samp that generates

$$(1^{Q \cdot N}, 1^L, \text{aux}_{\mathcal{A}}, \{C^k\}_{k \in [Q]}).$$

To prove the theorem, we show that

$$\begin{aligned} & \left(\text{aux}_{\mathcal{A}}, \text{obf}_{\text{off}} = (\text{mi-prFE.mpk}, \{\text{mi-prFE.ct}_{i,j}\}_{i \in [c], j \in [\lambda]}, \text{mi-prFE.sk}_U), \left\{ \text{obf}_{\text{on}}^k = \text{mi-prFE.ct}_{c+1}^k \right\}_{k \in [Q]} \right) \\ & \approx_c \left(\text{aux}_{\mathcal{A}}, \text{obf}_{\text{off}} = (\text{mi-prFE.mpk}, \{\text{mi-prFE.ct}_{i,j}\}_{i \in [c], j \in [\lambda]}, \text{mi-prFE.sk}_U), \right. \\ & \quad \left. \left\{ \text{obf}_{\text{on}}^k \leftarrow C\mathcal{T}_{\text{mi-prFE}, c+1} \right\}_{k \in [Q]} \right), \end{aligned} \quad (6.45)$$

holds assuming

$$\left(\text{aux}_{\mathcal{A}}, 1^L, \{C^1(i)\}_{i \in [N]}, \dots, \{C^Q(i)\}_{i \in [N]} \right) \approx_c \left(\text{aux}_{\mathcal{A}}, 1^L, \{\Delta_i^1\}_{i \in [N]}, \dots, \{\Delta_i^Q\}_{i \in [N]} \right) \quad (6.46)$$

where $C\mathcal{T}_{\text{mi-prFE}, c+1}$ denotes the set of binary strings with the same length as $\text{mi-prFE.ct}_{c+1}^k$ and

$$\begin{aligned} & (1^{Q \cdot N}, 1^L, \text{aux}_{\mathcal{A}}, \{C^k\}_{k \in [Q]}) \leftarrow \text{Samp}(1^\lambda), \\ & (\text{mi-prFE.mpk}, \text{mi-prFE.msk}) \leftarrow \text{mi-prFE.Setup}(1^\lambda, 1^{c+1}, \text{prm}), \end{aligned}$$

$$\begin{aligned}
\text{mi-prFE.sk}_U &\leftarrow \text{mi-prFE.KeyGen}(\text{mi-prFE.msk}, U), \\
\text{mi-prFE.ct}_{i,j} &\leftarrow \text{mi-prFE.Enc}_i(\text{mi-prFE.msk}, j) \text{ for } i \in [c], j \in [\lambda], \\
\Delta_1^k, \dots, \Delta_N^k &\leftarrow \{0, 1\} \text{ for } k \in [Q], \\
\text{mi-prFE.ct}_{c+1}^k &\leftarrow \text{mi-prFE.Enc}_{c+1}(\text{mi-prFE.msk}, C^k) \text{ for } k \in [Q].
\end{aligned}$$

We invoke the security of mi-prFE with sampler $\text{Samp}_{\text{mi-prFE}}$ that outputs

$$\left(\begin{array}{ll} \text{Function:} & \text{Universal Circuit } U \\ \text{Inputs:} & \{x_1^{j_1} = j_1, \dots, x_c^{j_c} = j_c, x_{c+1}^k = C^k\}_{j_1, \dots, j_c \in [\lambda], k \in [Q]} \\ \text{Auxiliary Information:} & \text{aux}_{\mathcal{A}} \end{array} \right) \quad (6.47)$$

Using the mi-prFE security with sampler $\text{Samp}_{\text{mi-prFE}}$, we have that

$$\begin{aligned}
&\left(\text{aux}_{\mathcal{A}}, \text{mi-prFE.mpk}, \{\text{mi-prFE.ct}_{i,j}\}_{i \in [c], j \in [\lambda]}, \text{mi-prFE.sk}_U, \{\text{mi-prFE.ct}_{c+1}^k\}_{k \in [Q]} \right) \\
&\approx_c \left(\text{aux}_{\mathcal{A}}, \text{mi-prFE.mpk}, \{\delta_{i,j}\}_{i \in [c], j \in [\lambda]}, \text{mi-prFE.sk}_U, \{\delta_{c+1}^k\}_{k \in [Q]} \right) \quad (6.48)
\end{aligned}$$

given

$$\left(\text{aux}_{\mathcal{A}}, \left\{ U(C^k, x_1^{j_1}, \dots, x_c^{j_c}) \right\}_{j_1, \dots, j_c \in [\lambda], k \in [Q]} \right) \approx_c \left(\text{aux}_{\mathcal{A}}, \{\Delta_i^k \leftarrow \{0, 1\}\}_{i \in [N], k \in [Q]} \right) \quad (6.49)$$

where $\delta_{i,j} \leftarrow \text{Sim}_i(\text{mi-prFE.msk})$ for $i \in [c], j \in [\lambda]$ and $\delta_{c+1}^k \leftarrow \mathcal{CT}_{\text{mi-prFE}, c+1}$ ¹⁶.

This implies

$$\begin{aligned}
&\left(\text{aux}_{\mathcal{A}}, \text{mi-prFE.mpk}, \{\text{mi-prFE.ct}_{i,j}\}_{i \in [c], j \in [\lambda]}, \text{mi-prFE.sk}_U, \{\text{mi-prFE.ct}_{c+1}^k\}_{k \in [Q]} \right) \\
&\approx_c \left(\text{aux}_{\mathcal{A}}, \text{mi-prFE.mpk}, \{\delta_{i,j}\}_{i \in [c], j \in [\lambda]}, \text{mi-prFE.sk}_U, \{\delta_{c+1}^k\}_{k \in [Q]} \right) \\
&\approx_c \left(\text{aux}_{\mathcal{A}}, \text{mi-prFE.mpk}, \{\text{mi-prFE.ct}_{i,j}\}_{i \in [c], j \in [\lambda]}, \text{mi-prFE.sk}_U, \{\delta_{c+1}^k\}_{k \in [Q]} \right) \quad (6.50)
\end{aligned}$$

given Equation (6.49), where the first indistinguishability follows from Equation (6.48)

¹⁶This follows from the fact that the simulator for the $c + 1$ -th slot outputs a random string as is required by Definition 6.11.

and the second also from Equation (6.48) by noting that replacing the last term with independently chosen random variable makes the distinguishing task of the distributions even harder. Equation (6.49) follows from Equation (6.46) and the definition of circuit U , hence the proof. ■

Instantiation. Instantiating the underlying mi-prFE scheme with constant arity $c + 1$ as in Theorem 6.28, we get the following theorem.

Theorem 6.37. *Assuming evasive LWE (Assumption 2.6) and LWE (Assumption 2.5), there exists a secure pPRIO scheme for fixed input domain supporting circuits of bounded size $L = \text{poly}(\lambda)$ with*

$$|\text{Obf}_{\text{off}}| = \text{poly}(L, \lambda), \quad |\text{Obf}_{\text{on}}| = \text{poly}(L, \lambda).$$

6.6.3 Construction for Flexible Input Domain

Here, we provide a construction of pPRIO for flexible input domain $\text{Flex.pPRIO} = \text{Flex.}(\text{ObfOff}, \text{ObfOn}, \text{Eval})$ for any circuits with binary output space from pPRIO for fixed input domain. The restriction that the output domain of the circuits being binary is removed in Section 6.6.4.

Building Blocks. Below, we list the ingredients for our construction.

1. A pseudorandom function $\text{PRF} : \{0, 1\}^\lambda \times \{0, 1\}^\lambda \rightarrow \{0, 1\}^\lambda$ with key space, input space and output space as $\{0, 1\}^\lambda$. The input to the PRF will be in the form of $\mathbf{x} \in [\lambda^c]$ for some constant $c \in \mathbb{N}$. Since we have $2^\lambda > \text{poly}(\lambda)$, the input space of the PRF is large enough to accommodate inputs $\mathbf{x} \in [\lambda^c]$ with an appropriate encoding. It is known that PRFs can be constructed from one-way functions.
2. A fixed input domain pPRIO scheme $\text{Fixed.}(\text{ObfOff}, \text{ObfOn}, \text{Eval})$ as constructed in Section 6.6.2 with the efficiency requirement defined in Definition 6.19 and the security defined as per Definition 6.20. The input domain of pPRIO should be $[\lambda^c]$ for some constant c .

$\text{Flex.Obf}(1^\lambda, C)$. For a circuit C such that $\text{size}(C) \leq L$ for $L = L(\lambda)$, the obfuscation

algorithm works as following.

Flex.ObfOff($1^\lambda, 1^L$). The offline phase of obfuscation algorithm does the following.

- Run $(\text{Fixed.Obf}_{\text{off},i}, \text{Fixed.st}_i) \leftarrow \text{Fixed.ObfOff}(1^\lambda, 1^L, \lambda^i)$ for $i \in [\lambda]$.
- Output $\text{obf}_{\text{off}} := \{\text{Fixed.Obf}_{\text{off},i}\}_{i \in [\lambda]}$ and $\text{st} = \{\text{Fixed.st}_i\}_{i \in [\lambda]}$.

Flex.ObfOn(st, C). The online phase of obfuscation algorithm takes as input a circuit $C : [N] \rightarrow \{0, 1\}$ and parses the input as $\text{st} = \{\text{Fixed.st}_i\}_{i \in [\lambda]}$. It then does the following:

- Find $c \in \mathbb{N}$ such that $\lambda^{c-1} < N \leq \lambda^c$.
- Sample $\text{sd} \leftarrow \{0, 1\}^\lambda$
- Construct a circuit $C[\text{sd}] : [\lambda^c] \rightarrow \{0, 1\}$ that is defined as, on input $\mathbf{x} \in [\lambda^c]$, $C[\text{sd}](\mathbf{x}) = \begin{cases} C(\mathbf{x}) & \text{if } \mathbf{x} \leq N \\ \text{PRF}(\text{sd}, \mathbf{x}) & \text{if } N < \mathbf{x} \leq \lambda^c \end{cases}$.
- Compute $\text{Fixed.Obf}_{\text{on},c} \leftarrow \text{Fixed.ObfOn}(\text{Fixed.st}_c, C[\text{sd}])$.
- Output $\text{obf}_{\text{on}} := \text{Fixed.Obf}_{\text{on},c}$.

Flex.Eval(Obf, \mathbf{x}). The evaluation algorithm does the following.

- Parse $\text{Obf} = (\text{obf}_{\text{off}}, \text{obf}_{\text{on}}) = (\{\text{Fixed.Obf}_{\text{off},i}\}_{i \in [\lambda]}, \text{Fixed.Obf}_{\text{on},c})$.
- Construct $\text{Fixed.Obf} = (\text{Fixed.Obf}_{\text{off},c}, \text{Fixed.Obf}_{\text{on},c})$ ¹⁷.
- Output $\text{Fixed.Eval}(\text{Fixed.Obf}, \mathbf{x})$.

Correctness. For $\text{Obf} = (\text{obf}_{\text{off}}, \text{obf}_{\text{on}}) = (\{\text{Fixed.Obf}_{\text{off},i}\}_{i \in [\lambda]}, \text{Fixed.Obf}_{\text{on},c})$, we have $\text{Fixed.Obf}_{\text{off},c} \leftarrow \text{Fixed.ObfOff}(1^\lambda, 1^L, \lambda^c)$ and $\text{Fixed.Obf}_{\text{on},c} \leftarrow \text{Fixed.ObfOn}(\text{Fixed.st}_c, C[\text{sd}])$ where $C[\text{sd}] : [\lambda^c] \rightarrow \{0, 1\}$. From the correctness of the underlying **Fixed.pPRIO** scheme we have

¹⁷We assume that we can obtain c from the length of \mathbf{x} .

$\text{Fixed.Eval}(\text{Fixed.Obf}, \mathbf{x}) = C[\text{sd}](\mathbf{x}) = C(\mathbf{x})$ for any input $\mathbf{x} \in [\lambda^c]$.

Efficiency. It is straightforward to see that the algorithms run in polynomial time in the input length and λ . The only non-trivial part is to bound the running time of $\text{Fixed.ObfOff}(1^\lambda, 1^L, \lambda^i)$ for $i \in [\lambda]$. The running time of this part can be bounded by $\text{poly}(\lambda, L, \log \lambda^i) = \text{poly}(\lambda, L)$ by the efficiency property of the underlying fixed input domain pPRIO as per Definition 6.19. In particular, we have

$$|\text{Obf}_{\text{off}}| = \text{poly}(L, \lambda), \quad |\text{Obf}_{\text{on}}| = \text{poly}(L, \lambda)$$

where $(\text{Obf}_{\text{off}}, \text{Obf}_{\text{on}}) \leftarrow \text{Obf}(1^\lambda, C)$ for circuit $C : [N] \rightarrow [M]$ whose size is bounded by $L = L(\lambda)$.

Security. We prove the security of the above construction. Before doing so, we prove the following useful lemma. The lemma essentially says that if a part of the auxiliary information is pseudorandom in the pre-condition distribution, then it is pseudorandom in the corresponding post-condition distribution where we apply pPRIO. Conceptually similar lemma is proven in Lemma 3.4 of [22] in the context of evasive LWE.

Lemma 6.38. *Let $\text{pPRIO} = (\text{Fixed.ObfOff}, \text{Fixed.ObfOn}, \text{Fixed.Eval})$ be a pPRIO scheme for fixed input domain and Samp be a PPT algorithm that takes as input 1^λ and outputs*

$$\left(1^{NQ}, 1^L, \text{aux} = (\text{aux}_1, \text{aux}_2) \in \{0, 1\}^* \times \mathcal{X}, \{C^j\}_{j \in [Q]}\right)$$

for some set \mathcal{X} . Let us assume that

$$\left((\text{aux}_1, \text{aux}_2), 1^L, \{C^j(i)\}_{i \in [N], j \in [Q]}\right) \approx_c \left((\text{aux}_1, \mathbf{x}), 1^L, \{\Delta_i^j\}_{i \in [N], j \in [Q]}\right)$$

holds for $\mathbf{x} \leftarrow \mathcal{X}, \Delta_i^j \leftarrow \{0, 1\}$ and assume the security of pPRIO with respect to Samp .

We then have

$$\left((\text{aux}_1, \text{aux}_2), \text{obf}_{\text{off}}, \{\text{obf}_{\text{on}}^j\}_{j \in [Q]}\right) \approx_c \left((\text{aux}_1, \mathbf{x}), \text{obf}_{\text{off}}, \{\delta^j\}_{j \in [Q]}\right), \quad (6.51)$$

where we have $(\text{obf}_{\text{off}}, \text{st}) \leftarrow \text{Fixed.ObfOff}(1^\lambda, 1^L, N)$, $\text{obf}_{\text{on}}^j \leftarrow \text{Fixed.ObfOn}(\text{st}, C^j)$ and $\delta^j \leftarrow \mathcal{CT}_{\text{ObfOn}}$ for $j \in [Q]$, where $\mathcal{CT}_{\text{ObfOn}}$ is the set of binary strings with the same length as obf_{on}^j .

Proof. From the assumption, we have

$$\left((\text{aux}_1, \text{aux}_2), 1^L, \{C^j(i)\}_{i \in [N], j \in [Q]} \right) \approx_c \left((\text{aux}_1, \mathbf{x}), 1^L, \{\Delta_i^j\}_{i \in [N], j \in [Q]} \right) \quad (6.52)$$

which implies $(\text{aux}_1, \text{aux}_2, 1^L) \approx_c (\text{aux}_1, \mathbf{x}, 1^L)$. This further implies

$$(\text{aux}_1, \text{aux}_2, 1^L, \{\Delta_i^j\}_{i \in [N], j \in [Q]}) \approx_c (\text{aux}_1, \mathbf{x}, 1^L, \{\Delta_i^j\}_{i \in [N], j \in [Q]}) \quad (6.53)$$

since adding independently sampled random terms does not make the task of indistinguishing the distributions easier. Equation (6.52) and Equation (6.53) implies $(\text{aux}_1, \text{aux}_2, 1^L, \{C^j(i)\}_{i \in [N], j \in [Q]}) \approx_c (\text{aux}_1, \text{aux}_2, 1^L, \{\Delta_i^j\}_{i \in [N], j \in [Q]})$. Applying pPRIO security definition with respect to Samp, we get

$$\left(\text{aux}_1, \text{aux}_2, \text{obf}_{\text{off}}, \{\text{obf}_{\text{on}}^j\}_{j \in [Q]} \right) \approx_c \left(\text{aux}_1, \text{aux}_2, \text{obf}_{\text{off}}, \{\delta^j\}_{j \in [Q]} \right).$$

Recall that we have $(\text{aux}_1, \text{aux}_2, 1^L) \approx_c (\text{aux}_1, \mathbf{x}, 1^L)$. This further implies

$$\left(\text{aux}_1, \text{aux}_2, \text{obf}_{\text{off}}, \{\delta^j\}_{j \in [Q]} \right) \approx_c \left(\text{aux}_1, \mathbf{x}, \text{obf}_{\text{off}}, \{\delta^j\}_{j \in [Q]} \right), \quad (6.54)$$

since obf_{off} can be sampled independently given $1^L, N$ and $\{\delta^j\}_{j \in [Q]}$.

From Section 6.6.3 and Equation (6.54) we deduce

$$\left(\text{aux}_1, \text{aux}_2, \text{obf}_{\text{off}}, \{\text{obf}_{\text{on}}^j\}_{j \in [Q]} \right) \approx_c \left(\text{aux}_1, \mathbf{x}, \text{obf}_{\text{off}}, \{\delta^j\}_{j \in [Q]} \right)$$

and hence the lemma. ■

Following theorem asserts the security of our construction.

Theorem 6.39. *Our pPRIO scheme for flexible input domain is secure.*

Proof. Let us consider a sampler Samp that outputs $(1^{\sum_{j \in [Q]} N_j}, 1^L, \text{aux}, \{C^j\}_{j \in [Q]})$.

To prove the theorem, we show that

$$(\text{aux}, \text{obf}_{\text{off}}, \text{obf}_{\text{on}}^1, \dots, \text{obf}_{\text{on}}^Q) \approx_c (\text{aux}, \text{obf}_{\text{off}}, \delta^1, \dots, \delta^Q), \quad (6.55)$$

holds assuming

$$(\text{aux}, \{C^1(i)\}_{i \in [N_1]}, \dots, \{C^Q(i)\}_{i \in [N_Q]}) \approx_c (\text{aux}, \{\Delta_i^1\}_{i \in [N_1]}, \dots, \{\Delta_i^Q\}_{i \in [N_Q]})$$

where

$$\text{obf}_{\text{off}} = \{\text{Fixed.Obf}_{\text{off},i}\}_{i \in [\lambda]}, (\text{Fixed.Obf}_{\text{off},i}, \text{Fixed.st}_i) \leftarrow \text{Fixed.ObfOff}(1^\lambda, 1^L, \lambda^i) \text{ for } i \in [\lambda],$$

$$\text{obf}_{\text{on}}^1 = \text{Fixed.Obf}_{\text{on},c_j},$$

$$\text{Fixed.Obf}_{\text{on},c_j} \leftarrow \text{Fixed.ObfOn}(\text{Fixed.st}_{c_j}, C^j[\text{sd}^j]) \text{ for } c_j \in \mathbb{N} \text{ s.t. } \lambda^{c_j-1} < N_j \leq \lambda^{c_j},$$

$$\delta^j \leftarrow C\mathcal{T}^j \text{ where } C\mathcal{T}^j \text{ is the set of binary strings with same length as } \text{obf}_{\text{on}}^j, \text{ for } j \in [Q]$$

$$\Delta_i^j \leftarrow \{0, 1\} \text{ for } j \in [Q], i \in [N_j]$$

We define S_k to be $S_k := \{j \in [Q] : \lambda^{k-1} < N_j \leq \lambda^k\}$. Since Samp is a PPT algorithm, there exists a constant c_{\max} such that $\lambda^{c_{\max}-1} < \max_{j \in [Q]} N_j \leq \lambda^{c_{\max}}$ and therefore we have $[Q] = \cup_{k \in [c_{\max}]} S_k$. Rearranging the terms, it suffices to show that

$$(\text{aux}, \text{obf}_{\text{off}}, \{\text{obf}_{\text{on}}^j\}_{j \in S_1}, \dots, \{\text{obf}_{\text{on}}^j\}_{j \in S_{c_{\max}}}) \approx_c (\text{aux}, \text{obf}_{\text{off}}, \{\delta^j\}_{j \in S_1}, \dots, \{\delta^j\}_{j \in S_{c_{\max}}})$$

holds assuming

$$\begin{aligned} & (\text{aux}, \{C^j(i)\}_{j \in S_1, i \in [N_j]}, \dots, \{C^j(i)\}_{j \in S_{c_{\max}}, i \in [N_j]}) \\ & \approx_c (\text{aux}, \{\Delta_i^j\}_{j \in S_1, i \in [N_j]}, \dots, \{\Delta_i^j\}_{j \in S_{c_{\max}}, i \in [N_j]}). \end{aligned} \quad (6.56)$$

To prove this, we prove the following lemma.

