Research Statement

My research interests are in the area of wireless networking, sensing and Internet-of-Things (IoT) technologies. I design and implement systems that improve the performance of wireless networks and related systems. First, I intend to build such solutions on theoretically sound principles based on a good understanding of wireless signals and second, I put a major emphasis on making such solutions practically deployable. My research efforts have led to multiple contributions in different areas spanning dynamic spectrum sharing, admission control in wireless networks and radio frequency (RF) based localization. My research appeared in some of the premier venues in systems and networking, [MOBICOM '13, IMC '14, INFOCOM '15, INFOCOM '17, INFOCOM '18, CoNEXT '14, CoNEXT '16, CoNEXT '18, NSDI '19]. In the following, I present a brief overview of my past and current research. I base the discussion in two parts. The first part deals with my dissertation research and primarily addresses problems in the wireless networking space. The second part includes my current experience in an industrial research lab working with wireless sensing technologies.

Dissertation Research

Although majority of the devices today are wireless enabled, with the exponentiating growth in the number of 'connected' devices it is unclear how incremental changes to the current network infrastructure can sustain such growth. The problem stems from the fact that the capacity of a wireless network is fundamentally capped by limited availability of physical resources, e.g., wireless spectrum. With innovations in applications like 4K video, demands for good user experience and an exponentiating user base: such capacity is already close to hit the ceiling. For instance, even though the capacity of LTE networks today is close to the Shannon limit, it is still unclear how to meet additional capacity demands in near future with the surge of data traffic.

To have a holistic view of the problem, my dissertation research stands astride two extreme ends of the network protocol stack: (a) the Application layer and (b) the Physical layer. Essentially, the general direction of my research looks at optimizing resource utilization to improve capacity / performance in these two layers. I made contributions to understanding and characterizing Quality of Experience (QoE) of different types Internet applications (video streaming, web browsing, VoIP etc.) in heterogeneous network conditions [CellNet '13, IMC '14, CoNEXT '16]. Second, I contributed to augmenting wireless capacity by harvesting unused radio frequency spectrum or white spaces [MOBICOM '13, CoNEXT '14, Hot Wireless '16, INFOCOM '17, INFOCOM '18, TCCN '19]. Third, I have also worked on network-side positioning of cellular band devices [INFOCOM '15].

► **Research Approach:** Many of such problems can be solved using analytical (whitebox) techniques that depend on meticulous mathematical modeling of the concerned system. However the complexity and scale of today's systems deem such whitebox solutions to be rather impractical due to lack of availability of granular information. Inspired by the recent success of *data-driven* (blackbox) analysis in many fields of computing, my dissertation research explores such approaches in modeling and optimizing performance of networked systems. In collaboration with the industry, much of my research efforts translated to building actual system prototypes that were passed on to respective product teams.

■ **Characterizing Quality of Experience:** Today's mobile application ecosystem is a complicated mess. Various entangled dimensions add to such complexity: carriers, networks, client devices and of course diverse types of application programs. Amidst such complexity what matters the most is Quality of Experience (QoE) perceived by end users. Guaranteeing good QoE for the user is a challenge in mobile apps due to their diverse resource requirements and the resource-constrained, variable nature of wireless networks and mobile devices.

A key question we ask is: how to provision a network given a set of diverse applications and client devices, or how to optimize the overall user experience for multiple users and applications in a given network? We develop a user feedback-based system called `Adapp` that learns the best network to use for an app when multiple networks are available [CellNet '13]. `Adapp` optimizes QoE at an individual user level on the client-side. To complement this effort, we also develop a network-side system called `ExBox` (**Ex**perience middle**Box**) that uses a new experiential capacity-based formulation to aid admission control and network selection [CoNEXT '16]. `ExBox` can handle multiple different types of applications in the network and dynamic network conditions. Both `Adapp` and `ExBox` have been tested on real network testbeds. My dissertation also contributes towards extensive QoE measurement studies for multiple applications across different commercial network providers [IMC '14] and hardware vendors in specific networking contexts.

■ **Optimizing Shared Spectrum Access Technologies:** With the explosion of mobile data, there is a growing realization that RF spectrum must be treated as an important resource that is in limited supply. Policy makers and researchers alike are promoting various forms of spectrum sharing models to improve spectrum utilization. This has promoted new spectrum sharing regimes where a diverse set of wireless technologies, including broadcast television, radar and various forms of wireless communications, co-exist in the same spectrum band while respecting specific spectrum rights. Just like any other resource with mismatched demand and supply, all steps towards better utilization have also increased the need for large scale spectrum monitoring.