Lemma 6.40. *For $c \in [0, c_{\max}]$, let us consider the following statement:*

$$\left(\begin{array}{c} \text{aux}, \text{obf}_{\text{off}}, \{\text{obf}_{\text{on}}^j\}_{j \in S_1}, \dots, \{\text{obf}_{\text{on}}^j\}_{j \in S_c}, \\ \{C^j[\text{sd}^j](i)\}_{j \in S_{c+1}, i \in [\lambda^{c+1}]}, \dots, \{C^j[\text{sd}^j](i)\}_{j \in S_{c_{\max}}, i \in [\lambda^{c_{\max}}]} \end{array} \right)$$

$$\approx_c \left(\text{aux}, \text{obf}_{\text{off}}, \{\delta^j\}_{j \in S_1}, \dots, \{\delta^j\}_{j \in S_c}, \{\Delta_i^j\}_{j \in S_{c+1}, i \in [\lambda^{c+1}]}, \dots, \{\Delta_i^j\}_{j \in S_{c_{\max}}, i \in [\lambda^{c_{\max}}]} \right). \quad (6.57)$$

Then, assuming the security of pPRIO, the above computational indistinguishability for $c = c^*$ implies that for $c^* + 1$ for all $c^* \in [0, c_{\max} - 1]$.

Proof. By setting

$$\begin{aligned} \text{aux}_1 &:= \left(\text{aux}, \{\text{Fixed.Obf}_{\text{off},k}\}_{k \in [\lambda] \setminus \{c^*+1\}} \right), \\ \text{aux}_2 &:= \left(\begin{array}{c} \{\text{obf}_{\text{on}}^j\}_{j \in S_1}, \dots, \{\text{obf}_{\text{on}}^j\}_{j \in S_{c^*}}, \\ \{C^j[\text{sd}^j](i)\}_{j \in S_{c^*+2}, i \in [\lambda^{c^*+2}]}, \dots, \{C^j[\text{sd}^j](i)\}_{j \in S_{c_{\max}}, i \in [\lambda^{c_{\max}}]} \end{array} \right), \\ \text{Circuits} &:= \{C^j[\text{sd}^j]\}_{j \in S_{c^*+1}} \end{aligned}$$

Equation (6.57) with $c = c^*$ implies

$$\left((\text{aux}_1, \text{aux}_2), \{C^j[\text{sd}^j](i)\}_{j \in S_{c^*+1}, i \in [\lambda^{c^*+1}]} \right) \approx_c \left((\text{aux}_1, \mathbf{x}), \{\Delta_i^j\}_{j \in S_{c^*+1}, i \in [\lambda^{c^*+1}]} \right),$$

where \mathbf{x} is a random string with the same length as aux_2 . By applying Theorem 6.38, we have

$$\left((\text{aux}_1, \text{aux}_2), \text{Obf}_{\text{off}, c^*+1}, \{\text{obf}_{\text{on}}^j\}_{j \in S_{c^*+1}} \right) \approx_c \left((\text{aux}_1, \mathbf{x}), \text{Obf}_{\text{off}, c^*+1}, \{\delta^j\}_{j \in S_{c^*+1}} \right).$$

By rearranging the terms, we can observe that the above equation is equivalent to Equation (6.57) with $c = c^* + 1$. ■

We then conclude the proof of Theorem 6.39 using Theorem 6.40. Our goal is to prove Equation (6.57) with $c = c_{\max}$, since it is equivalent to Equation (6.55). This is proven by the induction, where the induction step is already provided by Theorem 6.40. Therefore, it suffices to prove Equation (6.57) with $c = 0$. This immediately follows from Equation (6.56) and by the security of PRF, since the LHS distribution of Equation (6.57) with $c = 0$ is obtained by padding the LHS distribution of Equation (6.56) with PRF

values $\{\text{PRF}(\text{sd}^j, i)\}_{c \in [c_{\max}], j \in S_c, i \in [N_j+1, \lambda^c]}$, which are pseudorandom. This completes the proof of Theorem 6.39. \blacksquare

Instantiation. Instantiating the underlying pPRIO scheme as in Theorem 6.37, we get the following theorem.

Theorem 6.41. *Assuming evasive LWE (Assumption 2.6) and LWE (Assumption 2.5), there exists a secure pPRIO scheme for flexible input domain supporting circuits of bounded size $L = \text{poly}(\lambda)$ with*

$$|\text{Obf}_{\text{off}}| = \text{poly}(L, \lambda), \quad |\text{Obf}_{\text{on}}| = \text{poly}(L, \lambda).$$

6.6.4 Extending the Output Length.

In our construction of pPRIO with flexible input space in Section 6.6.3, we only consider the case of the output space of the obfuscated circuit being $\{0, 1\}$. We can extend the output space to be $\{0, 1\}^{\ell_{\text{out}}}$ for arbitrary polynomial ℓ_{out} as follows. We do not change the off-line phase of the obfuscation algorithm. When we run on-line phase of the obfuscation algorithm, we derive sub-circuits $C_1, \dots, C_{\ell_{\text{out}}}$ of C with binary outputs, where C_j takes on input i and outputs the j -th bit of $C(i)$. We then run $\text{ObfOn}(\text{st}, C_j)$ for all $j \in [\ell_{\text{out}}]$ and output them as the obfuscation of C . The evaluation algorithm then recovers each output separately and combines them. It is straightforward to see that the security of the construction is preserved with this modification. As for the efficiency, we still maintain the asymptotic efficiency of

$$|\text{Obf}_{\text{off}}| = \text{poly}(L, \lambda), \quad |\text{Obf}_{\text{on}}| = \ell_{\text{out}} \text{poly}(L, \lambda) = \text{poly}(L, \lambda)$$

where $(\text{Obf}_{\text{off}}, \text{Obf}_{\text{on}}) \leftarrow \text{Obf}(1^\lambda, C)$ for circuit C whose size is bounded by $L = L(\lambda)$.

6.7 INSTANTIATING THE RANDOM ORACLES USING prIO

Hohenberger, Sahai, and Waters [119] show that some applications of random oracles (ROM) can be made secure in the standard model by instantiating the hash functions using iO in a specific manner. In this section, we discuss that we can replace full-fledged iO with prIO in these applications. As a concrete example, we show that full-domain hash (FDH) signatures can be proven secure in the standard model if we instantiate the hash function using prIO in place of iO . We also refer to Section 6.8 for the application of prIO for instantiating random oracle in Sakai-Ohgishi-Kasahara ID-based NIKE.

6.7.1 Full-Domain Hash Signatures (Selectively Secure) from prIO

Ingredients. We make use of the following ingredients.

1. A one-way trapdoor permutation family (TDP).
2. Punctured PRFs
3. A prIO scheme.
 1. **Setup**(1^λ) : The setup algorithm takes as input the security parameter and does the following:
 - It runs the setup of the TDP to obtain a public index PK along with a trapdoor SK, yielding the map $g_{PK} : \{0, 1\}^n \rightarrow \{0, 1\}^n$ together with its inverse g_{SK}^{-1} .
 - It chooses a puncturable PRF key K for F where $F(K, \cdot) : \{0, 1\}^\lambda \rightarrow \{0, 1\}^n$. Then, it creates a prIO obfuscation of the of the program Full Domain Hash Figure 6.2. We refer to the obfuscated program as the function $H : \{0, 1\}^\lambda \rightarrow \{0, 1\}^n$. We need the truth table of the PRF to be pseudorandom against an adversary whose size is polynomial in κ , where κ is the parameter specified by our prIO. This can be achieved assuming the subexponential security for the PRF.
 - It outputs the verification key VK as the trapdoor index PK as well as the hash function $H(\cdot)$. The secret key is the trapdoor SK.
 2. **Sign**(SK, m): Output $\sigma = g_{SK}^{-1}(H(m))$.
 3. **Verify**(VK, m, σ) : Check if $g_{PK}(\sigma) = H(m)$ and output “Accept” if and only if this is true.

Correctness follows from the correctness of the TDP. We sketch security next.



Figure 6.2: Full Domain Hash

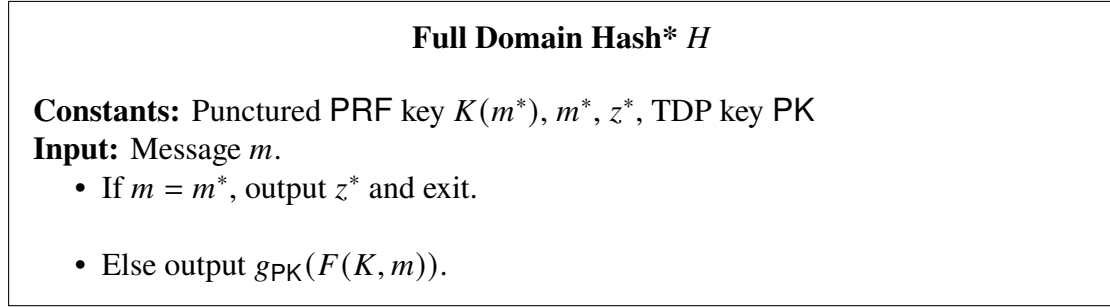


Figure 6.3: Full Domain Hash*

Security The security proof closely resembles that of [119], except that it can be simpler since the security guarantee of prIO is stronger than iO for the specific case of pseudorandom functionalities. In particular, prIO does not require two circuits to have identical truth tables in order to guarantee security, but allows different truth tables so long as they are pseudorandom (given auxiliary information).

Theorem 6.42. *If the obfuscation scheme in Section 6.5.2 is secure as per Definition 6.14 with respect to a parameter κ and samplers that output circuits defined as Figure 6.2 and 6.3, F is subexponentially secure punctured PRF, and the trapdoor permutation scheme TDP is one-way, then the above signature scheme is selectively secure.*

Proof. The proof follows a similar sequence of hybrids as [119] ¹⁸. In the first hybrid we move to using the obfuscation of the circuit in Figure 6.3 with z^* being a random point in $\{0, 1\}^n$. Since the truth tables of the two programs are pseudorandom given PK and SK (even against an adversary that runs in polynomial time in κ), indistinguishability holds by the guarantee of prIO . This allows to reduce the security to that of the TDP, exactly

¹⁸We can skip Hybrid 1 in [119] because we don't need the programs to be exactly equivalent as discussed above.

as in [119] since a valid signature would imply a preimage to z^* and other signatures can be simulated using the PRF key. ■

6.7.2 Discussion about Other Applications.

Hohenberger, Sahai, and Waters [119] demonstrate that the selective security of the full-domain hash (FDH) signature based on trapdoor permutations (TDP), the adaptive security of RSA FDH signatures [79], the selective security of BLS signatures, and the adaptive security of BLS signatures [45] can be proven in the standard model by carefully instantiating the underlying hash function by iO for each application. As shown in Section 6.7.1, the random oracle in the FDH signature can be instantiated using prIO instead of full-fledged iO . Similarly, we can instantiate the random oracle in selectively secure BLS signatures with prIO , following a strategy similar to that in [119]. At a high level, these proofs follow those in the random oracle model (ROM), where we use iO to obfuscate a derandomized version of the simulator for the hash function in ROM-based proofs. In these settings, the truth table of the simulated hash function is pseudorandom, allowing us to follow the same proof strategy using prIO .

For adaptively secure RSA FDH and BLS signatures, the situation is different. In these cases, Hohenberger et al. adopt an alternative proof strategy that deviates from the high level strategy of obfuscating the simulator for the proof in the ROM. This is due to the fact that the original proofs [45, 79] are incompatible with the conditions required for using iO , where the truth table of the hash functions must remain unchanged across game hops. To be compatible with iO , they introduce a structure for the hash function, making its truth table no longer pseudorandom. This prevents us from replacing the hash function with prIO following their approach.

To instantiate the hash function with prIO , we revert to the original ROM-based proof strategy [45, 79]. Unlike the iO -based approach, prIO -based proof does not require the truth table of the hash function to remain unchanged across game hops; it only requires

the truth table to be pseudorandom. This relaxed condition enables the use of the original ROM security proofs. We omit the details on these applications, since they are similar to and simpler than our random oracle instantiation for the Sakai-Ohgishi-Kasahara ID-based NIKE in Section 6.8.

6.8 ID-BASED NON-INTERACTIVE KEY EXCHANGE

In this section we show our construction of ID-NIKE scheme. The construction is the same as the ID-based NIKE system by Sakai, Ohgishi, and Kasahara [159] except that the hash function modeled as random oracle is replaced with an obfuscation of a PRF. This leads to the first instantiation of ID-NIKE without multi-linear maps [88] or indistinguishability obfuscation in the standard model.

6.8.1 Construction

Building Blocks. We require the following building blocks for our construction.

1. We use prIO scheme (iO, Eval) with input space $\mathcal{ID} \stackrel{\text{def}}{=} \{0, 1\}^n$ that is secure as per Definition 6.14 with parameter κ . By using our construction in Section 6.5.2 instantiated by mi-prFE construction in Section 6.3.2, we have $\kappa = \lambda^{n^2 \log \lambda}$.
2. Symmetric pairing $\mathcal{G} = (\mathbb{G}, \mathbb{G}_T, p, e, g)$ with prime order p that is equipped with an efficiently computable function $\text{MapToPoint} : \mathcal{D}_{\text{MTP}} \rightarrow \mathbb{G}$, where \mathcal{D}_{MTP} is an efficiently samplable domain. We assume that the DBDH assumption (Assumption 6.16) holds on \mathcal{G} .
We need that there is an efficiently computable *randomized* function MapToPoint^{-1} satisfying $\text{MapToPoint}(\text{MapToPoint}^{-1}(h)) = h$ for all $h \in \mathbb{G}$. We also need $\text{MapToPoint}^{-1}(h) \equiv x$ for $x \leftarrow \mathcal{D}_{\text{MTP}}$ and $h \leftarrow \mathbb{G}$. We denote the randomness space of the algorithm MapToPoint^{-1} by \mathcal{R}_{MTP} .
3. A subexponentially secure pseudorandom function $\text{PRF} : \{0, 1\}^\Lambda \times \mathcal{ID} \rightarrow \mathcal{D}_{\text{MTP}}$. Here, $\Lambda = \Lambda(\lambda)$ is set so that $2^{\Lambda^\delta} > \kappa^{\omega(1)}$ holds, where δ is defined as a constant such that there is no adversary with size 2^{λ^δ} and distinguishing advantage $2^{-\lambda^\delta}$ against PRF for all sufficiently large λ . An example choice would be to take $\Lambda := (n^2 \lambda)^{1/\delta}$.

We describe the construction below.

Setup(1^λ): On input the security parameter 1^λ , it chooses the description of symmetric

pairing groups $\mathcal{G} = (\mathbb{G}, \mathbb{G}_T, p, e, g)$ with prime order $p > 2^{2\lambda}$, a PRF key $\text{sd} \leftarrow \{0, 1\}^\Lambda$, and $\alpha \leftarrow \mathbb{Z}_p$ and computes $\tilde{C} \leftarrow iO(\text{PRF}(\text{sd}, \cdot))$. Here, the obfuscated circuit $\text{PRF}(\text{sd}, \cdot)$ is appropriately padded so that its size is the same as the circuit $F[\overline{\text{sd}}, \widetilde{\text{sd}}, g^\beta, g^\gamma]$ described in Figure 6.4. Finally, it outputs $\text{mpk} = (\mathcal{G}, \tilde{C})$ and $\text{msk} = \alpha$.

Ext(mpk, msk, id): On input $\text{mpk} = (\mathcal{G}, \tilde{C})$, $\text{msk} = \alpha$, and an identity id , it computes $H(\text{id})$, where the hash function $H : \mathcal{ID} \rightarrow \mathbb{G}$ is defined as $H(\text{id}) \stackrel{\text{def}}{=} \text{MapToPoint}(\tilde{C}(\text{id}))$. It then computes and outputs $\text{usk}_{\text{id}} = H(\text{id})^\alpha$.

Share(mpk, usk_{id_1} , id_2): Given the master public key mpk , a user secret key $\text{usk}_{\text{id}_1} = H(\text{id}_1)^\alpha$, and an identity $\text{id}_2 \in \mathcal{ID}$, it computes and outputs $K \stackrel{\text{def}}{=} e(\text{usk}_{\text{id}_1}, H(\text{id}_2))$ as the shared key.

Correctness. The correctness of the construction can be seen by $e(\text{usk}_{\text{id}_1}, H(\text{id}_2)) = e(H(\text{id}_1), H(\text{id}_2))^\alpha = e(\text{usk}_{\text{id}_2}, H(\text{id}_1))$.

6.8.2 Security Proof

Theorem 6.43. *The construction above is secure as per Definition 6.8 assuming that the DBDH assumption over \mathcal{G} holds (See Assumption 6.16), prIO is secure with parameter κ (as per Definition 6.14) and with respect to a sampler specified in the proof of Theorem 6.44, and PRF is subexponentially secure.*

Proof. We prove the security of the scheme via the following hybrids. Let us fix a PPT adversary \mathcal{A} . We denote the advantage of the adversary by ϵ and without loss of generality, we assume that \mathcal{A} makes exactly Q key extraction queries. For the sake of contradiction, let us assume that ϵ is non-negligible. In the following, we denote the advantage of the adversary in **Game-xx** by ϵ_{xx} .

Game-0. This is the security game for ID-NIKE. By definition, we have $\epsilon_0 = \epsilon$.

Game-1. In this game, we sample random function $R : \mathcal{ID} \rightarrow [2Q]$ at the end of the game independently from anything else in the game. We then change the adversary \mathcal{A} so that it outputs a random coin as its guess if the following does *not* hold:

$$R(\text{id}_1^*) = 1 \wedge R(\text{id}_2^*) = 2 \wedge R(\text{id}_1) \notin \{1, 2\} \wedge \dots \wedge R(\text{id}_Q) \notin \{1, 2\}, \quad (6.58)$$

where id_1^* and id_2^* are the challenge identities and $\text{id}_1, \dots, \text{id}_Q$ are the identities for which key extraction queries are made by \mathcal{A} . If the abort condition above is not satisfied, the output of \mathcal{A} is unchanged. Since Equation (6.58) holds with probability $\frac{1}{4Q^2} \cdot \left(1 - \frac{2}{2Q}\right)^Q = \Theta(\epsilon_0/Q^2)$, we have $\epsilon_1 = \Theta(\epsilon_0/Q^2)$.

Game-2. In this game, we choose R at the beginning of the game and stop the game immediately and force \mathcal{A} to output random guess if \mathcal{A} makes a key extraction query or challenge query that violates Equation (6.58). With this change, the distribution of the output by \mathcal{A} is unchanged and therefore we have $\epsilon_2 = \epsilon_1$.

Game-3. In this game, we replace R with a pseudorandom function $\overline{\text{PRF}} : \{0, 1\}^\lambda \times \mathcal{ID} \rightarrow [2Q]$. In more detail, the game samples $\overline{\text{sd}} \leftarrow \{0, 1\}^\lambda$ and uses $\overline{\text{PRF}}(\overline{\text{sd}}, \cdot)$ in place of the function $R(\cdot)$ when we check the abort condition. It is straightforward to see that $|\epsilon_3 - \epsilon_2| \leq \text{negl}(\lambda)$ by the security of PRF.

Game-4. In this game, we change how \tilde{C} is computed. In more detail, we sample $\beta, \gamma \leftarrow \mathbb{Z}_p$ and $\widetilde{\text{sd}} \leftarrow \{0, 1\}^\Lambda$ at the beginning of the game. We then set $F[\overline{\text{sd}}, \widetilde{\text{sd}}, g^\beta, g^\gamma]$ as Figure 6.4, where we use yet another pseudorandom function $\widetilde{\text{PRF}} : \{0, 1\}^\Lambda \times \mathcal{ID} \rightarrow \mathcal{R}_{\text{MTP}} \times \mathbb{Z}_p$ in the description of the circuit. In this game, \tilde{C} is computed as $\tilde{C} \leftarrow iO(F[\overline{\text{sd}}, \widetilde{\text{sd}}, g^\beta, g^\gamma])$. As we will show in Theorem 6.44, we have $|\epsilon_4 - \epsilon_3| \leq \text{negl}(\lambda)$ by the security of PRIO and by the subexponential security of PRF and $\widetilde{\text{PRF}}$.

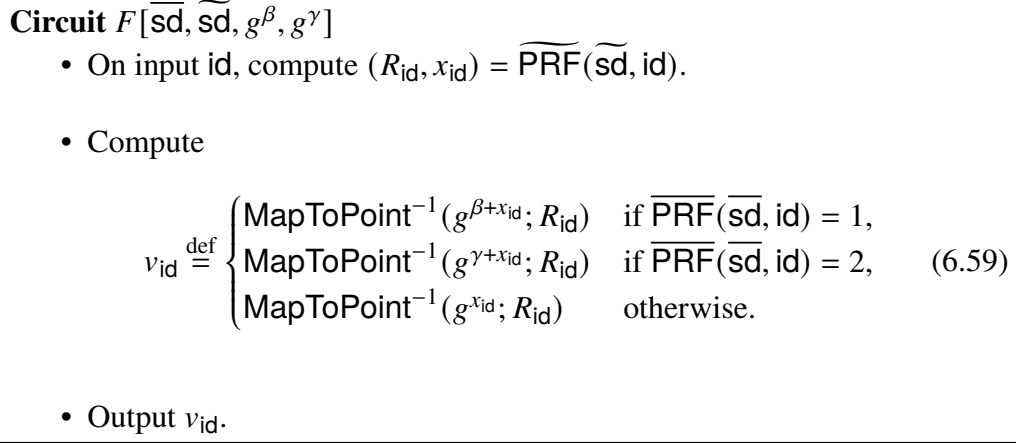


Figure 6.4: Description of the circuit $F[\overline{\text{sd}}, \widetilde{\text{sd}}, g^\beta, g^\gamma]$.

Game-5. In this game, we no longer use the exponents β and γ explicitly in answering the key extraction queries. In more detail, when the adversary makes a key extraction query to id , we proceed as follows. First, we check whether $\overline{\text{PRF}}(\overline{\text{sd}}, \cdot) \notin \{1, 2\}$ or not. If it does not hold, we abort the game as specified in **Game-2**. Otherwise, we compute $(R_{\text{id}}, x_{\text{id}}) = \widetilde{\text{PRF}}(\widetilde{\text{sd}}, \text{id})$ and return $\text{usk}_{\text{id}} \stackrel{\text{def}}{=} (g^\alpha)^{x_{\text{id}}}$ as the secret key.

We claim that this change does not alter the view of the adversary. To see this, we observe that

$$\begin{aligned} \text{H}(\text{id}) &= \text{MapToPoint}(\tilde{C}(\text{id})) = \text{MapToPoint}(F[\overline{\text{sd}}, \widetilde{\text{sd}}, g^\beta, g^\gamma](\text{id})) \\ &= \text{MapToPoint}(\text{MapToPoint}^{-1}(g^{x_{\text{id}}}; R_{\text{id}})) = g^{x_{\text{id}}} \end{aligned}$$

holds, where the first equation holds by the definition of H , the second by the correctness of $i\mathcal{O}$, the third by $\overline{\text{PRF}}(\overline{\text{sd}}, \text{id}) \notin \{1, 2\}$ and Equation (6.59), and the fourth by the property of MapToPoint^{-1} . This implies $(g^\alpha)^{x_{\text{id}}} = \text{H}(\text{id})^\alpha$ as desired. We therefore have $\epsilon_5 = \epsilon_4$.

Game-6. In this game, we change the way K^* is computed when $\text{coin} = 0$. In particular, we compute $T = e(g, g)^{\alpha\beta\gamma}$ at the beginning of the game and when the adversary

asks for the challenge key, we return

$$K^* = T \cdot e(g^\alpha, g^\beta)^{x_{\text{id}_2^*}} \cdot e(g^\alpha, g^\gamma)^{x_{\text{id}_1^*}} \cdot e(g, g^\alpha) \quad (6.60)$$

to the adversary, where $(R_{\text{id}_i^*}, x_{\text{id}_i^*}, y_{\text{id}_i^*}) = \widetilde{\text{PRF}}(\widetilde{\text{sd}}, \text{id}_i^*)$ for $i = 1, 2$.

We claim that this change does not alter the view of the adversary. To see this, recall that the game aborts unless $\overline{\text{PRF}}(\overline{\text{sd}}, \text{id}_1^*) = 1 \wedge \overline{\text{PRF}}(\overline{\text{sd}}, \text{id}_2^*) = 2$. The first condition implies $H(\text{id}_1^*) = \text{MapToPoint}(v_{\text{id}_1^*}) = g^{\beta+x_{\text{id}_1^*}}$ by Equation (6.59) and by the property of MapToPoint^{-1} . Similarly, we have $H(\text{id}_2^*) = g^{\gamma+x_{\text{id}_2^*}}$. Thus, we have

$$e(H(\text{id}_1^*), H(\text{id}_2^*))^\alpha = e(g^{\beta+x_{\text{id}_1^*}}, g^{\gamma+x_{\text{id}_2^*}})^\alpha = T \cdot e(g^\alpha, g^\beta)^{x_{\text{id}_2^*}} \cdot e(g^\alpha, g^\gamma)^{x_{\text{id}_1^*}} \cdot e(g, g^\alpha)^{x_{\text{id}_1^*} x_{\text{id}_2^*}}$$

as desired. We therefore have $\epsilon_6 = \epsilon_5$.

Note that in this game, we do not need to know α any more to efficiently simulate the game. Rather, it suffices to know $g^\alpha, g^\beta, g^\gamma$, and $T = e(g, g)^{\alpha\beta\gamma}$.

Game-7. In this game, we replace T with random $T \leftarrow \mathbb{G}_T$. We claim that the adversary cannot distinguish this game from the previous game. This can be seen by straightforward reduction to the DBDH assumption, where we simulate **Game-6** if $T = e(g, g)^{\alpha\beta\gamma}$ and **Game-7** if $T \leftarrow \mathbb{G}_T$. Indeed, we only need to know g^β and g^γ for simulating mpk , g^α for answering the key extraction queries, and $g^\alpha, g^\beta, g^\gamma$, and T for simulating the challenge key. We therefore have $|\epsilon_6 - \epsilon_7| \leq \text{negl}(\lambda)$.

Game-8. In this game, we change the challenge shared key K^* to be random group element in \mathbb{G}_T , regardless of whether $\text{coin} = 0$ or $\text{coin} = 1$. We claim that this is a conceptual change. This can be easily seen by observing that K^* computed as Equation (6.60) is distributed uniformly at random over \mathbb{G}_T if so is T . Therefore,

we have $\epsilon_8 = \epsilon_7$.

Clearly, we have $\epsilon_8 = 0$, since no information of `coin` is leaked to the adversary in **Game-8**. However, from the above discussion, we have $0 = \epsilon_8 \geq \Theta(\epsilon)/Q^2 - \text{negl}(\lambda)$ by the triangle inequality. This contradicts our assumption that ϵ is non-negligible. ■

To complete the proof of Theorem 6.43, it remains to prove Theorem 6.44.

Lemma 6.44. *Assuming that iO is a secure PRIO scheme as per Definition 6.14 with the parameter κ and PRF and $\widetilde{\text{PRF}}$ are subexponentially secure, we have $|\epsilon_3 - \epsilon_4| \leq \text{negl}(\lambda)$.*

Proof. To prove the lemma, it suffices to show $(iO(\text{PRF}(\text{sd}, \cdot)), \overline{\text{sd}}, \mathcal{G}) \approx_c (iO(F[\overline{\text{sd}}, \widetilde{\text{sd}}, g^\beta, g^\gamma]), \overline{\text{sd}}, \mathcal{G})$, since we can simulate the game in a way that we simulate **Game-3** if the given terms come from the left distribution and **Game-4** otherwise. Any adversary distinguishing the games can be turned into an adversary that distinguishes the distributions. By considering a sampler `Samp` that takes as input 1^λ and outputs the auxiliary information $\text{aux} \stackrel{\text{def}}{=} (\overline{\text{sd}}, \mathcal{G})$ and circuits $C_0 \stackrel{\text{def}}{=} \text{PRF}(\text{sd}, \cdot)$ and $C_1 \stackrel{\text{def}}{=} F[\overline{\text{sd}}, \widetilde{\text{sd}}, g^\beta, g^\gamma]$ and invoking the security of PRIO, we can see that it suffices to prove

$$\begin{aligned} (1^\kappa, \text{aux}, \{\text{PRF}(\text{sd}, \text{id})\}_{\text{id} \in \mathcal{ID}}) &\approx_c (1^\kappa, \text{aux}, \{\delta_{\text{id}} \leftarrow \mathcal{D}_{\text{MTP}}\}_{\text{id} \in \mathcal{ID}}) \\ &\approx_c \left(1^\kappa, \text{aux}, \{v_{\text{id}} = F[\overline{\text{sd}}, \widetilde{\text{sd}}, g^\beta, g^\gamma]\}_{\text{id} \in \mathcal{ID}} \right). \end{aligned} \quad (6.61)$$

The former indistinguishability holds by the subexponential security of PRF. In particular, by our choice of the parameter $2^{\Lambda^\delta} > \kappa^{\omega(1)}$, we can conclude that the adversary with running time $\text{poly}(\kappa)$ cannot distinguish the distributions with advantage more than $\text{negl}(\kappa)$. To prove the latter indistinguishability, we further consider the following distributions. In the following, we do not change the distribution of $\text{aux} = (\overline{\text{sd}}, \mathcal{G})$ and only focus on the distribution of $\{v_{\text{id}}\}_{\text{id}}$.

- This is the rightmost distribution of Equation (6.61), where v_{id} is computed as Equation (6.59) for $(R_{\text{id}}, x_{\text{id}}) = \widetilde{\text{PRF}}(\text{id})$.
- This is the same as the previous distribution except that $(R_{\text{id}}, x_{\text{id}})$ that is used for

computing v_{id} is chosen uniformly at random from their respective domains.

- $\{v_{\text{id}}\}_{\text{id}}$, where $v_{\text{id}} \leftarrow \mathcal{D}_{\text{MTP}}$. Observe that this is the middle distribution in Equation (6.61).

It suffices to prove that the first and the third distributions are indistinguishable given $(1^\kappa, \text{aux})$. We first observe that the first and the second distributions are computationally indistinguishable by the subexponential security of $\widetilde{\text{PRF}}$ and our choice of parameter Λ , where we have $2^{\Lambda^\delta} > \kappa^{\omega(1)}$. To conclude the proof, we show that the second and the third distributions are actually the same distribution. In particular, we show that v_{id} is uniformly distributed over \mathcal{D}_{MTP} for each id in the second distribution. Here, we consider the case of $\widetilde{\text{PRF}}(\overline{\text{sd}}, \text{id}) = 1$. The other cases follow similarly. By the assumption, we know that $g^{\beta+x_{\text{id}}}$ is distributed uniformly at random over \mathbb{G} . Then, by the property of MapToPoint^{-1} and by the fact that R_{id} is random, v_{id} is uniformly distributed over \mathcal{D}_{MTP} as desired. ■

Theorem 6.45. *Let $\kappa = \lambda^{n^2 \log \lambda}$. Assuming non-uniform κ -evasive LWE (Assumption 6.14), subexponentially secure PRF against non-uniform adversary, non-uniform sub-exponential LWE, and the DBDH assumption, there exists a secure ID-NIKE scheme.*

APPENDIX

6.A PSEUDORANDOM FE WITH STRONGER SECURITY

In this section, we explain how to extend [12] construction to achieve strengthened security notion for pseudorandom FE that we call non-uniform κ -prCT security (Definition 6.3), which is required for many of our applications. The stronger version of the security is not proven in [12], but will be required for the construction of mi-prFE in Section 6.3.

6.A.1 Proof for Non-Uniform κ -prCT Security

Theorem 6.46. *Let $\kappa = 2^{\lambda^c}$ for some constant c . Assuming non-uniform κ -evasive LWE (Assumption 6.14), subexponentially secure PRF against non-uniform adversary,*

and non-uniform sub-exponential LWE (Assumption 2.5), there exists a prFE scheme satisfying κ - prCT security as per Definition 6.3.

Proof. We prove that the construction in Section 5.3.2 with λ being replaced by appropriately chosen $\Lambda = \text{poly}(\lambda)$ satisfies the security notion. The reason why we need this scaled version of the security parameter is that we have to consider an adversary whose running time can be κ , which is exponential in λ . In particular, in the security proof, we require LWE and PRF to be secure even against an adversary that runs polynomial time in κ . To handle such an adversary, we rely on the subexponential security of LWE and PRF . By our assumption, there exists $0 < \delta < 1$ such that there is no adversary with size 2^{λ^δ} and distinguishing advantage $2^{-\lambda^\delta}$ against LWE and PRF for all sufficiently large λ . To satisfy the requirement, we generate PRF and LWE instances with respect to a larger security parameter Λ that satisfies $2^{\Lambda^\delta} \geq \kappa^{\omega(1)}$. An example choice of the parameter would be $\Lambda := \lambda^{(c+1)/\delta}$.

The overall structure of the proof is the same as that of Theorem 5.15.

We start with a sampler $\text{Samp}_{\text{prFE}}$ and an adversary \mathcal{A}_1 satisfying $\text{Size}(\text{Samp}_{\text{prFE}}) \leq \text{poly}(\lambda')$ and $\text{Size}(\mathcal{A}_1) \leq \text{poly}(\kappa)$ for $\lambda' < \kappa$.¹⁹ We denote the size of \mathcal{A}_1 by t and the distinguishing advantage for the distributions in Equation (5.19) by ϵ . Assuming non-uniform κ -evasive LWE with respect to Samp defined from $\text{Samp}_{\text{prFE}}$ as in the proof of Theorem 5.15, we obtain an adversary \mathcal{A}_0 whose size is $Q(\lambda')t$ and the distinguishing advantage against the distributions in Equation (5.22) is $\epsilon/Q(\lambda') - \text{negl}(\kappa)$ for some polynomial Q by applying Theorem 6.15. We then consider the same sequence of hybrids as that for the proof of Theorem 5.15. Note that here, the security parameter for the construction λ is replaced by Λ and Q_{msg} and Q_{key} are bounded by $\text{poly}(\lambda')$, since the size of the sampler is $\text{poly}(\lambda')$. By the definition of the hybrids, the adversary has the distinguishing advantage $\epsilon/Q(\lambda') - \text{negl}(\kappa)$ for Hyb_0 and Hyb_8 . Furthermore, we argue

¹⁹Here, we deviate from our convention that the adversary runs in time polynomial in its input length. The input length of \mathcal{A}_1 is $\text{poly}(\lambda')$, but its running time is $\text{poly}(\kappa)$, which may be super-polynomial in λ' .

that the distinguishing advantage between Hyb_0 and Hyb_7 is only $\text{negl}(\kappa)$. We inspect this in the following:

- The changes from Hyb_0 to Hyb_1 , from Hyb_2 to Hyb_3 , from Hyb_4 to Hyb_5 , and from Hyb_6 to Hyb_7 are statistical, where each statistical difference is bounded by $\text{poly}(\lambda')/2^{-\Lambda}$. We have $\text{poly}(\lambda')/2^{-\Lambda} \leq \text{poly}(\kappa)/2^{-\Lambda} = \text{negl}(\kappa)$ by our choice of Λ .
- The change from Hyb_1 to Hyb_2 is computational, which is dependent on the hardness of LWE. Since the size of \mathcal{A}_0 is bounded by $\text{poly}(\kappa)$, the distinguishing advantage between Hyb_1 and Hyb_2 should be bounded by $\text{negl}(\kappa)$ by the subexponential hardness of LWE and by our choice of Λ .
- The change from Hyb_3 to Hyb_4 is computational, which is dependent on the security of PRF. Since the size of \mathcal{A}_0 is bounded by $\text{poly}(\kappa)$, the distinguishing advantage between Hyb_3 and Hyb_4 should be bounded by $\text{negl}(\kappa)$ by the subexponential security of PRF.
- The changes from Hyb_5 to Hyb_6 is conceptual and thus they are equivalent.

We therefore conclude that the distinguishing advantage of \mathcal{A}_0 against Hyb_7 and Hyb_8 should be $\epsilon/Q(\lambda') - \text{negl}(\kappa)$. Then, from \mathcal{A}_0 , it is straightforward to extract a distinguisher \mathcal{A}'_0 against the distributions in Equation (5.20) with the same advantage and almost the same size. This concludes the proof of the theorem. \blacksquare

Theorem 6.47. *Let $\kappa = 2^{\lambda^c}$ for some constant c . Assuming non-uniform κ -evasive LWE (Assumption 6.14), subexponentially secure PRF against non-uniform adversary, and non-uniform sub-exponential LWE (Assumption 2.5), there exists a prFE scheme for function class $\mathcal{F}_{L(\lambda), \ell(\lambda), \text{dep}(\lambda)} = \{f : \{0, 1\}^L \rightarrow \{0, 1\}^\ell\}$ satisfying κ -prCT security as per Definition 6.3 with efficiency*

$$|\text{mpk}| = L \cdot \text{poly}(\text{dep}, \lambda), \quad |\text{sk}_f| = \ell \cdot \text{poly}(\text{dep}, \lambda), \quad |\text{ct}| = L \cdot \text{poly}(\text{dep}, \lambda).$$

where $\text{dep} = \text{poly}(\lambda)$ is the depth bound on the functions supported by the scheme.

CHAPTER 7

EVASIVE LWE: CLASSES, ATTACKS, AND REPERCUSSIONS

The evasive LWE assumption, introduced by [169, 163], has emerged as a powerful framework to argue the security of advanced cryptographic primitives. Over time, several variants of this assumption have been formulated, depending on the visibility of randomness and structural components. We recall the general formulation of the evasive LWE assumption. Let Samp be any PPT algorithm that, on input 1^λ , outputs a matrix \mathbf{P} and auxiliary information aux containing all randomness used by Samp . Let \mathbf{B} be sampled uniformly at random from $\mathbb{Z}_q^{n \times m}$, and \mathbf{s} be short vector sampled uniformly at random from \mathbb{Z}_q^n .

The assumption postulates that:

$$\text{if } (\mathbf{B}, \mathbf{P}, \mathbf{s}^\top \mathbf{B}, \mathbf{s}^\top \mathbf{P}, \text{aux}) \approx_c (\mathbf{B}, \mathbf{P}, \$, \$, \text{aux}),$$

$$\text{then } (\mathbf{B}, \mathbf{P}, \mathbf{s}^\top \mathbf{B}, \mathbf{B}^{-1}(\mathbf{P}), \text{aux}) \approx_c (\mathbf{B}, \mathbf{P}, \$, \mathbf{B}^{-1}(\mathbf{P}), \text{aux}),$$

where $\mathbf{B}^{-1}(\mathbf{P})$ denotes a short preimage of \mathbf{P} under \mathbf{B} .

This chapter systematically explores the different formulations of the evasive LWE assumption. We first define the public-coin version, where the sampler's randomness is fully exposed. We then discuss two flavors of the private-coin version: the *binding* variant, where the matrices \mathbf{B} and \mathbf{P} remain publicly known but the sampler's internal randomness is hidden, and the *hiding* variant, where \mathbf{B} , \mathbf{P} , or both are partially or fully hidden from the adversary's view. Later, we also discuss the *circular* variant of evasive LWE, where the distributions in the assumption additionally includes some circular encodings.

For each variant, we present the formal definitions, followed by known attacks that highlight their limitations. We also discuss the implications of these attacks on the prior versions of our constructions and the assumptions. We now proceed to formally define each variant, starting with the public-coin evasive LWE assumption.