Our key question is: how effectively can we monitor available spectrum opportunities across time and space? One key technology in this space is spectrum databases that store spectrum availability information. We augment existing spectrum database technologies to improve their accuracies in a cost-effective fashion [CoNEXT '14]. The idea is to supplement existing model-based techniques with actual spectrum measurements. We also contribute towards making the spectrum measurements themselves scalable by developing techniques to perform spectrum sensing on mobile devices [MOBICOM '13, Hot Wireless '16]. These efforts culminate building a spectrum database system called SpecSense that can schedule and collect measurements from a distributed system of spectrum sensors in order to estimate spatio-temporal patterns in spectrum availability [INFOCOM '17, INFOCOM '18, TCCN '19].

■ **Network-Side Localization of Cellular Devices:** We address the problem of network-side localization where cellular operators are interested in localizing cellular devices by means of signal strength measurements alone [INFOCOM '15]. While fingerprinting-based approaches have been used recently to address this problem, they require significant amount of geo-tagged ('labeled') measurement data that is expensive for the operator to collect. Our goal is to use semi-supervised and unsupervised machine learning techniques to reduce or eliminate this effort without compromising the accuracy of localization. Our experimental results in a university campus (6 sq. km) demonstrate that sub-100 m median localization accuracy is achievable with very little or no labeled data so long as enough training is possible with 'unlabeled' measurements. This provides an opportunity for the operator to improve the model with time.

Industrial Research

After finishing my PhD, I joined NEC Laboratories America (2017) as a researcher in the mobile communications/networking group. The opportunity to work in an industrial lab helped me approach problems with a fresh perspective. Second, being in a group that is prominent in the academic research community helped me further hone my problem solving skills. At NEC Labs, my primary responsibility is to explore solutions and applications related to wireless sensing technologies. In particular, I am interested in using autonomous mobile platforms (e.g., robots, drones) as a sensing infrastructure. Public safety, for instance, fire fighting, disaster relief, emergency response applications seemed like a natural usecase for such sensing infrastructures. I developed two major solutions for such scenarios. The first one is to establish a cellular communication (LTE) infrastructure using a drone, that very much emulates a 'flying base station'. The second solution is to use the drone to find locations of such first responders when they enter GPS-denied environments. The spatial ubiquity achieved by a drone coupled with the versatility of RF signals as a sensing modality has made it possible to culminate the two into various autonomous mobile sensing applications. My tenure at NEC Labs led to publishing two papers in reputed venues [CoNEXT '18, NSDI '19] and a third one in submission.

■ **Drone-based Networking:** SkyLiTE (LTE from the Sky), is the first ever platform built to provide *on-demand* cellular connectivity using *untethered* drones. Untethered drones make the solution practical in terms of the drone's flexibility to move to locations that are otherwise inaccessible by the ground. However, this poses a host of technical challenges. For instance, such communication stacks are not designed to be deployed on 'mobile' base stations due to rapidly fluctuating link quality. I designed SkyRAN [CoNEXT '18], the Radio Access Network (RAN) part of the SkyLiTE system. SkyRAN places the drone mounted LTE *eNodeB* in the 3D aerial space that optimizes overall connectivity to the users (UEs) on the ground. The solution involves estimating the users' locations as well as their throughput, given the eNodeB's aerial position. The solution is adaptive to the mobility of the users and the drone continuously repositions itself to maintain optimal connectivity. SkyRAN's design is applicable beyond LTE and canbe used to provide Wi-Fi or millimeter wave (60 GHz) connectivity as well. SkyRAN has a big impact when it comes to deploying communication infrastructure in case of natural disasters or crisis situations.