7.1 PUBLIC-COIN EVASIVE LWE

The public-coin evasive LWE was introduced by [169]. In this version, the sampler Samp reveals all the random coins used during sampling.

Definition 7.1 (Public-Coin Evasive LWE). Let Samp be a PPT algorithm that, on input 1^λ , outputs

$$\mathbf{A}' \in \mathbb{Z}_q^{n \times m'}, \quad \mathbf{P} \in \mathbb{Z}_q^{n \times t}, \quad \text{aux} \in \{0, 1\}^*.$$

We define the following advantage functions:

$$\begin{aligned} \text{Adv}_{\mathcal{A}_0}^{\text{PRE}}(\lambda) &:= \Pr[\mathcal{A}_0(\mathbf{A}', \mathbf{s}^\top \mathbf{A}' + \mathbf{e}', \mathbf{s}^\top \mathbf{B} + \mathbf{e}, \mathbf{s}^\top \mathbf{P} + \mathbf{e}'', \text{aux}) = 1] \\ &\quad - \Pr[\mathcal{A}_0(\mathbf{A}', \mathbf{c}, \mathbf{c}_0, \mathbf{c}', \text{aux}) = 1], \end{aligned} \quad (7.1)$$

$$\begin{aligned} \text{Adv}_{\mathcal{A}_1}^{\text{POST}}(\lambda) &:= \Pr[\mathcal{A}_1(\mathbf{A}', \mathbf{s}^\top \mathbf{A}' + \mathbf{e}', \mathbf{s}^\top \mathbf{B} + \mathbf{e}, \mathbf{K}, \text{aux}) = 1] \\ &\quad - \Pr[\mathcal{A}_1(\mathbf{A}', \mathbf{c}, \mathbf{c}_0, \mathbf{K}, \text{aux}) = 1], \end{aligned} \quad (7.2)$$

where

$$\begin{aligned} (\mathbf{A}', \mathbf{P}, \text{aux}) &\leftarrow \text{Samp}(1^\lambda), \mathbf{B} \leftarrow \mathbb{Z}_q^{n \times m}, \quad \mathbf{s} \leftarrow \mathbb{Z}_q^n, \\ \mathbf{c} &\leftarrow \mathbb{Z}_q^m, \quad \mathbf{c}_0 \leftarrow \mathbb{Z}_q^t, \quad \mathbf{c}' \leftarrow \mathbb{Z}_q^{m'}, \\ \mathbf{e} &\leftarrow \mathcal{D}_{\mathbb{Z}_q, \chi}^m, \quad \mathbf{e}' \leftarrow \mathcal{D}_{\mathbb{Z}_q, \chi}^{m'}, \quad \mathbf{e}'' \leftarrow \mathcal{D}_{\mathbb{Z}_q, \chi}^t, \\ \mathbf{K} &\leftarrow \mathbf{B}^{-1}(\mathbf{P}, \tau) \text{ with standard deviation } \tau = O(\sqrt{m \log q}). \end{aligned}$$

We say that the *evasive* LWE assumption holds if for every PPT Samp , \mathcal{A}_0 , \mathcal{A}_1 , there

exists another PPT algorithm and a polynomial $Q(\cdot)$ such that

$$\text{Adv}_{\mathcal{A}_0}^{\text{PRE}}(\lambda) \geq \frac{\text{Adv}_{\mathcal{A}_1}^{\text{POST}}(\lambda)}{Q(\lambda)} - \text{negl}(\lambda).$$

Remark 42 (Noise Magnitudes [169]). For simplicity, we stated the assumption with all the LWE error terms $\mathbf{e}, \mathbf{e}', \mathbf{e}''$ having the same Gaussian parameter χ . It is straightforward to adapt the assumption and the scheme to a quantitatively weaker variant where the error terms in the post-condition have a larger Gaussian parameter than those in the pre-condition.

Attack on the public-coin evasive LWE and repercussions [19] presents an attack on a stronger version of public-coin evasive LWE, where the noise magnitude in the pre-condition is larger than the noise magnitude in the post-condition. We note that their attack *does not* affect the underlying assumption of the cpABE scheme of [169], which we use as a blackbox tool in Section 3.7 and Section 4.5.3.

7.2 PRIVATE-COIN BINDING EVASIVE LWE.

In the binding setting, Samp is private-coin, meaning its internal randomness remains hidden, but the matrices \mathbf{B} and \mathbf{P} are explicitly revealed.

Definition 7.2 (Private-Coin Binding Evasive LWE, [66, Definition 8]). Let $m, n, k, \ell > 0$ be integers and let q be a modulus. Let $\tau, \sigma, \sigma' > 0$. Let Samp be an algorithm that, on input 1^λ , outputs a matrix $\mathbf{P} \in \mathbb{Z}_q^{n \times k}$, a matrix $\mathbf{S} \in \mathbb{Z}_q^{\ell \times n}$, and auxiliary information aux. Let

$$\begin{aligned} \mathbf{B} &\leftarrow \mathbb{Z}_q^{n \times m}, \mathbf{K} \leftarrow \mathbf{B}^{-1}(\mathbf{P}, \tau), (\mathbf{S}, \mathbf{P}, \text{aux}) \leftarrow \text{Samp}(1^\lambda), \\ \mathbf{E} &\leftarrow \mathcal{D}_{\mathbb{Z}_q^{\ell \times m}, \sigma}, \quad \mathbf{E}' \leftarrow \mathcal{D}_{\mathbb{Z}_q^{\ell \times k}, \sigma'}, \mathbf{C} \leftarrow \mathbb{Z}_q^{\ell \times m}, \quad \mathbf{C}' \leftarrow \mathbb{Z}_q^{\ell \times k}. \end{aligned}$$

For PPT distinguishers \mathcal{A}_0 and \mathcal{A}_1 , we define the following functions:

$$\text{Adv}_{\mathcal{A}_0}^{\text{Pre}}(\lambda) := \Pr[\mathcal{A}_0(\mathbf{B}, \mathbf{P}, \mathbf{S}\mathbf{B} + \mathbf{E}, \mathbf{S}\mathbf{P} + \mathbf{E}', \text{aux}) = 1]$$

$$- \Pr[\mathcal{A}_0(\mathbf{B}, \mathbf{P}, \mathbf{C}, \mathbf{C}', \text{aux}) = 1], \quad (7.3)$$

$$\begin{aligned} \text{Adv}_{\mathcal{A}_1}^{\text{Post}}(\lambda) &:= \Pr[\mathcal{A}_1(\mathbf{B}, \mathbf{P}, \mathbf{SB} + \mathbf{E}, \mathbf{K}, \text{aux}) = 1] \\ &- \Pr[\mathcal{A}_1(\mathbf{B}, \mathbf{P}, \mathbf{C}, \mathbf{K}, \text{aux}) = 1]. \end{aligned} \quad (7.4)$$

We say that the binding evasive LWE assumption $\text{evLWE}(q, n, m, k, \ell, \text{Samp}, \tau, \sigma, \sigma')$ holds if there exists a polynomial Q such that $\deg_\lambda(Q) \leq c \cdot \deg_\lambda(|\text{Samp}|)$ (for some universal constant $c > 0$), and for every PPT distinguisher \mathcal{A}_1 there exists a PPT distinguisher \mathcal{A}_0 such that

$$\text{Adv}_{\mathcal{A}_0}^{\text{Pre}}(\lambda) \geq \frac{\text{Adv}_{\mathcal{A}_1}^{\text{Post}}(\lambda)}{Q(\lambda)} - \text{negl}(\lambda),$$

and $\text{time}(\mathcal{A}_0) \leq \text{time}(\mathcal{A}_1) \cdot Q(\lambda)$.

7.2.1 Attacks

Attack by [19, 121] The attacks proposed by [19] and [121], on the binding variant of evasive LWE, largely overlaps. They design the sampler of their counter-example by relying on the prior version of our scheme in Chapter 5¹ and using the tools from [120]. We give a high level overview of the attack next, majorly borrowed from Section 5.1.4. Construct a sampler Samp that does the following

- Choose a function f and an input vector $\mathbf{x} \in \{0, 1\}^L$ such that $f(\mathbf{x})$ is pseudorandom.
- Samples a GSW secret key \mathbf{s} and a PRF seed $\text{sd} \leftarrow \{0, 1\}^\lambda$. It then computes a GSW ciphertext, $\mathbf{X} = \mathbf{A}_{\text{fhe}} \mathbf{R} - (\mathbf{x}, \text{sd}) \otimes \mathbf{G}$, using public key $\mathbf{A}_{\text{fhe}} = (\bar{\mathbf{A}}_{\text{fhe}} \quad \bar{\mathbf{s}}^\top \bar{\mathbf{A}}_{\text{fhe}} + \mathbf{e}_{\text{fhe}}^\top)$ and randomness \mathbf{R} , – followed by a BGG^+ encoding of \mathbf{X} using randomness \mathbf{s} as $\mathbf{c}_{\text{att}}^\top := \mathbf{s}^\top (\mathbf{A}_{\text{att}} - \mathbf{X} \otimes \mathbf{G}) + \mathbf{e}_{\text{att}}^\top$. It additionally computes $\mathbf{c}_{\mathbf{B}}^\top := \mathbf{s}^\top \mathbf{B} + \mathbf{e}_{\mathbf{B}}^\top$.
- Samples a nonce $\mathbf{r} \leftarrow \{0, 1\}^\lambda$ and defines function $F[f, \mathbf{r}]$, with f and \mathbf{r} hardwired, as

$$F[f, \mathbf{r}](\mathbf{x}, \text{sd}) = f(\mathbf{x}) \lfloor q/2 \rfloor + \text{PRF}(\text{sd}, \mathbf{r}).$$

It then computes the FHE evaluation circuit VEval_F w.r.t. the function $F[f, \mathbf{r}]$ (this can be computed using the knowledge of $F[f, \mathbf{r}]$). Next, it computes the matrix $\mathbf{H}_{\mathbf{A}_{\text{att}}}^F$ for the circuit VEval_F using the public matrix \mathbf{A}_{att} . It sets $\mathbf{A}_F = \mathbf{A}_{\text{att}} \cdot \mathbf{H}_{\mathbf{A}_{\text{att}}}^F$

¹please refer to Section 5.1.4 for the prior version of our prFE scheme

and samples $\mathbf{K} \leftarrow \mathbf{B}^{-1}(\mathbf{A}_F)$.

- It outputs $(\mathbf{s}, \mathbf{P} = \mathbf{A}_F, \text{aux} = (\mathbf{X}, \mathbf{c}_{\text{att}}, \mathbf{c}_B, f, \mathbf{r}, \mathbf{A}_{\text{att}}))$.

Next, we discuss why the above sampler satisfies pre-condition but breaks post-condition and thus violating the private-coin binding evasive LWE assumption.

- *Pre-condition holds.* To argue that the pre-condition holds, it suffices to show the pseudorandomness of $(\mathbf{c}_B, \mathbf{c}_{\text{att}}, \mathbf{X}, \mathbf{c}_P)$, where $\mathbf{c}_P = \mathbf{sP} + \mathbf{e}_P$ with fresh Gaussian noise \mathbf{e}_P . The pseudorandomness is argued via the following sequence of hybrids.

Hyb₀. The adversary is given $(\mathbf{c}_B, \mathbf{c}_{\text{att}}, \mathbf{X}, \mathbf{c}_P)$, where all components are computed honestly by the sampler.

Hyb₁. Modify \mathbf{c}_P to be computed as $\mathbf{c}_P = \mathbf{c}_{\text{att}}\mathbf{H}_{\mathbf{A}_{\text{att}}, \mathbf{X}}^F + f(\mathbf{x}) + \mathbf{e}_P$. We use the observation that $\mathbf{c}_{\text{att}}\mathbf{H}_{\mathbf{A}_{\text{att}}, \mathbf{X}}^F + f(\mathbf{x}) + \mathbf{e}_P = \mathbf{sP} - \mathbf{e}_{\text{fhe}} + \mathbf{e}_{\text{att}}\mathbf{H}_{\mathbf{A}_{\text{att}}, \mathbf{X}}^F + \mathbf{e}_P$, where we set \mathbf{e}_P to be large enough to smudges the error terms $\mathbf{e}_{\text{fhe}} + \mathbf{e}_{\text{att}}\mathbf{H}_{\mathbf{A}_{\text{att}}, \mathbf{X}}^F$. So, we have $\text{Hyb}_0 \approx_s \text{Hyb}_1$.

Hyb₂. Replace $\mathbf{c}_B, \mathbf{c}_{\text{att}}$, and the GSW public key \mathbf{A}_{fhe} with uniformly random matrices. Since the GSW public key \mathbf{A}_{fhe} is an LWE sample with respect to the secret key \mathbf{s} , and $\mathbf{c}_B, \mathbf{c}_{\text{att}}$ are BGG^+ encodings, this replacement is computationally indistinguishable by the hardness of LWE.

Hyb₃. Replace the GSW ciphertext \mathbf{X} with a uniformly random matrix. This hybrid is statistically indistinguishable from Hyb_2 using leftover hash lemma.

Hyb₄. Finally, replace \mathbf{c}_P with a uniformly random string. This follows from the pseudorandomness of $f(\mathbf{x})$, and the fact that \mathbf{x} is not used elsewhere once \mathbf{X} has been replaced by a random matrix.

- *Post-condition fails.* The attack builds upon the ideas from [120] to correlate the least significant bit of $f(\mathbf{x})$ with the error term \mathbf{e}_{fhe} from the GSW homomorphic encryption. Specifically, they carefully design the circuit implementing f such that

$$f(\mathbf{x}) \equiv \mathbf{e}_{\text{fhe}}^\top \pmod{2}.$$

Once this correlation is established, the attack proceeds as follows. given $\mathbf{c}_B, \mathbf{c}_{\text{att}}, \mathbf{X}$, and \mathbf{K} ,

- Compute $\mathbf{v} = \mathbf{c}_B^\top \mathbf{K} - \mathbf{c}_{\text{att}}^\top \mathbf{H}_{\mathbf{A}_{\text{att}}, \mathbf{X}}^F \pmod{2} = f(\mathbf{x}) + \mathbf{e}_B^\top \mathbf{K} - (\mathbf{e}_{\text{fhe}}^\top \mathbf{R}_F + \mathbf{e}_{\text{att}}^\top \mathbf{H}_{\mathbf{A}_{\text{att}}, \mathbf{X}}^F)$. In this equation, the small error terms \mathbf{e}_{fhe} and \mathbf{e}_{att} cancel each other modulo 2.
- Find vectors \mathbf{e}_B and \mathbf{e}_{att} such that

$$\mathbf{e}_B^\top \mathbf{K} - \mathbf{e}_{\text{att}}^\top \mathbf{H}_{\mathbf{A}_{\text{att}}, \mathbf{X}}^F = \mathbf{v} \pmod{2}.$$

When the given terms are structured one can use sufficiently many equations to recover the error terms \mathbf{e}_B and \mathbf{e}_{att} . On the other hand, had the term $\mathbf{e}_B^T \mathbf{K}$ been truly random, such a system of equations would not admit any solution. This leads to a distinguishing strategy.

Attack 2 [85] The attacks proposed by [85] do not directly affect our construction, but they break the evasive LWE assumption underlying the proof of our initial construction when considered for *arbitrary* samplers. Specifically, they construct a counterexample via a malicious sampler that outputs a uniformly random matrix \mathbf{P} , a short-norm Gaussian matrix \mathbf{S} , and auxiliary information $\mathbf{aux} = \mathbf{SP} - 2\mathbf{T} \bmod q$, where \mathbf{T} consists of entries uniformly sampled from $[0, \lfloor q/2 \rfloor]$. They show the pre-condition $(\mathbf{B}, \mathbf{P}, \mathbf{SB} + \mathbf{E}, \mathbf{SP} + \mathbf{E}', \mathbf{aux}) \approx_c (\mathbf{B}, \mathbf{P}, \$, \$, \mathbf{aux})$, where $\$$ represents random, holds by relying on LWE with small secrets and noise flooding. They then give a distinguisher for the post-condition, which asserts that $(\mathbf{SB} + \mathbf{E}, \mathbf{D}, \mathbf{aux} = \mathbf{SP} - 2\mathbf{T}) \approx_c (\$, \mathbf{D}, \mathbf{aux})$. The distinguisher computes $\mathbf{W} = (\mathbf{CD} - \mathbf{aux}) \bmod q \bmod 2$ and checks whether the first row of \mathbf{W} lies in the row span of $\mathbf{D} \bmod 2$. In the real distribution, where $\mathbf{C} = \mathbf{SB} + \mathbf{E}$, the structure of \mathbf{ED} leaks through, and \mathbf{W} aligns with the span of $\mathbf{D} \bmod 2$. In the ideal distribution, where \mathbf{C} is uniform, \mathbf{W} appears uniform.

Effects on our Constructions and Assumptions. The attacks proposed by [19, 121] on the private-coin binding evasive LWE breaks our initial construction of the prFE scheme. The attacks proposed by [85] do not directly affect our constructions, but they break the evasive LWE assumption underlying the proof of our initial construction when considered for *arbitrary* samplers. In Section 5.1.4, we discuss in detail the fix to our initial construction and refining of our underlying assumption. Briefly, we follow the approach of [19] to restrict the sampler (please see Assumption 2.6 for the refined definition of evasive LWE) for which evasive LWE is conjectured true. To the best of our understanding, the original intuition of Evasive LWE still holds if malicious samplers such as the above are avoided.

7.3 PRIVATE-COIN HIDING EVASIVE LWE.

The hiding setting of the private-coin evasive LWE assumption is the strongest – here the sampler Samp is private-coin, and the matrices \mathbf{B}, \mathbf{P} are partially or completely hidden from the view of the adversary. In this variant, the assumption includes the requirement that $(\mathbf{P}, \text{aux}) \approx_c (\mathbf{P} + \mathbf{R}, \text{aux})$ where \mathbf{R} is a bounded-norm matrix. The indistinguishability between \mathbf{P} and $\mathbf{P} + \mathbf{R}$ guarantees that \mathbf{P} cannot be efficiently approximated, even with access to auxiliary information.

Definition 7.3 (Private-Coin Hiding Evasive LWE, [66, Definition 9]). Let $m, n, k, \ell > 0$ be integers and let q be a modulus. Let $\tau, \sigma, \sigma' > 0$. Let Samp be an algorithm that, on input 1^λ , outputs a matrix $\mathbf{P} \in \mathbb{Z}_q^{n \times k}$, a matrix $\mathbf{S} \in \mathbb{Z}_q^{\ell \times n}$, and auxiliary information aux . Let

$$\begin{aligned} \mathbf{B} &\leftarrow \mathbb{Z}_q^{n \times m}, \mathbf{K} \leftarrow \mathbf{B}^{-1}(\mathbf{P}, \tau), (\mathbf{S}, \mathbf{P}, \text{aux}) \leftarrow \text{Samp}(1^\lambda), \\ \mathbf{E} &\leftarrow \mathcal{D}_{\mathbb{Z}_q^{\ell \times m}, \sigma}, \mathbf{E}' \leftarrow \mathcal{D}_{\mathbb{Z}_q^{\ell \times k}, \sigma'}, \mathbf{C} \leftarrow \mathbb{Z}_q^{\ell \times m}, \mathbf{C}' \leftarrow \mathbb{Z}_q^{\ell \times k}, \mathbf{R} \leftarrow \mathcal{U}([\kappa])^{n \times k}. \end{aligned}$$

For PPT distinguishers \mathcal{A}' , \mathcal{A}'' , and \mathcal{A} , define the following functions:

$$\begin{aligned} \text{Adv}_{\mathcal{A}'}^{\text{Pre1}}(\lambda) &:= \Pr[\mathcal{A}'(\mathbf{SB} + \mathbf{E}, \mathbf{SP} + \mathbf{E}', \text{aux}) = 1] \\ &\quad - \Pr[\mathcal{A}'(\mathbf{C}, \mathbf{C}', \text{aux}) = 1], \end{aligned} \tag{7.5}$$

$$\begin{aligned} \text{Adv}_{\mathcal{A}''}^{\text{Pre2}}(\lambda) &:= \Pr[\mathcal{A}''(\mathbf{P}, \text{aux}) = 1] \\ &\quad - \Pr[\mathcal{A}''(\mathbf{P} + \mathbf{R}, \text{aux}) = 1], \end{aligned} \tag{7.6}$$

$$\begin{aligned} \text{Adv}_{\mathcal{A}}^{\text{Post}}(\lambda) &:= \Pr[\mathcal{A}(\mathbf{SB} + \mathbf{E}, \mathbf{K}, \text{aux}) = 1] \\ &\quad - \Pr[\mathcal{A}(\mathbf{C}, \mathbf{K}, \text{aux}) = 1]. \end{aligned} \tag{7.7}$$

We say that the hiding evasive LWE assumption $\text{evLWE}(q, m, n, k, \ell, \text{Samp}, \kappa, \tau, \sigma, \sigma')$ holds if there exists a polynomial Q such that $\deg_\lambda(Q) \leq c \cdot \deg_\lambda(|\text{Samp}|)$ (for some universal constant $c > 0$), and for every PPT distinguisher \mathcal{A} there exist PPT distinguishers

\mathcal{A}' , \mathcal{A}'' such that

$$\text{Adv}_{\mathcal{A}'}^{\text{Pre1}}(\lambda) + \text{Adv}_{\mathcal{A}''}^{\text{Pre2}}(\lambda) \geq \frac{\text{Adv}_{\mathcal{A}}^{\text{Post}}(\lambda)}{Q(\lambda)} - \text{negl}(\lambda),$$

and

$$\text{time}(\mathcal{A}') \leq \text{time}(\mathcal{A}) \cdot Q(\lambda), \quad \text{time}(\mathcal{A}'') \leq \text{time}(\mathcal{A}) \cdot Q(\lambda).$$

7.3.1 Attacks and Repercussions

Attack by [164] [164] noted that the private-coin evasive LWE assumption was already prone to heuristic obfuscation-based counterexamples. The attack works as follows: Let aux be an obfuscation of the following program $\Pi_{\mathbf{P}, \tau}$, which has the matrix \mathbf{P} and a corresponding trapdoor τ hard-wired into it. On input (\mathbf{C}, \mathbf{D}) , where $\mathbf{C} \in \mathbb{Z}_q^{2m \times m}$ and $\mathbf{D} \in \mathbb{Z}_q^{m \times 2m}$, the program attempts to find a short matrix \mathbf{S}_0 such that $\mathbf{CD} \approx \mathbf{S}_0 \mathbf{P} \bmod q$. If such a short \mathbf{S}_0 exists, the program outputs 1; otherwise, it outputs 0.

- *Pre-Condition Holds.* Under the LWE assumption, $(\mathbf{SB} + \mathbf{E}, \mathbf{SP} + \mathbf{E}')$ is pseudorandom. Thus, the pre-condition (indistinguishability from random) remains satisfied.
- *Post-Condition Broken.* The adversary receives $(\mathbf{SB} + \mathbf{E}, \mathbf{B}^{-1}(\mathbf{P}), \text{aux})$. It computes $(\mathbf{SB} + \mathbf{E})\mathbf{B}^{-1}(\mathbf{P}) \approx \mathbf{SP}$, and feeds $(\mathbf{SB} + \mathbf{E}, \mathbf{B}^{-1}(\mathbf{P}))$ into the obfuscated program. Using the embedded trapdoor τ , the program verifies whether \mathbf{SP} is close to some $\mathbf{S}_0 \mathbf{P}$ and outputs 1 if so. For random matrices \mathbf{C} , the product \mathbf{CD} behaves randomly, and the program outputs 0. Thus, the adversary can distinguish valid from random instances, violating the post-condition.

Attacks by [66] The obfuscation-based counterexample of [66] shows that the private-coin evasive LWE assumption can be broken even without relying on full-fledged obfuscation, using only null-iO and the hardness of LWE. In this attack, the sampler Samp generates a tall matrix $\mathbf{S} \in \mathbb{Z}_q^{m_P \times n}$ and a wide matrix $\mathbf{P} \in \mathbb{Z}_q^{n \times m_P}$, both uniformly random, and outputs auxiliary information aux containing a null-iO obfuscation of a circuit $C_{\mathbf{W}}$. Here, $\mathbf{W} = \mathbf{SP} + \mathbf{E}''$ with \mathbf{E}'' sampled as short noise. The obfuscated circuit $C_{\mathbf{W}}$ is designed to check, on input matrices $(\mathbf{M}_1, \mathbf{M}_2)$, whether the product $\mathbf{M}_1 \mathbf{M}_2$ closely approximates the matrix \mathbf{W} . It outputs 1 if they are close and 0 otherwise.

- *Post-Condition Broken.* In the post-condition, the adversary receives $(\mathbf{SB} +$

$\mathbf{E}, \mathbf{B}^{-1}(\mathbf{P})$). By computing $(\mathbf{S}\mathbf{B} + \mathbf{E})\mathbf{B}^{-1}(\mathbf{P})$, the adversary effectively recovers $\mathbf{S}\mathbf{P}$ up to small noise. Feeding this into the obfuscated circuit results in the circuit outputting 1 with overwhelming probability. This allows the adversary to distinguish between real and random samples, thereby breaking the post-condition.

- *Pre-Condition Holds.* In the pre-condition, the auxiliary information \mathbf{aux} becomes useless for distinguishing. By the hardness of LWE, the matrix $\mathbf{S}\mathbf{P} + \mathbf{E}''$ is computationally indistinguishable from a uniformly random matrix \mathbf{R} . Consequently, the circuit $C_{\mathbf{R}}$ behaves essentially as the constant-zero function. By the security of null-iO, the obfuscation of $C_{\mathbf{R}}$ can be safely replaced with an obfuscation of a padded zero circuit, which leaks no meaningful information to the adversary.

[66] also present several other simple algebraic attacks that target the hiding variant of private-coin evasive LWE.

Attack by [124] [124] first shows that the counterexample presented by [164] is invalid. Specifically, it demonstrates that the pre-condition distributions in the VWW counterexample are actually distinguishable, contrary to their claim. Since evasive LWE only concerns the case where pre-condition distributions are indistinguishable, this renders the original VWW attack ineffective. Nevertheless, [124] constructs a new valid counterexample that breaks the underlying assumption of [164], using LWE and the existence of instance-hiding witness encryption (ihWE). Briefly, instance-hiding WE is a strengthening of witness encryption where encryptions with respect to two different unsatisfiable instances are computationally indistinguishable.

The counterexample is as follows : The sampler Samp outputs a random tall matrix \mathbf{S} , a wide matrix \mathbf{P} , and auxiliary information \mathbf{aux} consisting of an ihWE encryption of $\mathbf{S}\mathbf{P} + \mathbf{E}''$, for short noise \mathbf{E}'' . In the pre-condition, by LWE hardness and the instance-hiding property of WE, the ciphertext \mathbf{aux} hides any information about \mathbf{S} or \mathbf{P} , and thus indistinguishability is preserved. However, in the post-condition, when the adversary receives $(\mathbf{S}\mathbf{B} + \mathbf{E}, \mathbf{B}^{-1}(\mathbf{P}))$, it can use these to approximate $\mathbf{S}\mathbf{P}$. Feeding this approximation into the ihWE decryption algorithm successfully recovers the hidden random string

embedded inside the ciphertext, allowing the adversary to distinguish real samples from random ones and thus break the post-condition.

Effects on our Constructions and Assumptions. We note that none of the initial versions of our constructions in Chapters 4, 5, or 6 relied on the hiding variant of the private-coin evasive LWE assumption. In all our schemes, the matrix \mathbf{B} is publicly available, and the matrix \mathbf{P} is publicly computable using the auxiliary information. Thus, initial versions of our constructions and the underlying assumptions *did not* get affected by the attacks discussed in this section.

7.4 CIRCULAR SMALL-SECRET EVASIVE LWE

Here we define the circular small-secret evasive LWE (evcsLWE) variant introduced by [122].

Assumption 7.1 (evcsLWE). Let $\mathcal{S}(1^\lambda; \text{aux})$ be an algorithm that, given randomness aux , outputs

$$\mathbf{A}_{\text{circ}} \leftarrow \mathbb{Z}_q^{(n+1) \times (m(n+1)^2 \lceil \log_2 q \rceil^2 + 1)m}, \bar{\mathbf{A}}' \in \mathbb{Z}_q^{n \times m'}, \mathbf{P} \in \mathbb{Z}_q^{n \times J}, \sigma, \sigma', \sigma_{-1}, \sigma_{\text{post}}, \sigma_{\text{pre}}$$

where $m \geq m_0(n, q)$ and $\sigma_{-1} \geq \sigma_0(n, m)$ and $\sigma_{\text{post}} \geq \sigma_{\text{pre}}$. Suppose

$$\begin{aligned} \bar{\mathbf{A}}_{\text{fhe}} &\leftarrow \mathbb{Z}_q^{n \times m}, \quad (\mathbf{B}, \tau) \leftarrow \text{TrapGen}(1^n, 1^m, q), \quad \mathbf{K} \leftarrow \mathbf{B}^{-1}(\mathbf{P}), \\ \mathbf{e}_{\text{fhe}} &\leftarrow \mathcal{D}_{\mathbb{Z}, \sigma}^m, \quad \mathbf{e}_{\text{circ}} \leftarrow \mathcal{D}_{\mathbb{Z}, \sigma'}^{(m(n+1)^2 \lceil \log_2 q \rceil^2 + 1)m}, \quad \mathbf{e}' \leftarrow \mathcal{D}_{\mathbb{Z}, \sigma'}^{m'}, \quad \mathbf{e}_{\mathbf{B}} \in \mathbb{Z}^m, \quad \mathbf{e}_{\mathbf{P}} \in \mathbb{Z}^J, \\ \delta_{\text{fhe}} &\leftarrow \mathbb{Z}_q^m, \quad \delta_{\text{circ}} \leftarrow \mathbb{Z}_q^{(m(n+1)^2 \lceil \log_2 q \rceil^2 + 1)m}, \quad \delta' \leftarrow \mathbb{Z}_q^{m'}, \quad \delta_{\mathbf{B}} \leftarrow \mathbb{Z}_q^m, \quad \delta_{\mathbf{P}} \leftarrow \mathbb{Z}_q^J, \\ \mathbf{r} &\leftarrow \mathcal{D}_{\mathbb{Z}, \sigma}^n, \quad \mathbf{s} \leftarrow (\mathbf{r}^\top, -1)^\top, \quad \mathbf{R} \leftarrow \{0, 1\}^{m \times (n+1) \lceil \log_2 q \rceil m}, \quad \Delta \leftarrow \mathbb{Z}_q^{(n+1) \times (n+1) \lceil \log_2 q \rceil m}, \\ \mathbf{S} &= \begin{pmatrix} \bar{\mathbf{A}}_{\text{fhe}} \\ \mathbf{r}^\top \bar{\mathbf{A}}_{\text{fhe}} + \mathbf{e}_{\text{fhe}}^\top \end{pmatrix} \mathbf{R} - (\text{bits}(\mathbf{s})) \otimes \mathbf{G} \end{aligned}$$

In the precondition, the entries of $\mathbf{e}_{\mathbf{B}}, \mathbf{e}_{\mathbf{P}}$ are independent and follow $\mathcal{D}_{\mathbb{Z}, \sigma_{\text{pre}}}$, and

$\text{evcsLWE}_{\text{pre}}^S$ states that

$$\left\{ \begin{pmatrix} 1^\lambda, \text{aux}, \bar{\mathbf{A}}_{\text{fhe}}, \mathbf{B}, \mathbf{r}^\top \bar{\mathbf{A}}_{\text{fhe}} + \mathbf{e}_{\text{fhe}}^\top, \mathbf{S}, \\ \mathbf{s}^\top (\mathbf{A}_{\text{circ}} - (1, \text{bits}(\mathbf{S})) \otimes \mathbf{G}) + \mathbf{e}_{\text{circ}}^\top \\ \mathbf{r}^\top \bar{\mathbf{A}}' + (\mathbf{e}')^\top, \mathbf{r}^\top \mathbf{B} + \mathbf{e}_{\mathbf{B}}^\top, \mathbf{r}^\top \mathbf{P} + \mathbf{e}_{\mathbf{P}}^\top \end{pmatrix} \right\}_{\lambda \in \mathbb{N}} \approx \left\{ \begin{pmatrix} 1^\lambda, \text{aux}, \bar{\mathbf{A}}_{\text{fhe}}, \mathbf{B}, \delta_{\text{fhe}}^\top, \Delta, \\ \delta_{\text{circ}}^\top, \\ (\delta')^\top, \delta_{\mathbf{B}}^\top, \delta_{\mathbf{P}}^\top \end{pmatrix} \right\}_{\lambda \in \mathbb{N}}.$$

In the postcondition, the entries of $\mathbf{e}_{\mathbf{B}}$ are independent and follow $\mathcal{D}_{\mathbb{Z}, \sigma_{\text{post}}}$, and

$\text{evcsLWE}_{\text{post}}^S$ states that

$$\left\{ \begin{pmatrix} 1^\lambda, \text{aux}, \bar{\mathbf{A}}_{\text{fhe}}, \mathbf{B}, \mathbf{r}^\top \bar{\mathbf{A}}_{\text{fhe}} + \mathbf{e}_{\text{fhe}}^\top, \mathbf{S}, \\ \mathbf{s}^\top (\mathbf{A}_{\text{circ}} - (1, \text{bits}(\mathbf{S})) \otimes \mathbf{G}) + \mathbf{e}_{\text{circ}}^\top \\ \mathbf{r}^\top \bar{\mathbf{A}}' + (\mathbf{e}')^\top, \mathbf{r}^\top \mathbf{B} + \mathbf{e}_{\mathbf{B}}^\top, \mathbf{K} \end{pmatrix} \right\}_{\lambda \in \mathbb{N}} \approx \left\{ \begin{pmatrix} 1^\lambda, \text{aux}, \bar{\mathbf{A}}_{\text{fhe}}, \mathbf{B}, \delta_{\text{fhe}}^\top, \Delta, \\ \delta_{\text{circ}}^\top, \\ (\delta')^\top, \delta_{\mathbf{B}}^\top, \mathbf{K} \end{pmatrix} \right\}_{\lambda \in \mathbb{N}}.$$

The *evasive circular small-secret LWE assumption* states that $\text{evcsLWE}_{\text{pre}}^S$ implies $\text{evcsLWE}_{\text{post}}^S$ for all efficient samplers S .

Remark 43. In our assumption (Assumption 5.18) used in constructing prFE (Section 5.3.4), the FHE encoding \mathbf{S} additionally encodes the attribute \mathbf{x} , whereas in above it only encode the FHE secret key \mathbf{s} . This is created by the difference between ABE and FE, since the attribute \mathbf{x} can be public in the former and must be hidden in the latter.

Attacks and repercussions. [19] presents an attack on the circular small-secret evasive LWE variant introduced by [122], which is very similar to the attack on the private-coin binding evasive assumption. To safeguard our construction of prFE (Section 5.3.4) based on a variant of this assumption, we implement the modulus reduction fix as discussed in Section 5.1.4 and also further refine the assumption itself. In particular, we conjecture that our sampler, used in the security proof of our prFE , for natural classes of functions, lies within the secure class of samplers \mathcal{SC} for which the evasive circular small-secret LWE assumption holds.

7.5 CONTRIVED FUNCTIONALITY ATTACKS.

While we prove the existence of our prFE and prIO schemes assuming LWE and (a variant) of private-coin binding of Evasive LWE, the work of [65] and [19] showed that there exist contrived, somehow “self-referential”, classes of pseudorandom functionalities for which pseudorandom obfuscation and pseudorandom FE cannot exist.

Impossibility of Pseudorandom Obfuscation. Here we describe the impossibility result of [65] for pseudorandom obfuscation (PRO). Recall that in pseudorandom obfuscation (PRO), for a PRF family $\mathcal{F} = \{f_K\}_{K \in \mathcal{K}}$ and an auxiliary input function $\mathbf{aux} : \mathcal{K} \rightarrow \{0, 1\}^*$, the definition requires:

- If the truth-tables $\{f_K(x)\}_{x \in \mathcal{X}}$, together with \mathbf{aux}_K , are pseudorandom, that is,

$$\{f_K(x)\}_{x \in \mathcal{X}}, \mathbf{aux}_K \approx_c \{u_x\}_{x \in \mathcal{X}}, \mathbf{aux}_K,$$

where each $u_x \leftarrow \mathcal{Y}$ is uniform,

- then the obfuscations $\text{PRO.Obf}(f_K)$ and $\text{PRO.Obf}(f_{K'})$ must be indistinguishable in the presence of \mathbf{aux}_K :

$$\text{PRO.Obf}(f_K), \mathbf{aux}_K \approx_c \text{PRO.Obf}(f_{K'}), \mathbf{aux}_K,$$

where K, K' are independent uniformly random keys.

The counterexample idea is to let \mathbf{aux}_K contain a witness encryption (WE) ciphertext for the NP statement φ asserting that “there exists a small program P computing f_K .” The pre-condition of PRO holds via the pseudorandomness of \mathcal{F} and the semantic security of the witness encryption scheme: when the function outputs are pseudorandom (or replaced with random values), the instance φ corresponds to a “no-instance” and the WE ciphertexts become computationally indistinguishable from random strings.

However, the post-condition of PRO becomes vulnerable: the obfuscated program $\text{PRO.Obf}(f_K)$ provides a short program P that serves as a valid witness for φ , allowing the adversary to decrypt the auxiliary input \mathbf{aux}_K . In contrast, without access to $\text{PRO.Obf}(f_K)$, the authors use a counting argument to show that the probability of

correctly decrypting aux_K is negligible. This yields a distinguishing attack that violates the PRO security guarantee. They further extend their counter example to the case where there is no auxilliary input.

Impossibility of Pseudorandom Functional Encryption. [19] adapts the counterexamples from [65] against the existence of PRO for all function families to show that there exist pseudorandom function families for which multi-challenge PRFE is not possible.

Repercussions and Countermeasures. As discussed in [19], this impossibility result can be seen as analogous to the known impossibilities for the random oracle model (ROM) [68] and virtual black-box (VBB) obfuscation [34]. In more detail, despite these impossibility results, the practical usefulness of ROM in cryptographic protocols, and of VBB obfuscation for restricted classes of functionalities [167, 69], is widely accepted. Similarly, the pseudorandom functionalities used in our applications – such as computing blind garbled circuits or functional encryption ciphertexts—are quite natural and do not fall prey to the contrived counterexamples constructed for PRO impossibility.

Furthermore, we note that the above counterexamples do not apply to the single-challenge pseudorandom functional encryption setting, where we set $Q_{\text{msg}} = 1$ in Definition 5.13. Thus, even for general functionalities, single-challenge prFE remains feasible. Consequently, our applications stated in Theorem 5.8, Theorem 5.9, and Corollary 5.9.1 remain safe against these attacks.

7.6 OUR PERSPECTIVE.

We view these attacks as an important step forward in our understanding of evasive LWE. Note that evasive LWE should be seen as a family of assumptions parametrized by the description of the sampler and choice of error distributions, which, if invoked in full generality, is now known to be false, even in the public-coin setting. However,

as discussed in [19], the original intuition by Wee [169] and Tsabary [163] about the security of evasive LWE can be recovered by taking precautions to identify and respect a “safe zone” for evasive LWE – thus, by refining/restraining the formulation of evasive LWE, even the original construction (presented in the first online posting of this work) is secure. In addition, we modify our original construction of prFE to implement a modulus reduction step, which negates the effect of choosing contrived circuits in the construction – for this modified construction, we do not even know of any attack using malicious samplers. We also remark that in the real world, the circuit representation of functions is chosen by the key generator who is an honest party (it holds the master secret key), and we view the counter-examples emerging from such circuit choices as indicating the limits of the assumption rather than the security of the existing constructions.

While counter-examples can (and should) create concern about indiscriminately using the assumption, one school of thought might be to completely discard the assumption by labeling it false and meaningless. Another school of thought might be that counter-examples are also progress, and that they shed light on if/how the assumption should be refined to regain security, and constructions from these refined assumptions are still meaningful (especially in the absence of alternatives).

Our perspective is that disciplined new conjectures are important to make meaningful progress on problems that have resisted solutions from standard assumptions. In this context, we believe that the evasive LWE assumption has played a crucial role in enabling constructions that could not be built from plain LWE despite significant effort by the community over several years. By examining carefully the nature of counter-examples and seeing whether there are meaningful lessons to learn, we can hope to arrive at stable versions that allow to expand the boundaries of cryptography. These constructions and their proofs could yield insights that would eventually enable candidates from standard assumptions (as happened in the world of pairings, for instance) [165, 103, 18, 105, 132].