■ **Infrastructure-free Indoor Localization:** I designed TrackIO [NSDI '19] that includes a suite of technologies providing real-time localization in GPS-denied environments without reliance on any pre-deployed localization infrastructure or exhaustive calibration of the given environment. The technology is based on a precise RF ranging modality called the Ultra Wideband (UWB). UWB provides a good trade-off when it comes to indoor penetrability (30 – 40 meters) versus ranging precision ($\approx$ 10 centimeters). The solution uses a single drone (with GPS access) equipped with a UWB radio that moves around the perimeter of a building and localizes mobile

UWB nodes inside. The drone continuously estimates the distances of the nodes from itself based on UWB's precise *time of flight* (ToF) information, while it's trajectory forms what is known as a *Synthetic Aperture Radar* (SAR) to localize the nodes in an absolute sense. `TrackIO` started as a summer intern project and soon became big as I continued to make the system scale with more mobile nodes and added location based features (e.g., indoor geo-tagged videos). Algorithms are designed to localize the set of mobile nodes together with as few ranging measurements as possible. A 'drone-free' version of `TrackIO` is also designed and implemented. The practicality and versatility of the `TrackIO` solution created a huge traction from within and outside NEC as mentioned in the following section.

Broader Impact

My tenure both in the academic setting, as a doctoral candidate, and in the industry, as a researcher, has produced impactful work. The impact goes beyond publishing papers, writing successful grant applications or granted patents. Based on my work, we got opportunity to talk to the government and policy-makers with initial promises for trialing out our solution.

► **Funding, Proposal Writing**: My work on spectrum databases and mobile spectrum sensing systems has produced two successfully funded proposals to the National Science Foundation (USD 344,946 and USD 800,000). This also resulted in successful collaboration with the industry (Dr. Milind Buddhikot, Bell Laboratories, NJ) and academia (Prof. Sumit Roy, University of Washington, Seattle). Given the reach of cellular Internet, this research will positively impact a broad spectrum of the population. Research on leveraging unused spectrum can potentially improve cellular bandwidth and improve connectivity; this will reduce the cost of expensive cellular data plans making mobile Internet more affordable.

► **Technology Transfer, Patents**: The `ExBox` system that I built in collaboration with HP Labs resulted in two granted patents. With applications like $4K$ video, Augmented Reality (AR)/Virtual Reality (VR) getting more mainstream, my work on modeling quality of experience (QoE) becomes more relevant. The work was taken up by Aruba Networks (acquired by Hewlett Packard Enterprise) to be integrated with their admission control solutions for enterprise wireless LANs. My work on network-side localization has been transferred to Huawei Technologies.

► **Technology Incubation, Startup**: `TrackIO` is one of the most successful solutions that addresses a pressing pain point of our society: user location in unknown environments. We received appreciation from different verticals including fire counties, police departments and construction companies. `TrackIO` technology was selected as a finalist in *Under Fire 2019*, a highly attended firefighting technology innovation showcase. We have already received *Letters of Intent* (LoI) from atleast four Fire Counties who see potential in our technology and are interested to be a part of the pilot run (which, I believe, will lead to early stage investments). The technology is equally appreciated by the business units in NEC Japan. In the United States, TrackIO is currently in the process of being spun-off from NEC Laboratories and is in the incubation stage under NEC-X, a start-up accelerator program run by NEC in the Silicon Valley, California.

Future Research Directions

Wireless networking is a rapidly progressing field of research that goes beyond the realm of communications. Wireless systems are being envisioned as sensing modalities that transcends traditional interfaces for communication and gives it a brand new dimension to interact with and interpret the physical world. Take for instance, millimeter waves (mm-wave). While mm-wave technology promises a new horizon for `5G` networks with bandwidths never thought of before, it also opens up opportunities for performing complicated real-word sensing tasks, like intricate gesture recognition, soil and crop monitoring and health diagnostics. My long term research goal is to bridge the gap between communication and sensing, design holistic solutions that touch our everyday lives. In the following, I outline an agenda for my future research directions.

■ Wireless Networking

I will continue to explore technologies that will help improve the current state of wireless networks in terms of capacity, coverage and deployability to meet future traffic demands. Additionally, my experience in setting up wireless testbeds, running experiments on controlled as well as on production networks, collecting network measurements give me a solid foundation to carry forward and closely manage experimental/testbed-related research activities.

► **Networking in Challenged Environments:** Recent years have seen tremendous growth in natural disaster scenarios in the country - be it terrible floods (Kerala'18,'19), landslides (Uttarakhand'13), earthquakes or forest fires (Bandipur'19). Such incidents play havoc on the infrastructure of the country not to mention circumstances when telecommunication networks are completely taken down. My expertise on drone technology along with relevant communications modality can help create networking solutions for such challenged environments. Many lives are lost due to inadequate tracking of people or their inability to communicate with the emergency responders. With a platform similar to `SkyLiTE` as the foundation, I want to develop diverse systems that work in such challenged environments lending a hand of hope to my fellow citizens. Such technology will not only help common people communicate during crisis situations but also assist search and rescue teams to track humans, work efficiently with respect to food and first aid dispatch.