CHAPTER 8

CONCLUSIONS

This thesis explores how evasive **LWE** can unlock expressive and fine-grained encryption primitives that remain out of reach under standard lattice assumptions. Motivated by the need for controlled access and secure computation in modern distributed systems, we explore a suite of fine-grained encryption primitives that are efficient, quantum-resilient, and provably secure under explicit hardness assumptions.

We begin with broadcast, trace, and revoke (**TR**) systems, aiming to distribute encrypted content while preventing access by revoked users and supporting tracing of pirate decoders. Our constructions achieve optimal parameters and embed user identities directly into secret keys, thereby eliminating the need for index-to-identity mappings and enhancing user anonymity. In the public trace setting, we rely only on polynomial-hardness assumptions, namely compact functional encryption and key-policy attribute-based encryption (both instantiable under well-understood assumptions). In the secret trace setting, we achieve the first optimal-size **TR** system with embedded identities from assumptions outside Obfustopia—specifically, **LWE**, lockable obfuscation, and a ciphertext-policy **ABE** built from evasive and tensor **LWE**. Both variants extend to support super-polynomial revocation lists under subexponential **LWE**.

Next, we construct the first attribute-based encryption (**ABE**) scheme for Turing machines with unbounded collusion resistance from lattice-based assumptions. Our construction builds on **LWE**, evasive **LWE**, and a new circular tensor **LWE** assumption, allowing encryption over unbounded-length inputs and decryption time that scales with the input. As a stepping stone, we also obtain the first ciphertext-policy **ABE** for circuits of unbounded depth from the same assumptions.

We then present the first compact functional encryption (FE) for pseudorandom functionalities from LWE and private-coin evasive LWE, achieving compactness without relying on pairings or indistinguishability obfuscation. This construction leads to optimal key-policy and ciphertext-policy ABE for unbounded-depth circuits under the same assumptions. We further compile this construction into multi-input FE and indistinguishability obfuscation (iO) for pseudorandom functionalities.

Finally, we address recent counterexamples to evasive LWE, which show that its most general forms can fail under malicious samplers or contrived functionalities. We argue that the assumption remains meaningful when used in well-defined “safe zones”.

Future Directions. This thesis opens up several promising directions for future research.

Below, we highlight some key directions:

- *Quantum-Safe Public Traceable Broadcast, Tace and Revoke with Optimality.* Our construction of optimal broadcast, trace and revoke (TR) in the public trace setting currently relies on quantum-insecure assumptions – specifically, compact functional encryption schemes that are not known to be quantum-safe. An important open problem is to construct TR systems in the public trace setting under quantum-safe assumptions, ideally from lattice-based primitives alone.
- *Applications of New Tools.* We introduced several new cryptographic tools such as functional encryption for pseudorandom functionalities (prFE), multi-input functional encryption for pseudorandom functionalities (mi-prFE) and indistinguishability obfuscation for pseudorandom functionalities (PRIO). Exploring new applications of these tools in broader cryptographic contexts could further demonstrate their utility.
- *Weaker Notions of prCT Security.* Due to known impossibility results, our compact functional encryption for pseudorandom functionalities does not achieve prCT security in the multi-challenge setting. Investigating relaxed or alternative security notions that are still meaningful for the applications we study or the new applications would be valuable.
- *Constructions from Falsifiable Assumptions.* Several of our constructions rely on variants of evasive LWE, which are currently non-falsifiable. An important goal is to base these primitives on well-studied, falsifiable assumptions. This may involve developing new reductions from existing lattice assumptions or refining our current assumptions to admit falsifiability.

BIBLIOGRAPHY

- [1] Abdalla, M., Benhamouda, F., Gay, R.: From single-input to multi-client inner-product functional encryption. In: Galbraith, S.D., Moriai, S. (eds.) ASIACRYPT 2019, Part III. LNCS, vol. 11923, pp. 552–582. Springer, Cham, Switzerland, Kobe, Japan (Dec 8–12, 2019). https://doi.org/10.1007/978-3-030-34618-8_19
- [2] Abdalla, M., Benhamouda, F., Kohlweiss, M., Waldner, H.: Decentralizing inner-product functional encryption. In: Lin, D., Sako, K. (eds.) PKC 2019, Part II. LNCS, vol. 11443, pp. 128–157. Springer, Cham, Switzerland, Beijing, China (Apr 14–17, 2019). https://doi.org/10.1007/978-3-030-17259-6_5
- [3] Abdalla, M., Catalano, D., Fiore, D., Gay, R., Ursu, B.: Multi-input functional encryption for inner products: Function-hiding realizations and constructions without pairings. In: Shacham, H., Boldyreva, A. (eds.) CRYPTO 2018, Part I. LNCS, vol. 10991, pp. 597–627. Springer, Cham, Switzerland, Santa Barbara, CA, USA (Aug 19–23, 2018). https://doi.org/10.1007/978-3-319-96884-1_20
- [4] Abdalla, M., Gay, R., Raykova, M., Wee, H.: Multi-input inner-product functional encryption from pairings. In: Coron, J.S., Nielsen, J.B. (eds.) EUROCRYPT 2017, Part I. LNCS, vol. 10210, pp. 601–626. Springer, Cham, Switzerland, Paris, France (Apr 30 – May 4, 2017). https://doi.org/10.1007/978-3-319-56620-7_21
- [5] Agrawal, S.: Indistinguishability Obfuscation Without Multilinear Maps: New Techniques for Bootstrapping and Instantiation. In: EUROCRYPT. LNCS, vol. 11476, pp. 191–225. Springer (2019), https://doi.org/10.1007/978-3-030-17653-2_7
- [6] Agrawal, S., Boneh, D., Boyen, X.: Efficient lattice (H)IBE in the standard model. In: Gilbert, H. (ed.) EUROCRYPT 2010. LNCS, vol. 6110, pp. 553–572. Springer Berlin Heidelberg, Germany, French Riviera (May 30 – Jun 3, 2010). https://doi.org/10.1007/978-3-642-13190-5_28
- [7] Agrawal, S., Boneh, D., Boyen, X.: Lattice basis delegation in fixed dimension and shorter-ciphertext hierarchical IBE. In: Rabin, T. (ed.) CRYPTO 2010. LNCS, vol. 6223, pp. 98–115. Springer Berlin Heidelberg, Germany, Santa Barbara, CA, USA (Aug 15–19, 2010). https://doi.org/10.1007/978-3-642-14623-7_6
- [8] Agrawal, S., Goyal, R., Tomida, J.: Multi-input quadratic functional encryption from pairings. In: Malkin, T., Peikert, C. (eds.) CRYPTO 2021, Part IV. LNCS, vol. 12828, pp. 208–238. Springer, Cham, Switzerland, Virtual Event (Aug 16–20,

- 2021). https://doi.org/10.1007/978-3-030-84259-8_8
- [9] Agrawal, S., Goyal, R., Tomida, J.: Multi-party functional encryption. In: TCC. LNCS, vol. 13043, pp. 224–255. Springer (2021), https://doi.org/10.1007/978-3-030-90453-1_8
 - [10] Agrawal, S., Goyal, R., Tomida, J.: Multi-input quadratic functional encryption: Stronger security, broader functionality. In: TCC. LNCS, vol. 13747, pp. 711–740. Springer (2022), https://doi.org/10.1007/978-3-031-22318-1_25
 - [11] Agrawal, S., Kumari, S., Yadav, A., Yamada, S.: Broadcast, trace and revoke with optimal parameters from polynomial hardness. In: Annual International Conference on the Theory and Applications of Cryptographic Techniques. pp. 605–636. Springer (2023). https://doi.org/10.1007/978-3-031-30620-4_20
 - [12] Agrawal, S., Kumari, S., Yamada, S.: Compact Pseudorandom Functional Encryption from Evasive LWE. IACR Cryptology ePrint Archive (2024)
 - [13] Agrawal, S., Kumari, S., Yamada, S.: Attribute Based Encryption for Turing Machines from Lattices. In: Crypto (2024). https://doi.org/10.1007/978-3-031-68382-4_11
 - [14] Agrawal, S., Kumari, S., Yamada, S.: Pseudorandom Multi-Input Functional Encryption and Applications. IACR Cryptology ePrint Archive (2024)
 - [15] Agrawal, S., Maitra, M.: FE and iO for Turing machines from minimal assumptions. In: Beimel, A., Dziembowski, S. (eds.) TCC 2018, Part II. LNCS, vol. 11240, pp. 473–512. Springer, Cham, Switzerland, Panaji, India (Nov 11–14, 2018). https://doi.org/10.1007/978-3-030-03810-6_18
 - [16] Agrawal, S., Maitra, M., Vempati, N.S., Yamada, S.: Functional encryption for Turing machines with dynamic bounded collusion from LWE. In: Malkin, T., Peikert, C. (eds.) CRYPTO 2021, Part IV. LNCS, vol. 12828, pp. 239–269. Springer, Cham, Switzerland, Virtual Event (Aug 16–20, 2021). https://doi.org/10.1007/978-3-030-84259-8_9
 - [17] Agrawal, S., Maitra, M., Yamada, S.: Attribute based encryption (and more) for nondeterministic finite automata from LWE. In: Boldyreva, A., Micciancio, D. (eds.) CRYPTO 2019, Part II. LNCS, vol. 11693, pp. 765–797. Springer, Cham, Switzerland, Santa Barbara, CA, USA (Aug 18–22, 2019). https://doi.org/10.1007/978-3-030-26951-7_26
 - [18] Agrawal, S., Maitra, M., Yamada, S.: Attribute based encryption for deterministic finite automata from DLIN. In: Hofheinz, D., Rosen, A. (eds.) TCC 2019, Part II. LNCS, vol. 11892, pp. 91–117. Springer, Cham,

- Switzerland, Nuremberg, Germany (Dec 1–5, 2019). https://doi.org/10.1007/978-3-030-36033-7_4
- [19] Agrawal, S., Modi, A., Yadav, A., Yamada, S.: Evasive LWE: Attacks, Variants & Obfuscopia. IACR Cryptology ePrint Archive (2025)
 - [20] Agrawal, S., Pellet-Mary, A.: Indistinguishability obfuscation without maps: Attacks and fixes for noisy linear FE. In: EUROCRYPT. LNCS, vol. 12105, pp. 110–140. Springer (2020), https://doi.org/10.1007/978-3-030-45721-1_5
 - [21] Agrawal, S., Pellet-Mary, A.: Indistinguishability obfuscation without maps: Attacks and fixes for noisy linear FE. In: Canteaut, A., Ishai, Y. (eds.) EUROCRYPT 2020, Part I. LNCS, vol. 12105, pp. 110–140. Springer, Cham, Switzerland, Zagreb, Croatia (May 10–14, 2020). https://doi.org/10.1007/978-3-030-45721-1_5
 - [22] Agrawal, S., Rossi, M., Yadav, A., Yamada, S.: Constant input attribute based (and predicate) encryption from evasive and tensor LWE. In: Handschuh, H., Lysyanskaya, A. (eds.) CRYPTO 2023, Part IV. LNCS, vol. 14084, pp. 532–564. Springer, Cham, Switzerland, Santa Barbara, CA, USA (Aug 20–24, 2023). https://doi.org/10.1007/978-3-031-38551-3_17
 - [23] Agrawal, S., Tomida, J., Yadav, A.: Attribute-based multi-input FE (and more) for attribute-weighted sums. In: Handschuh, H., Lysyanskaya, A. (eds.) CRYPTO 2023, Part IV. LNCS, vol. 14084, pp. 464–497. Springer, Cham, Switzerland, Santa Barbara, CA, USA (Aug 20–24, 2023). https://doi.org/10.1007/978-3-031-38551-3_15
 - [24] Agrawal, S., Wichs, D., Yamada, S.: Optimal broadcast encryption from lwe and pairings in the standard model. In: TCC (2020). https://doi.org/10.1007/978-3-030-64375-1_6
 - [25] Agrawal, S., Yadav, A., Yamada, S.: Multi-input attribute based encryption and predicate encryption. In: Dodis, Y., Shrimpton, T. (eds.) CRYPTO 2022, Part I. LNCS, vol. 13507, pp. 590–621. Springer, Cham, Switzerland, Santa Barbara, CA, USA (Aug 15–18, 2022). https://doi.org/10.1007/978-3-031-15802-5_21
 - [26] Agrawal, S., Yamada, S.: CP-ABE for circuits (and more) in the symmetric key setting. In: Pass, R., Pietrzak, K. (eds.) TCC 2020, Part I. LNCS, vol. 12550, pp. 117–148. Springer, Cham, Switzerland, Durham, NC, USA (Nov 16–19, 2020). https://doi.org/10.1007/978-3-030-64375-1_5
 - [27] Agrawal, S., Yamada, S.: Optimal broadcast encryption from pairings and LWE. In: Canteaut, A., Ishai, Y. (eds.) EUROCRYPT 2020, Part I. LNCS, vol. 12105, pp. 13–43. Springer, Cham, Switzerland, Zagreb, Croatia (May 10–14, 2020).

https://doi.org/10.1007/978-3-030-45721-1_2

- [28] Ananth, P., Brakerski, Z., Segev, G., Vaikuntanathan, V.: From selective to adaptive security in functional encryption. In: CRYPTO (2015). https://doi.org/10.1007/978-3-662-48000-7_32
- [29] Ananth, P., Jain, A.: Indistinguishability obfuscation from compact functional encryption. In: Gennaro, R., Robshaw, M.J.B. (eds.) CRYPTO 2015, Part I. LNCS, vol. 9215, pp. 308–326. Springer Berlin Heidelberg, Germany, Santa Barbara, CA, USA (Aug 16–20, 2015). https://doi.org/10.1007/978-3-662-47989-6_15
- [30] Ananth, P., Jain, A., Sahai, A.: Achieving compactness generically: Indistinguishability obfuscation from non-compact functional encryption. IACR Cryptology ePrint Archive, 2015:730 (2015)
- [31] Ananth, P., Vaikuntanathan, V.: Optimal bounded-collusion secure functional encryption. In: Theory of Cryptography Conference. pp. 174–198. Springer (2019). https://doi.org/10.1007/978-3-030-36030-6_8
- [32] Ananth, P.V., Sahai, A.: Functional encryption for Turing machines. In: Kushilevitz, E., Malkin, T. (eds.) TCC 2016-A, Part I. LNCS, vol. 9562, pp. 125–153. Springer Berlin Heidelberg, Germany, Tel Aviv, Israel (Jan 10–13, 2016). https://doi.org/10.1007/978-3-662-49096-9_6
- [33] Attrapadung, N., Libert, B., de Panafieu, E.: Expressive key-policy attribute-based encryption with constant-size ciphertexts. In: Catalano, D., Fazio, N., Gennaro, R., Nicolosi, A. (eds.) PKC 2011. LNCS, vol. 6571, pp. 90–108. Springer Berlin Heidelberg, Germany, Taormina, Italy (Mar 6–9, 2011). https://doi.org/10.1007/978-3-642-19379-8_6
- [34] Barak, B., Goldreich, O., Impagliazzo, R., Rudich, S., Sahai, A., Vadhan, S.P., Yang, K.: On the (im)possibility of obfuscating programs. In: Kilian, J. (ed.) CRYPTO 2001. LNCS, vol. 2139, pp. 1–18. Springer Berlin Heidelberg, Germany, Santa Barbara, CA, USA (Aug 19–23, 2001). https://doi.org/10.1007/3-540-44647-8_1
- [35] Bellare, M., Hoang, V.T., Rogaway, P.: Adaptively secure garbling with applications to one-time programs and secure outsourcing. In: Wang, X., Sako, K. (eds.) ASIACRYPT 2012. LNCS, vol. 7658, pp. 134–153. Springer Berlin Heidelberg, Germany, Beijing, China (Dec 2–6, 2012). https://doi.org/10.1007/978-3-642-34961-4_10
- [36] Bellare, M., Hoang, V.T., Rogaway, P.: Foundations of garbled circuits. In: Yu, T., Danezis, G., Gligor, V.D. (eds.) ACM CCS 2012. pp. 784–796. ACM Press, Raleigh, NC, USA (Oct 16–18, 2012). <https://doi.org/10.1145/2382196.2382279>

- [37] Bellare, M., Rogaway, P.: Random oracles are practical: A paradigm for designing efficient protocols. In: Proceedings of the 1st ACM Conference on Computer and Communications Security. pp. 62–73 (1993). <https://doi.org/10.1145/168588.168596>
- [38] Bethencourt, J., Sahai, A., Waters, B.: Ciphertext-policy attribute-based encryption. In: IEEE Symposium on Security and Privacy. pp. 321–334 (2007). <https://doi.org/10.1109/SP.2007.11>
- [39] Bitansky, N., Paneth, O., Rosen, A.: On the cryptographic hardness of finding a Nash equilibrium. In: Guruswami, V. (ed.) 56th FOCS. pp. 1480–1498. IEEE Computer Society Press, Berkeley, CA, USA (Oct 17–20, 2015). <https://doi.org/10.1109/FOCS.2015.94>
- [40] Bitansky, N., Vaikuntanathan, V.: Indistinguishability obfuscation from functional encryption. FOCS (2015). <https://doi.org/10.1109/FOCS.2015.20>, <http://eprint.iacr.org/2015/163>
- [41] Bitansky, N., Vaikuntanathan, V.: Indistinguishability obfuscation from functional encryption. Journal of the ACM **65**(6), 39:1–39:37 (2018). <https://doi.org/10.1145/3234511>
- [42] Boneh, D., Franklin, M.K.: Identity-based encryption from the Weil pairing. In: Kilian, J. (ed.) CRYPTO 2001. LNCS, vol. 2139, pp. 213–229. Springer Berlin Heidelberg, Germany, Santa Barbara, CA, USA (Aug 19–23, 2001). https://doi.org/10.1007/3-540-44647-8_13
- [43] Boneh, D., Gentry, C., Gorbunov, S., Halevi, S., Nikolaenko, V., Segev, G., Vaikuntanathan, V., Vinayagamurthy, D.: Fully key-homomorphic encryption, arithmetic circuit ABE and compact garbled circuits. In: Nguyen, P.Q., Oswald, E. (eds.) EUROCRYPT 2014. LNCS, vol. 8441, pp. 533–556. Springer Berlin Heidelberg, Germany, Copenhagen, Denmark (May 11–15, 2014). https://doi.org/10.1007/978-3-642-55220-5_30
- [44] Boneh, D., Gentry, C., Waters, B.: Collusion resistant broadcast encryption with short ciphertexts and private keys. In: CRYPTO (2005). https://doi.org/10.1007/11535218_16
- [45] Boneh, D., Lynn, B., Shacham, H.: Short signatures from the Weil pairing. In: Boyd, C. (ed.) ASIACRYPT 2001. LNCS, vol. 2248, pp. 514–532. Springer Berlin Heidelberg, Germany, Gold Coast, Australia (Dec 9–13, 2001). https://doi.org/10.1007/3-540-45682-1_30
- [46] Boneh, D., Sahai, A., Waters, B.: Functional encryption: Definitions and challenges. In: Theory of Cryptography: 8th Theory of Cryptography Conference, TCC 2011, Providence, RI, USA, March 28–30, 2011. Proceedings 8. pp. 253–273.

- Springer (2011). https://doi.org/10.1007/978-3-642-19571-6_16
- [47] Boneh, D., Waters, B.: Constrained pseudorandom functions and their applications. In: Sako, K., Sarkar, P. (eds.) ASIACRYPT 2013, Part II. LNCS, vol. 8270, pp. 280–300. Springer Berlin Heidelberg, Germany, Bangalore, India (Dec 1–5, 2013). https://doi.org/10.1007/978-3-642-42045-0_15
 - [48] Boneh, D., Waters, B., Zhandry, M.: Low overhead broadcast encryption from multilinear maps. In: CRYPTO (2014). https://doi.org/10.1007/978-3-662-44371-2_12
 - [49] Boneh, D., Zhandry, M.: Multiparty key exchange, efficient traitor tracing, and more from indistinguishability obfuscation. *Algorithmica* **79**(4), 1233–1285 (2017). https://doi.org/10.1007/978-3-662-44371-2_27
 - [50] Boyen, X., Waters, B.: Anonymous hierarchical identity-based encryption (without random oracles). In: CRYPTO (2006). https://doi.org/10.1007/11818175_17
 - [51] Boyle, E., Goldwasser, S., Ivan, I.: Functional signatures and pseudorandom functions. In: Krawczyk, H. (ed.) PKC 2014. LNCS, vol. 8383, pp. 501–519. Springer Berlin Heidelberg, Germany, Buenos Aires, Argentina (Mar 26–28, 2014). https://doi.org/10.1007/978-3-642-54631-0_29
 - [52] Brakerski, Z.: Fully homomorphic encryption without modulus switching from classical GapSVP. In: Annual Cryptology Conference. pp. 868–886. Springer (2012). https://doi.org/10.1007/978-3-642-32009-5_50
 - [53] Brakerski, Z., Döttling, N., Garg, S., Malavolta, G.: Factoring and pairings are not necessary for io: Circular-secure LWE suffices. *Cryptology ePrint Archive* (2020)
 - [54] Brakerski, Z., Döttling, N., Garg, S., Malavolta, G.: Candidate iO from homomorphic encryption schemes. *Journal of Cryptology* **36**(3), 27 (2023). <https://doi.org/10.1007/s00145-023-09471-5>
 - [55] Brakerski, Z., Gentry, C., Vaikuntanathan, V.: (Leveled) fully homomorphic encryption without bootstrapping. *ACM Transactions on Computation Theory (TOCT)* **6**(3), 1–36 (2014). <https://doi.org/10.1145/2090236.2090262>
 - [56] Brakerski, Z., Komargodski, I., Segev, G.: Multi-input functional encryption in the private-key setting: Stronger security from weaker assumptions. In: EUROCRYPT. Springer (2016). https://doi.org/10.1007/978-3-662-49896-5_30
 - [57] Brakerski, Z., Langlois, A., Peikert, C., Regev, O., Stehlé, D.: Classical hardness of learning with errors. In: Boneh, D., Roughgarden, T., Feigenbaum, J. (eds.) 45th ACM STOC. pp. 575–584. ACM Press, Palo Alto, CA, USA (Jun 1–4, 2013). <https://doi.org/10.1145/2488608.2488680>

- [58] Brakerski, Z., Lombardi, A., Segev, G., Vaikuntanathan, V.: Anonymous IBE, leakage resilience and circular security from new assumptions. In: Nielsen, J.B., Rijmen, V. (eds.) EUROCRYPT 2018, Part I. LNCS, vol. 10820, pp. 535–564. Springer, Cham, Switzerland, Tel Aviv, Israel (Apr 29 – May 3, 2018). https://doi.org/10.1007/978-3-319-78381-9_20
- [59] Brakerski, Z., Segev, G.: Function-private functional encryption in the private-key setting. *Journal of Cryptology* **31**(1), 202–225 (2018). <https://doi.org/10.1007/s00145-017-9255-y>
- [60] Brakerski, Z., Tsabary, R., Vaikuntanathan, V., Wee, H.: Private constrained PRFs (and more) from LWE. In: Kalai, Y., Reyzin, L. (eds.) TCC 2017, Part I. LNCS, vol. 10677, pp. 264–302. Springer, Cham, Switzerland, Baltimore, MD, USA (Nov 12–15, 2017). https://doi.org/10.1007/978-3-319-70500-2_10
- [61] Brakerski, Z., Vaikuntanathan, V.: Efficient fully homomorphic encryption from (standard) LWE. In: Ostrovsky, R. (ed.) 52nd FOCS. pp. 97–106. IEEE Computer Society Press, Palm Springs, CA, USA (Oct 22–25, 2011). <https://doi.org/10.1109/FOCS.2011.12>
- [62] Brakerski, Z., Vaikuntanathan, V.: Fully homomorphic encryption from ring-LWE and security for key dependent messages. In: Rogaway, P. (ed.) CRYPTO 2011. LNCS, vol. 6841, pp. 505–524. Springer Berlin Heidelberg, Germany, Santa Barbara, CA, USA (Aug 14–18, 2011). https://doi.org/10.1007/978-3-642-22792-9_29
- [63] Brakerski, Z., Vaikuntanathan, V.: Efficient fully homomorphic encryption from (standard) LWE. *SIAM Journal on computing* **43**(2), 831–871 (2014). <https://doi.org/10.1109/FOCS.2011.12>
- [64] Brakerski, Z., Vaikuntanathan, V.: Lattice-inspired broadcast encryption and succinct ciphertext-policy ABE. In: 13th Innovations in Theoretical Computer Science Conference (ITCS 2022). Schloss Dagstuhl-Leibniz-Zentrum für Informatik (2022)
- [65] Branco, P., Döttling, N., Jain, A., Malavolta, G., Mathialagan, S., Peters, S., Vaikuntanathan, V.: Pseudorandom Obfuscation and Applications. *Cryptology ePrint Archive*, Paper 2024/1742 (2024), <https://eprint.iacr.org/2024/1742>
- [66] Brzuska, C., Ünal, A., Woo, I.K.: Evasive LWE assumptions: Definitions, Classes, and Counterexamples. In: International Conference on the Theory and Application of Cryptology and Information Security. pp. 418–449. Springer (2024). https://doi.org/10.1007/978-981-96-0894-2_14
- [67] Brzuska, C., Farshim, P., Mittelbach, A.: Indistinguishability obfuscation and

- UCEs: The case of computationally unpredictable sources. In: Garay, J.A., Gennaro, R. (eds.) CRYPTO 2014, Part I. LNCS, vol. 8616, pp. 188–205. Springer Berlin Heidelberg, Germany, Santa Barbara, CA, USA (Aug 17–21, 2014). https://doi.org/10.1007/978-3-662-44371-2_11
- [68] Canetti, R., Goldreich, O., Halevi, S.: The random oracle methodology, revisited. *J. ACM* **51**(4), 557–594 (2004). <https://doi.org/10.1145/1008731.1008734>
- [69] Canetti, R., Rothblum, G.N., Varia, M.: Obfuscation of hyperplane membership. In: Micciancio, D. (ed.) TCC 2010. LNCS, vol. 5978, pp. 72–89. Springer Berlin Heidelberg, Germany, Zurich, Switzerland (Feb 9–11, 2010). https://doi.org/10.1007/978-3-642-11799-2_5
- [70] Caro, A.D., Iovino, V., Jain, A., O’Neill, A., Paneth, O., Persiano, G.: On the achievability of simulation-based security for functional encryption. In: CRYPTO (2013). https://doi.org/10.1007/978-3-642-40084-1_29
- [71] Cash, D., Hofheinz, D., Kiltz, E., Peikert, C.: Bonsai trees, or how to delegate a lattice basis. In: Gilbert, H. (ed.) EUROCRYPT 2010. LNCS, vol. 6110, pp. 523–552. Springer Berlin Heidelberg, Germany, French Riviera (May 30 – Jun 3, 2010). https://doi.org/10.1007/978-3-642-13190-5_27
- [72] Chen, Y., Vaikuntanathan, V., Waters, B., Wee, H., Wichs, D.: Traitor-tracing from LWE made simple and attribute-based. In: Beimel, A., Dziembowski, S. (eds.) TCC 2018, Part II. LNCS, vol. 11240, pp. 341–369. Springer, Cham, Switzerland, Panaji, India (Nov 11–14, 2018). https://doi.org/10.1007/978-3-030-03810-6_13
- [73] Cho, C., Döttling, N., Garg, S., Gupta, D., Miao, P., Polychroniadou, A.: Laconic oblivious transfer and its applications. In: Katz, J., Shacham, H. (eds.) CRYPTO 2017, Part II. LNCS, vol. 10402, pp. 33–65. Springer, Cham, Switzerland, Santa Barbara, CA, USA (Aug 20–24, 2017). https://doi.org/10.1007/978-3-319-63715-0_2
- [74] Chor, B., Fiat, A., Naor, M.: Tracing traitors. In: CRYPTO (1994). https://doi.org/10.1007/3-540-48658-5_25
- [75] Chotard, J., Dufour Sans, E., Gay, R., Phan, D.H., Pointcheval, D.: Decentralized multi-client functional encryption for inner product. In: Peyrin, T., Galbraith, S. (eds.) ASIACRYPT 2018, Part II. LNCS, vol. 11273, pp. 703–732. Springer, Cham, Switzerland, Brisbane, Queensland, Australia (Dec 2–6, 2018). https://doi.org/10.1007/978-3-030-03329-3_24
- [76] Clear, M., McGoldrick, C.: Multi-identity and multi-key leveled FHE from learning with errors. In: Gennaro, R., Robshaw, M.J.B. (eds.) CRYPTO 2015, Part II. LNCS, vol. 9216, pp. 630–656. Springer Berlin Heidelberg, Germany,

- Santa Barbara, CA, USA (Aug 16–20, 2015). https://doi.org/10.1007/978-3-662-48000-7_31
- [77] Cocks, C.: An Identity Based Encryption Scheme Based on Quadratic Residues. In: IMA Int. Conf. (2001). https://doi.org/10.1007/3-540-45325-3_32
 - [78] Cohen, A., Holmgren, J., Nishimaki, R., Vaikuntanathan, V., Wichs, D.: Watermarking cryptographic capabilities. In: Wichs, D., Mansour, Y. (eds.) 48th ACM STOC. pp. 1115–1127. ACM Press, Cambridge, MA, USA (Jun 18–21, 2016). <https://doi.org/10.1145/2897518.2897651>
 - [79] Coron, J.S.: On the exact security of full domain hash. In: Bellare, M. (ed.) CRYPTO 2000. LNCS, vol. 1880, pp. 229–235. Springer Berlin Heidelberg, Germany, Santa Barbara, CA, USA (Aug 20–24, 2000). https://doi.org/10.1007/3-540-44598-6_14
 - [80] Coron, J., Lepoint, T., Tibouchi, M.: Practical multilinear maps over the integers. In: CRYPTO (2013). https://doi.org/10.1007/978-3-642-40041-4_26
 - [81] Datta, P., Okamoto, T., Tomida, J.: Full-hiding (unbounded) multi-input inner product functional encryption from the k -Linear assumption. In: Abdalla, M., Dahab, R. (eds.) PKC 2018, Part II. LNCS, vol. 10770, pp. 245–277. Springer, Cham, Switzerland, Rio de Janeiro, Brazil (Mar 25–29, 2018). https://doi.org/10.1007/978-3-319-76581-5_9
 - [82] Devadas, L., Quach, W., Vaikuntanathan, V., Wee, H., Wichs, D.: Succinct LWE sampling, random polynomials, and obfuscation. In: Theory of Cryptography: 19th International Conference, TCC 2021, Raleigh, NC, USA, November 8–11, 2021, Proceedings, Part II 19. pp. 256–287. Springer (2021). https://doi.org/10.1007/978-3-030-90453-1_9
 - [83] Döttling, N., Garg, S.: Identity-based encryption from the Diffie-Hellman assumption. In: Katz, J., Shacham, H. (eds.) CRYPTO 2017, Part I. LNCS, vol. 10401, pp. 537–569. Springer, Cham, Switzerland, Santa Barbara, CA, USA (Aug 20–24, 2017). https://doi.org/10.1007/978-3-319-63688-7_18
 - [84] Döttling, N., Garg, S., Hajiabadi, M., Masny, D.: New constructions of identity-based and key-dependent message secure encryption schemes. In: Abdalla, M., Dahab, R. (eds.) PKC 2018, Part I. LNCS, vol. 10769, pp. 3–31. Springer, Cham, Switzerland, Rio de Janeiro, Brazil (Mar 25–29, 2018). https://doi.org/10.1007/978-3-319-76578-5_1
 - [85] Döttling, N., Jain, A., Malavolta, G., Mathialagan, S., Vaikuntanathan, V.: Simple and General Counterexamples for Private-Coin Evasive LWE. Cryptology ePrint Archive (2025)

- [86] Fiat, A., Naor, M.: Broadcast encryption. In: Stinson, D.R. (ed.) CRYPTO (1993). https://doi.org/10.1007/3-540-48329-2_40
- [87] Francati, D., Friolo, D., Malavolta, G., Venturi, D.: Multi-key and Multi-input Predicate Encryption (for Conjunctions) from Learning with Errors. *Journal of Cryptology* **37**(3), 24 (2024). <https://doi.org/10.1007/s00145-024-09504-7>
- [88] Freire, E.S.V., Hofheinz, D., Paterson, K.G., Striecks, C.: Programmable hash functions in the multilinear setting. In: Canetti, R., Garay, J.A. (eds.) CRYPTO 2013, Part I. LNCS, vol. 8042, pp. 513–530. Springer Berlin Heidelberg, Germany, Santa Barbara, CA, USA (Aug 18–22, 2013). https://doi.org/10.1007/978-3-642-40041-4_28
- [89] Garg, S., Gentry, C., Halevi, S.: Candidate multilinear maps from ideal lattices. In: EUROCRYPT (2013). https://doi.org/10.1007/978-3-642-38348-9_1
- [90] Garg, S., Gentry, C., Halevi, S., Raykova, M., Sahai, A., Waters, B.: Candidate indistinguishability obfuscation and functional encryption for all circuits. In: 54th FOCS. pp. 40–49. IEEE Computer Society Press, Berkeley, CA, USA (Oct 26–29, 2013). <https://doi.org/10.1109/FOCS.2013.13>
- [91] Garg, S., Gentry, C., Halevi, S., Raykova, M., Sahai, A., Waters, B.: Candidate indistinguishability obfuscation and functional encryption for all circuits. *SIAM Journal on Computing* **45**(3), 882–929 (2016)
- [92] Garg, S., Pandey, O., Srinivasan, A.: Revisiting the cryptographic hardness of finding a nash equilibrium. In: Robshaw, M., Katz, J. (eds.) CRYPTO 2016, Part II. LNCS, vol. 9815, pp. 579–604. Springer Berlin Heidelberg, Germany, Santa Barbara, CA, USA (Aug 14–18, 2016). https://doi.org/10.1007/978-3-662-53008-5_20
- [93] Garg, S., Pandey, O., Srinivasan, A., Zhandry, M.: Breaking the sub-exponential barrier in obfustopia. In: EUROCRYPT (2017). https://doi.org/10.1007/978-3-319-56617-7_6
- [94] Garg, S., Srinivasan, A.: Single-key to multi-key functional encryption with polynomial loss. In: TCC (2016). https://doi.org/10.1007/978-3-662-53644-5_16
- [95] Gay, R., Jain, A., Lin, H., Sahai, A.: Indistinguishability obfuscation from simple-to-state hard problems: New assumptions, new techniques, and simplification. In: EUROCRYPT (2021). https://doi.org/10.1007/978-3-030-77883-5_4
- [96] Gay, R., Pass, R.: Indistinguishability obfuscation from circular security. In: Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory

- of Computing. pp. 736–749 (2021). <https://doi.org/10.1145/3406325.3451070>
- [97] Gentry, C.: Fully homomorphic encryption using ideal lattices. In: Proceedings of the forty-first annual ACM symposium on Theory of computing. pp. 169–178 (2009). <https://doi.org/10.1145/1536414.1536440>
 - [98] Gentry, C., Peikert, C., Vaikuntanathan, V.: Trapdoors for hard lattices and new cryptographic constructions. In: Ladner, R.E., Dwork, C. (eds.) 40th ACM STOC. pp. 197–206. ACM Press, Victoria, BC, Canada (May 17–20, 2008). <https://doi.org/10.1145/1374376.1374407>
 - [99] Gentry, C., Sahai, A., Waters, B.: Homomorphic encryption from learning with errors: Conceptually-simpler, asymptotically-faster, attribute-based. In: Canetti, R., Garay, J.A. (eds.) CRYPTO 2013, Part I. LNCS, vol. 8042, pp. 75–92. Springer Berlin Heidelberg, Germany, Santa Barbara, CA, USA (Aug 18–22, 2013). https://doi.org/10.1007/978-3-642-40041-4_5
 - [100] Goldreich, O., Goldwasser, S., Micali, S.: How to construct random functions (extended abstract). In: 25th FOCS. pp. 464–479. IEEE Computer Society Press, Singer Island, Florida (Oct 24–26, 1984). <https://doi.org/10.1109/SFCS.1984.715949>
 - [101] Goldwasser, S., Gordon, S.D., Goyal, V., Jain, A., Katz, J., Liu, F.H., Sahai, A., Shi, E., Zhou, H.S.: Multi-input functional encryption. In: Nguyen, P.Q., Oswald, E. (eds.) EUROCRYPT 2014. LNCS, vol. 8441, pp. 578–602. Springer Berlin Heidelberg, Germany, Copenhagen, Denmark (May 11–15, 2014). https://doi.org/10.1007/978-3-642-55220-5_32
 - [102] Goldwasser, S., Kalai, Y.T., Popa, R.A., Vaikuntanathan, V., Zeldovich, N.: How to run Turing machines on encrypted data. In: Canetti, R., Garay, J.A. (eds.) CRYPTO 2013, Part II. LNCS, vol. 8043, pp. 536–553. Springer Berlin Heidelberg, Germany, Santa Barbara, CA, USA (Aug 18–22, 2013). https://doi.org/10.1007/978-3-642-40084-1_30
 - [103] Gong, J., Waters, B., Wee, H.: ABE for DFA from k-Lin. In: Advances in Cryptology–CRYPTO 2019: 39th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 18–22, 2019, Proceedings, Part II 39. pp. 732–764. Springer (2019). https://doi.org/10.1007/978-3-030-26951-7_25
 - [104] Gong, J., Wee, H.: Adaptively secure ABE for DFA from k-Lin and more. In: Annual International Conference on the Theory and Applications of Cryptographic Techniques. pp. 278–308. Springer (2020). https://doi.org/10.1007/978-3-030-45727-3_10
 - [105] González, A., Zacharakis, A.: Fully-succinct publicly verifiable delegation from

- constant-size assumptions. In: Theory of Cryptography Conference. pp. 529–557. Springer (2021). https://doi.org/10.1007/978-3-030-90459-3_18
- [106] Gorbunov, S., Vaikuntanathan, V., Wee, H.: Functional encryption with bounded collusions via multi-party computation. In: Safavi-Naini, R., Canetti, R. (eds.) CRYPTO 2012. LNCS, vol. 7417, pp. 162–179. Springer Berlin Heidelberg, Germany, Santa Barbara, CA, USA (Aug 19–23, 2012). https://doi.org/10.1007/978-3-642-32009-5_11
 - [107] Gorbunov, S., Vaikuntanathan, V., Wee, H.: Attribute based encryption for circuits. In: STOC (2013). <https://doi.org/10.1145/2488608.2488677>
 - [108] Gorbunov, S., Vaikuntanathan, V., Wee, H.: Predicate encryption for circuits from LWE. In: Gennaro, R., Robshaw, M.J.B. (eds.) CRYPTO 2015, Part II. LNCS, vol. 9216, pp. 503–523. Springer Berlin Heidelberg, Germany, Santa Barbara, CA, USA (Aug 16–20, 2015). https://doi.org/10.1007/978-3-662-48000-7_25
 - [109] Gorbunov, S., Vaikuntanathan, V., Wichs, D.: Leveled fully homomorphic signatures from standard lattices. In: Proceedings of the forty-seventh annual ACM symposium on Theory of computing. pp. 469–477 (2015). <https://doi.org/10.1145/2746539.2746576>
 - [110] Goyal, R., Koppula, V., Waters, B.: Semi-adaptive security and bundling functionalities made generic and easy. In: Hirt, M., Smith, A.D. (eds.) TCC 2016-B, Part II. LNCS, vol. 9986, pp. 361–388. Springer Berlin Heidelberg, Germany, Beijing, China (Oct 31 – Nov 3, 2016). https://doi.org/10.1007/978-3-662-53644-5_14
 - [111] Goyal, R., Koppula, V., Waters, B.: Lockable obfuscation. In: Umans, C. (ed.) 58th FOCS. pp. 612–621. IEEE Computer Society Press, Berkeley, CA, USA (Oct 15–17, 2017). <https://doi.org/10.1109/FOCS.2017.62>
 - [112] Goyal, R., Koppula, V., Waters, B.: Collusion resistant traitor tracing from learning with errors. In: STOC (2018). <https://doi.org/10.1145/3188745.3188844>
 - [113] Goyal, R., Koppula, V., Waters, B.: New approaches to traitor tracing with embedded identities. In: TCC (2019). https://doi.org/10.1007/978-3-030-36033-7_6
 - [114] Goyal, R., Quach, W., Waters, B., Wichs, D.: Broadcast and trace with n^ϵ ciphertext size from standard assumptions. In: Crypto (2019), <https://eprint.iacr.org/2019/636>
 - [115] Goyal, R., Vusirikala, S., Waters, B.: Collusion resistant broadcast and trace from positional witness encryption. In: PKC (2019). <https://doi.org/10.1007/>