► **Understanding `5G` Deployments:** `5G` is an umbrella term meant to bring together a host of technologies (mm-waves, massive MIMO, multi-user MIMO, spectrum sharing), possibilities in which such technologies are combined together and deployment strategies. I want to contribute towards the ongoing `5G` research in India, particularly in terms of deployment model. Given the short range, directionality and highly lossy nature of mm-waves, placing such base stations would bring a new set of challenges. How does such links perform under mobile user scenarios? Are such links more suitable for high bandwidth wireless backhauls? What should be the strategy for using such mm-wave base stations in conjunction to existing 4G/LTE infrastructure? Second, experiments need to be carried out in mm-wave (e.g., 28 GHz, 39 GHz) as well as sub-6 GHz bands based on application scenarios. Third, my experience with spectrum sensing will help me understand spectrum opportunities in India (e.g., in DTV spectrum) to use unused or unclaimed spectrum whitespaces.

► **Data Driven Modeling of Network Performance:** As networks grow more complex, modeling, understanding and debugging network performance issues become harder. For instance, there is now little understanding whether a poor video streaming experience on a smartphone should be attributed to capacity issues in the last mile wireless link or in the cellular network core, Internet congestion, CDN provider or a combination of multiple such factors. Disentangling such dependencies is not a straightforward problem in a realistic setting as majority of the network components lack sufficient visibility to the end user application so that traditional modeling approaches can be applied. I will exploit crowdsourcing techniques and data driven modeling methodologies to address such modeling problems. The idea is to collect scenario-specific data for application performance via crowd-sourcing and then use such large-scale data to draw inferences about performance dependencies. Such approaches will be critical in building mobile applications and provisioning networks in the future.

■ Wireless Sensing

Over the next decade, commercial devices like smartphones or IoT systems will be augmented with new radio interfaces for communication that will double as new modalities for wireless sensing. While Google is introducing mm-wave technology (60 GHz) in its latest *Pixel*-brand of smartphones, Apple integrated an ultra wideband (UWB) radio in their latest iPhone-11 release. Specifically, I want to embrace such pervasiveness of RF based sensing coupled with advances in signal processing, machine learning and neural networks that will lead to the invention of fundamentally new primitives in the sensing space spanning an array of application verticals.

► **Towards a holistic 'Sensing+Communication' Interface:** At a fundamental level, today's technologies decouple sensing and communication in the sense that applications are just exposed to data bits from the communication interface. However, sensing brings a whole new context to such applications that has the potential to make them even more richer and proactive. This requires us to rethink such interfaces and understand their interaction. I want to contribute towards bridging such gap between communication and sensing. Take for example, a wireless car key-fob, used to remotely unlock/start cars. Typically, such key-fobs use *bluetooth* technology to communicate with the controller present in the car. This communication has been proved to be insecure leading to a recent rise of car thefts. Imagine that in the near future car key-fobs communicate using an UWB interface that doubles as a sensor to accurately estimate the distance of the car (range) from the key. The UWB communication protocol coupled with such ranging will make it necessary for the driver to be physically present in the vicinity of the car in order to unlock it, improving the security manyfold.

► **Indoor Localization:** Improvement in RF based localization is in tandem with advances in wireless sensing technologies. At any rate, indoor localization still remains largely an unsolved problem in majority of real world scenarios. With a considerable amount of experience in this domain, I want to innovate and improve further on this technology in terms of algorithms, systems, cost and deployability. My goal is to achieve a robust localization primitive with as less pre-deployed infrastructure as possible, at an affordable cost and without the necessity of any intricate calibration effort. Building over such a primitive and integrating it to appropriate Human-Computer Interaction (HCI) interfaces, I want to employ such technology in different pressing problems in our society. For instance, integrating it to a *smart wheelchair* can assist disabled/blind people navigate indoors or when integrated with critical medical equipments (e.g., life-support systems, portable X-Rays, CT-scanners or ECG machines) inside a hospital can result in better accountability and availability. Overall, fine grained localization and tracking is still a 'holy grail' for a lot of such applications that has tremendous social traction.