- [116] Goyal, V., Pandey, O., Sahai, A., Waters, B.: Attribute-based encryption for fine-grained access control of encrypted data. In: Juels, A., Wright, R.N., De Capitani di Vimercati, S. (eds.) ACM CCS 2006. pp. 89–98. ACM Press, Alexandria, Virginia, USA (Oct 30 – Nov 3, 2006). <https://doi.org/10.1145/1180405.1180418>, available as Cryptology ePrint Archive Report 2006/309
- [117] Hofheinz, D., Jager, T., Khurana, D., Sahai, A., Waters, B., Zhandry, M.: How to generate and use universal samplers. In: Cheon, J.H., Takagi, T. (eds.) ASIACRYPT 2016, Part II. LNCS, vol. 10032, pp. 715–744. Springer Berlin Heidelberg, Germany, Hanoi, Vietnam (Dec 4–8, 2016). https://doi.org/10.1007/978-3-662-53890-6_24
- [118] Hohenberger, S., Sahai, A., Waters, B.: Full domain hash from (leveled) multilinear maps and identity-based aggregate signatures. In: Canetti, R., Garay, J.A. (eds.) CRYPTO 2013, Part I. LNCS, vol. 8042, pp. 494–512. Springer Berlin Heidelberg, Germany, Santa Barbara, CA, USA (Aug 18–22, 2013). https://doi.org/10.1007/978-3-642-40041-4_27
- [119] Hohenberger, S., Sahai, A., Waters, B.: Replacing a random oracle: Full domain hash from indistinguishability obfuscation. In: Nguyen, P.Q., Oswald, E. (eds.) EUROCRYPT 2014. LNCS, vol. 8441, pp. 201–220. Springer Berlin Heidelberg, Germany, Copenhagen, Denmark (May 11–15, 2014). https://doi.org/10.1007/978-3-642-55220-5_12
- [120] Hopkins, S.B., Jain, A., Lin, H.: Counterexamples to new circular security assumptions underlying iO. In: Malkin, T., Peikert, C. (eds.) CRYPTO 2021, Part II. LNCS, vol. 12826, pp. 673–700. Springer, Cham, Switzerland, Virtual Event (Aug 16–20, 2021). https://doi.org/10.1007/978-3-030-84245-1_23
- [121] Hsieh, Y.C., Jain, A., Lin, H.: Lattice-Based Post-Quantum iO from Circular Security with Random Opening Assumption (Part II: zeroizing attacks against private-coin evasive LWE assumptions). Cryptology ePrint Archive (2025)
- [122] Hsieh, Y.C., Lin, H., Luo, J.: Attribute-based encryption for circuits of unbounded depth from lattices. In: 2023 IEEE 64th Annual Symposium on Foundations of Computer Science (FOCS). pp. 415–434. IEEE (2023). <https://doi.org/10.1109/FOCS57990.2023.00031>
- [123] Hsieh, Y.C., Lin, H., Luo, J.: A general framework for lattice-based ABE using evasive inner-product functional encryption. In: Joye, M., Leander, G. (eds.) EUROCRYPT 2024, Part II. LNCS, vol. 14652, pp. 433–464. Springer, Cham, Switzerland, Zurich, Switzerland (May 26–30, 2024). https://doi.org/10.1007/978-3-031-58723-8_15

- [124] Huang, T.H., Hung, W.H., Yamada, S.: A Note on Obfuscation-based Attacks on Private-coin Evasive LWE. Cryptology ePrint Archive (2025)
- [125] Hubáček, P., Yorgev, E.: Hardness of continuous local search: Query complexity and cryptographic lower bounds. In: Klein, P.N. (ed.) 28th SODA. pp. 1352–1371. ACM-SIAM, Barcelona, Spain (Jan 16–19, 2017). <https://doi.org/10.1137/1.9781611974782.88>
- [126] Jain, A., Lin, H., Lou, P., Sahai, A.: Polynomial-time cryptanalysis of the subspace flooding assumption for post-quantum *iO*. In: Hazay, C., Stam, M. (eds.) EUROCRYPT 2023, Part I. LNCS, vol. 14004, pp. 205–235. Springer, Cham, Switzerland, Lyon, France (Apr 23–27, 2023). https://doi.org/10.1007/978-3-031-30545-0_8
- [127] Jain, A., Lin, H., Luo, J.: On the optimal succinctness and efficiency of functional encryption and attribute-based encryption. In: Hazay, C., Stam, M. (eds.) EUROCRYPT 2023, Part III. LNCS, vol. 14006, pp. 479–510. Springer, Cham, Switzerland, Lyon, France (Apr 23–27, 2023). https://doi.org/10.1007/978-3-031-30620-4_16
- [128] Jain, A., Lin, H., Matt, C., Sahai, A.: How to Leverage Hardness of Constant-Degree Expanding Polynomials over \mathbb{R} to build *iO*. In: EUROCRYPT (2019). https://doi.org/10.1007/978-3-030-17653-2_9
- [129] Jain, A., Lin, H., Sahai, A.: Indistinguishability obfuscation from well-founded assumptions. In: Khuller, S., Williams, V.V. (eds.) 53rd ACM STOC. pp. 60–73. ACM Press, Virtual Event, Italy (Jun 21–25, 2021). <https://doi.org/10.1145/3406325.3451093>
- [130] Jain, A., Lin, H., Sahai, A.: Indistinguishability obfuscation from LPN over \mathbb{F}_p , DLIN, and PRGs in NC^0 . In: Dunkelman, O., Dziembowski, S. (eds.) EUROCRYPT 2022, Part I. LNCS, vol. 13275, pp. 670–699. Springer, Cham, Switzerland, Trondheim, Norway (May 30 – Jun 3, 2022). https://doi.org/10.1007/978-3-031-06944-4_23
- [131] Jin, Z., Kalai, Y.T., Lombardi, A., Mathialagan, S.: Universal SNARGs for NP from Proofs of Correctness. Cryptology ePrint Archive, Paper 2024/2015 (2024), <https://eprint.iacr.org/2024/2015>
- [132] Kalai, Y., Lombardi, A., Vaikuntanathan, V., Wichs, D.: Boosting batch arguments and RAM delegation. In: Proceedings of the 55th Annual ACM Symposium on Theory of Computing. pp. 1545–1552 (2023). <https://doi.org/10.1145/3564246.3585200>
- [133] Katz, J., Sahai, A., Waters, B.: Predicate encryption supporting disjunctions, polynomial equations, and inner products. Journal of cryptology **26**, 191–224

- (2013). <https://doi.org/10.1007/s00145-012-9119-4>
- [134] Kiayias, A., Papadopoulos, S., Triandopoulos, N., Zacharias, T.: Delegatable pseudorandom functions and applications. In: Sadeghi, A.R., Gligor, V.D., Yung, M. (eds.) ACM CCS 2013. pp. 669–684. ACM Press, Berlin, Germany (Nov 4–8, 2013). <https://doi.org/10.1145/2508859.2516668>
 - [135] Kim, S., Wu, D.J.: Collusion resistant trace-and-revoke for arbitrary identities from standard assumptions. In: ASIACRYPT (2020). https://doi.org/10.1007/978-3-030-64834-3_3
 - [136] Kitagawa, F., Nishimaki, R., Tanaka, K., Yamakawa, T.: Adaptively secure and succinct functional encryption: Improving security and efficiency, simultaneously. In: Boldyreva, A., Micciancio, D. (eds.) CRYPTO 2019, Part III. LNCS, vol. 11694, pp. 521–551. Springer, Cham, Switzerland, Santa Barbara, CA, USA (Aug 18–22, 2019). https://doi.org/10.1007/978-3-030-26954-8_17
 - [137] Kitagawa, F., Tanaka, K.: Key dependent message security and receiver selective opening security for identity-based encryption. In: Abdalla, M., Dahab, R. (eds.) PKC 2018, Part I. LNCS, vol. 10769, pp. 32–61. Springer, Cham, Switzerland, Rio de Janeiro, Brazil (Mar 25–29, 2018). https://doi.org/10.1007/978-3-319-76578-5_2
 - [138] Komargodski, I., Segev, G.: From minicrypt to obfustopia via private-key functional encryption. In: Coron, J.S., Nielsen, J.B. (eds.) EUROCRYPT 2017, Part I. LNCS, vol. 10210, pp. 122–151. Springer, Cham, Switzerland, Paris, France (Apr 30–May 4, 2017). https://doi.org/10.1007/978-3-319-56620-7_5
 - [139] Komargodski, I., Segev, G.: From minicrypt to obfustopia via private-key functional encryption. *Journal of Cryptology* **33**(2), 406–458 (2020). https://doi.org/10.1007/978-3-319-56620-7_5
 - [140] Koppula, V., Lewko, A.B., Waters, B.: Indistinguishability obfuscation for Turing machines with unbounded memory. In: Servedio, R.A., Rubinfeld, R. (eds.) 47th ACM STOC. pp. 419–428. ACM Press, Portland, OR, USA (Jun 14–17, 2015). <https://doi.org/10.1145/2746539.2746614>
 - [141] Libert, B., Titu, R.: Multi-client functional encryption for linear functions in the standard model from LWE. In: Galbraith, S.D., Moriai, S. (eds.) ASIACRYPT 2019, Part III. LNCS, vol. 11923, pp. 520–551. Springer, Cham, Switzerland, Kobe, Japan (Dec 8–12, 2019). https://doi.org/10.1007/978-3-030-34618-8_18
 - [142] Lin, H., Luo, J.: Compact adaptively secure ABE from k -Lin: Beyond NC^1 and towards NL. In: Canteaut, A., Ishai, Y. (eds.) EUROCRYPT 2020, Part III. LNCS, vol. 12107, pp. 247–277. Springer, Cham, Switzerland, Zagreb, Croatia

- (May 10–14, 2020). https://doi.org/10.1007/978-3-030-45727-3_9
- [143] Lin, H., Luo, J.: Succinct and adaptively secure ABE for ABP from k -Lin. In: Moriai, S., Wang, H. (eds.) ASIACRYPT 2020, Part III. LNCS, vol. 12493, pp. 437–466. Springer, Cham, Switzerland, Daejeon, South Korea (Dec 7–11, 2020). https://doi.org/10.1007/978-3-030-64840-4_15
 - [144] Lin, H., Pass, R., Seth, K., Telang, S.: Indistinguishability obfuscation with non-trivial efficiency. In: Public-Key Cryptography–PKC 2016, pp. 447–462. Springer (2016). https://doi.org/10.1007/978-3-662-49387-8_17
 - [145] Lindell, Y., Pinkas, B.: A proof of security of Yao’s protocol for two-party computation. *Journal of Cryptology* **22**(2), 161–188 (Apr 2009). <https://doi.org/10.1007/s00145-008-9036-8>
 - [146] Mathialagan, S., Peters, S., Vaikuntanathan, V.: Adaptively Sound Zero-Knowledge Snarks for UP. In: Annual International Cryptology Conference. pp. 38–71. Springer (2024). https://doi.org/10.1007/978-3-031-68403-6_2
 - [147] Mathialagan, S., Peters, S., Vaikuntanathan, V.: Using PRTT Obfuscation for Unlevelled FHE. Personal Communication (2024)
 - [148] Micciancio, D., Peikert, C.: Trapdoors for lattices: Simpler, tighter, faster, smaller. In: Pointcheval, D., Johansson, T. (eds.) EUROCRYPT 2012. LNCS, vol. 7237, pp. 700–718. Springer Berlin Heidelberg, Germany, Cambridge, UK (Apr 15–19, 2012). https://doi.org/10.1007/978-3-642-29011-4_41
 - [149] Mukherjee, P., Wichs, D.: Two round multiparty computation via multi-key FHE. In: Fischlin, M., Coron, J.S. (eds.) EUROCRYPT 2016, Part II. LNCS, vol. 9666, pp. 735–763. Springer Berlin Heidelberg, Germany, Vienna, Austria (May 8–12, 2016). https://doi.org/10.1007/978-3-662-49896-5_26
 - [150] Naor, M., Pinkas, B.: Efficient trace and revoke schemes. *International Journal of Information Security* **9**(6), 411–424 (2010). <https://doi.org/10.1007/s10207-010-0121-2>
 - [151] Nishimaki, R., Wichs, D., Zhandry, M.: Anonymous traitor tracing: How to embed arbitrary information in a key. In: Fischlin, M., Coron, J.S. (eds.) EUROCRYPT 2016, Part II. LNCS, vol. 9666, pp. 388–419. Springer Berlin Heidelberg, Germany, Vienna, Austria (May 8–12, 2016). https://doi.org/10.1007/978-3-662-49896-5_14
 - [152] Pippenger, N., Fischer, M.J.: Relations among complexity measures. *Journal of the ACM (JACM)* **26**(2) (1979). <https://doi.org/10.1145/322123.322138>
 - [153] Quach, W., Wee, H., Wichs, D.: Laconic function evaluation and applications. In: 2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS).

- pp. 859–870. IEEE (2018). <https://doi.org/10.1109/FOCS.2018.00086>
- [154] Ragavan, S., Vafa, N., Vaikuntanathan, V.: Indistinguishability Obfuscation from Bilinear Maps and LPN Variants. Cryptology ePrint Archive (2024)
 - [155] Regev, O.: On lattices, learning with errors, random linear codes, and cryptography. *Journal of the ACM* **56**(6), 34:1–34:40 (2009). <https://doi.org/10.1145/1568318.1568324>
 - [156] Sahai, A., Seyalioglu, H.: Worry-free encryption: functional encryption with public keys. In: Al-Shaer, E., Keromytis, A.D., Shmatikov, V. (eds.) *ACM CCS 2010*. pp. 463–472. ACM Press, Chicago, Illinois, USA (Oct 4–8, 2010). <https://doi.org/10.1145/1866307.1866359>
 - [157] Sahai, A., Waters, B.: Fuzzy Identity-Based Encryption. In: *EUROCRYPT* (2005). https://doi.org/10.1007/11426639_27
 - [158] Sahai, A., Waters, B.: How to use indistinguishability obfuscation: deniable encryption, and more. In: Shmoys, D.B. (ed.) *46th ACM STOC*. pp. 475–484. ACM Press, New York, NY, USA (May 31 – Jun 3, 2014). <https://doi.org/10.1145/2591796.2591825>
 - [159] Sakai, R., Ohgishi, K., Kasahara, M.: Cryptosystems based on pairing. *SCIS 2000* (2000)
 - [160] Takashima, K.: Expressive attribute-based encryption with constant-size ciphertexts from the decisional linear assumption. In: *SCN* (2014). https://doi.org/10.1007/978-3-319-10879-7_17
 - [161] Tomida, J.: Tightly secure inner product functional encryption: Multi-input and function-hiding constructions. In: Galbraith, S.D., Moriai, S. (eds.) *ASIACRYPT 2019, Part III*. LNCS, vol. 11923, pp. 459–488. Springer, Cham, Switzerland, Kobe, Japan (Dec 8–12, 2019). https://doi.org/10.1007/978-3-030-34618-8_16
 - [162] Tsabary, R.: Fully secure attribute-based encryption for t-CNF from LWE. In: *CRYPTO* (2019). https://doi.org/10.1007/978-3-030-26948-7_3
 - [163] Tsabary, R.: Candidate witness encryption from lattice techniques. In: Dodis, Y., Shrimpton, T. (eds.) *CRYPTO 2022, Part I*. LNCS, vol. 13507, pp. 535–559. Springer, Cham, Switzerland, Santa Barbara, CA, USA (Aug 15–18, 2022). https://doi.org/10.1007/978-3-031-15802-5_19
 - [164] Vaikuntanathan, V., Wee, H., Wichs, D.: Witness encryption and null-IO from evasive LWE. In: Agrawal, S., Lin, D. (eds.) *ASIACRYPT 2022, Part I*. LNCS, vol. 13791, pp. 195–221. Springer, Cham, Switzerland, Taipei, Taiwan (Dec 5–9, 2022). https://doi.org/10.1007/978-3-031-22963-3_7

- [165] Waters, B.: Functional encryption for regular languages. In: Safavi-Naini, R., Canetti, R. (eds.) CRYPTO 2012. LNCS, vol. 7417, pp. 218–235. Springer Berlin Heidelberg, Germany, Santa Barbara, CA, USA (Aug 19–23, 2012). https://doi.org/10.1007/978-3-642-32009-5_14
- [166] Waters, B., Wee, H., Wu, D.J.: Multi-authority ABE from lattices without random oracles. In: Kiltz, E., Vaikuntanathan, V. (eds.) TCC 2022, Part I. LNCS, vol. 13747, pp. 651–679. Springer, Cham, Switzerland, Chicago, IL, USA (Nov 7–10, 2022). https://doi.org/10.1007/978-3-031-22318-1_23
- [167] Wee, H.: On obfuscating point functions. In: Gabow, H.N., Fagin, R. (eds.) 37th ACM STOC. pp. 523–532. ACM Press, Baltimore, MA, USA (May 22–24, 2005). <https://doi.org/10.1145/1060590.1060669>
- [168] Wee, H.: Broadcast encryption with size $N^{1/3}$ and more from k -lin. In: Malkin, T., Peikert, C. (eds.) CRYPTO 2021, Part IV. LNCS, vol. 12828, pp. 155–178. Springer, Cham, Switzerland, Virtual Event (Aug 16–20, 2021). https://doi.org/10.1007/978-3-030-84259-8_6
- [169] Wee, H.: Optimal broadcast encryption and CP-ABE from evasive lattice assumptions. In: Dunkelman, O., Dziembowski, S. (eds.) EUROCRYPT 2022, Part II. LNCS, vol. 13276, pp. 217–241. Springer, Cham, Switzerland, Trondheim, Norway (May 30 – Jun 3, 2022). https://doi.org/10.1007/978-3-031-07085-3_8
- [170] Wee, H.: Circuit abe with $\text{poly}(\text{depth}, \lambda)$ -sized ciphertexts and keys from lattices. In: Crypto (2024). https://doi.org/10.1007/978-3-031-68382-4_6
- [171] Wee, H., Wichs, D.: Candidate obfuscation via oblivious LWE sampling. In: Canteaut, A., Standaert, F.X. (eds.) EUROCRYPT 2021, Part III. LNCS, vol. 12698, pp. 127–156. Springer, Cham, Switzerland, Zagreb, Croatia (Oct 17–21, 2021). https://doi.org/10.1007/978-3-030-77883-5_5
- [172] Wichs, D., Zirdelis, G.: Obfuscating compute-and-compare programs under LWE. In: Umans, C. (ed.) 58th FOCS. pp. 600–611. IEEE Computer Society Press, Berkeley, CA, USA (Oct 15–17, 2017). <https://doi.org/10.1109/FOCS.2017.61>
- [173] Yao, A.C.C.: Protocols for secure computations extended abstract. In: 23rd FOCS. vol. 28 (1982). <https://doi.org/10.1109/SFCS.1982.38>
- [174] Zhandry, M.: New techniques for traitor tracing: Size $N^{1/3}$ and more from pairings. In: CRYPTO (2020). https://doi.org/10.1007/978-3-030-56784-2_22

CURRICULUM VITAE

NAME Simran Kumari

DATE OF BIRTH 08 October 1996

EDUCATION QUALIFICATIONS

2019	M.Sc.		
	Institution	National Institute of Technology, Warangal	
	Specialization	Mathematics and Scientific Computing	
2017	B.Sc.		
	Institution	Delhi University, New Delhi	
	Specialization	Mathematics	

DOCTORAL COMMITTEE

Chairperson

Prof. Krishna Moorthy Sivalingam
CSE Department, IIT Madras

Guide

Prof. Shweta Agrawal
CSE Department, IIT Madras

Members

Prof. John Augustine
CSE Department, IIT Madras

Prof. Aishwarya Thiruvengadam
CSE Department, IIT Madras

Prof. Venkata Koppula
CSE Department, IIT Delhi