► **Large Scale Monitoring Infrastructure:** Future applications (e.g., smart cities) will depend on huge amounts of spatio-temporal measurement data of different modalities collected from a diverse array of sensors/IoT devices. We need better understanding about how to deploy such sensors in a geo-spatial sense or how to schedule them across space-time for performing specific sensing tasks. The latter is important as it may be impractical for a sensor to continuously make measurements due to energy or bandwidth limitations. Second, such sensor data may contribute to real time decision making in certain applications. Hence optimizing end-to-end latency is a vital requirement in such a monitoring infrastructure. Such platforms can be readily adapted to a number of applications, particularly in a smart city context. Municipality garbage-bins can be integrated with radio frequency based identification (RFID) tags that makes waste collection more efficient. Other important applications include air quality or drinking water monitoring, understanding car traffic congestion bottlenecks in cities and so on.

My vision is to create the necessary toolbox combining algorithms, sensor systems, network protocols and big data techniques that will be an essential component of future mobile networks – 5G and beyond – supporting upcoming applications such as IoT or smart cities. To deliver on this vision, I will go beyond narrowly focused academic setting and work with industry to understand realistic applications and constraints.

References

[**NSDI '19**] Ashutosh Dhekne, **Ayon Chakraborty**, Karthik Sundaresan, Sampath Rangarajan. *TrackIO: Tracking First Responders Inside-Out*, in USENIX NSDI 2019.

[**TCCN '19**] Arani Bhattacharya, **Ayon Chakraborty**, Samir Das, Himanshu Gupta, and Petar M. Djuric. *Spectrum Patrolling with Crowdsourced Spectrum Sensors*, in IEEE Transactions on Cognitive Communications and Networking (2019).

[**CoNEXT '18**] **Ayon Chakraborty**, Eugene Chai, Karthik Sundaresan, Amir Khojastepour, Sampath Rangarajan. *SkyRAN: A Self-Organizing LTE RAN in the Sky*, in ACM SIGCOMM CoNEXT 2018.

[**INFOCOM '18**] **Ayon Chakraborty**, Arani Bhattacharya, Snigdha Kamal, Samir Das, Himanshu Gupta and Petar Djuric. *Spectrum Patrolling with Crowdsourced Spectrum Sensors*, in IEEE INFOCOM 2018.

[**INFOCOM '17**] **Ayon Chakraborty**, Shaifur Rahman, Himanshu Gupta and Samir Das. *SpecSense: Crowdsensing for Efficient Querying of Spectrum Occupancy*, in IEEE INFOCOM 2017.

[**CoNEXT '16**] **Ayon Chakraborty**, Shruti Sanadhya, Samir Das, Dongho Kim and Kyu-Han Kim. *ExBox: Experience Management Middlebox for Wireless Networks*, in ACM SIGCOMM CoNEXT 2016.

[**Hot Wireless '16**] **Ayon Chakraborty**, Udit Gupta and Samir Das. *Benchmarking Resource Usage for Spectrum Sensing on Commodity Mobile Devices*, in ACM HotWireless 2016.

[**INFOCOM '15**] **Ayon Chakraborty**, Luis Ortiz and Samir Das. *Network-side Positioning of Cellular-band Devices with Minimal Effort*, in IEEE INFOCOM 2015.

[**CoNEXT '14**] **Ayon Chakraborty** and Samir Das. *Measurement-Augmented Spectrum Databases for White Space Spectrum*, in ACM SIGCOMM CoNEXT 2014.

[**IMC '14**] Fatima Zarinni, **Ayon Chakraborty**, Vyas Sekar, Samir Das and Phillipa Gill. *A First Look at Performance in Mobile Virtual Network Operators*, in ACM SIGCOMM IMC 2014. **Best paper award nominee.**

[**MOBICOM '13**] **Ayon Chakraborty**, Samir Das and Milind Buddhikot. *Radio Environment Mapping with Mobile Devices in the TV White Space*, ACM MOBICOM 2013 (Extended Abstract). **Finalist in ACM Student Research Competition.**

[**CellNet '13**] **Ayon Chakraborty** and Samir Das. *Adapp: An Adaptive Network Selection Framework for Smartphone Applications*, in ACM CellNet 2013 (co-held with ACM MobiSys'13).