Saurabh Kalikar

CONTACT INFORMATION

BSB331, PACE Lab
IIT Madras, Chennai 600 036

+91-9043756015
kalikar.saurabh@gmail.com

RESEARCH INTERESTS

Parallel and Distributed Computing, Program Analysis

EDUCATION

MS + PhD, Computer Science and Engineering Jan 2014 to Present

IIT Madras, India.

- Thesis Topic: *Optimizing Hierarchical Locks for Scalable Synchronization*
- Advisor: Dr. Rupesh Nasre

BTech, Computer Science and Engineering 2008 to 2012

Government College of Engineering, Amravati, Maharashtra, India.

PUBLICATIONS

1. Saurabh Kalikar, Rupesh Nasre. "DomLock: A New Multi-Granularity Locking Technique for Hierarchies". In *Proceedings of the* 21^{st} *ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming,* PPoPP 2016 held at Barcelona, Spain.
   - **Distinguished paper award at PPoPP'16.**
   - Artifacts successfully evaluated. (http://pace.cse.iitm.ac.in/tools.php).
2. Saurabh Kalikar, Rupesh Nasre. "DomLock: A New Multi-Granularity Locking Technique for Hierarchies". In *ACM Transactions on Parallel Computing (TOPC).*
   - **Invited article**.
3. Ganesh K, Saurabh Kalikar, Rupesh Nasre. "Multi-granularity Locking in Hierarchies with Synergistic Hierarchical and Fine-Grained Locks". In *Proceedings of the* 24^{th} *European Conference on Parallel and Distributed Computing,* EuroPar 2018 held at Turin, Italy.
4. Saurabh Kalikar, Rupesh Nasre. "NumLock: Towards Optimal Multi-Granularity Locking in Hierarchies". In *Proceedings of the* 47^{th} *International Conference on Parallel Processing,* ICPP 2018 held at Eugene, USA.
5. Saurabh Kalikar, Rupesh Nasre. "Toggle: Contention-Aware Task Scheduler for Concurrent Hierarchical Operations". In *Proceedings of the* 25^{th} *European Conference on Parallel and Distributed Computing,* EuroPar 2019, held at Gottingen Germany.

PROFESSIONAL EXPERIENCE

- Research Intern July 2019 to Present
  - **Microsoft Research, Bangalore**.
  - Working on Apache-Spark compiler to optimize big-data jobs. Spark's *catalyst* compiler generates the execution plans with lots of data shuffling across stages. We are enhancing Spark compiler with a richer type analysis and introducing new logical and physical operators for generating RDD code with lesser data shuffles.
- Research Intern Jul 2018 to Oct 2018
  - **Microsoft Research, Bangalore**.
  - We designed a source-to-source optimizing compiler for Distributed Big-Data jobs. For SQL-based big-data jobs, our compiler performs a whole query program analysis and identifies the scope for the optimization. The execution plans generated by our compiler are on an average 1.6× faster than the state-of-the-art big-data compilers.
- Software Engineer Mar 2013 to Dec 2013
  - Cognizant Technology Solutions India Private Ltd., Pune.
  - Developed an ASP.NET application for a pharmaceutical project and worked on the integration of front end and back end, with MS-SQL Server database.

Honors and Awards

- Institute Research Award, IIT Madras, Sept 2019.
- Distinguished paper award at PPoPP 2016.
- Won the HiPC'16 Student Parallel Programming Challenge-Intel Track (Team of 3).
- Invited to attend Google's 4^{th} PhD Student Summit on Compiler and Programming Technology, 5-7 December, 2016, Munich, Germany.
- Travel grants
  - ACM-India student travel grant for attending EuroPar 2019 (INR 60,000)
  - Kris Gopalakrishnan Endowment Student Travel Grant for attending EuroPar 2018 (INR 1,50,000).
  - ACM SIGPLAN Professional Activities Committee (PAC) travel grant for attending PPoPP 2016 (USD 1000).
  - ACM Programming Languages Mentoring Workshop (PLMW) scholarship for attending POPL 2015 (USD 550).

Graduate Courses

- Concurrent Programming
- Program Analysis
- Parallel Computer Architecture
- High-Performance Parallel Computing
- Advanced Data Structures and Algorithms
- Mathematical Concepts for Computer Science
- Digital Design Verification
- Indexing and Searching in Large Datasets
- Computer Architecture

Software Skills

- Computer Programming: C, C++, Java, C#, Scala
- Parallel Programming: Pthreads, CUDA, OpenMP
- Compiler Framework: LLVM

Projects

Toggle: Contention-Aware Task Scheduler for Concurrent Hierarchical Operations (PhD research)

- We propose *Toggle*, a task scheduler which communicates with lock manager for task scheduling.
- Our scheduler predicts possible lock conflicts among the tasks and schedules them to a single thread thus minimizing the lock contention overhead.
- We design hash-based concurrent data-structures for task management and a thread communication protocol for the scheduler.
- We test the performance of Toggle with well-known STMBench7 benchmark and obtain on an average 22% performance gain with our interval-based hierarchical locking techniques.

NumLock: Towards Optimal Multi-Granularity Locking in Hierarchies (PhD research)

- In multi-granularity locking, there exist multiple ways to lock a set of requested entities.
- Each of these locking alternatives has different locking cost and concurrency cost.
- We propose a novel technique, *NumLock*, which provides a model to choose the best locking option for a set of requests at runtime with negligible overhead.
- We evaluate NumLock with STMBench7 benchmark, and compare against intention locks and DomLock and observe that NumLock outperforms the existing locking approaches.

DomLock: A new multi-granularity locking technique for hierarchies (PhD research)

- We propose a new multi-granularity locking technique for hierarchical data structures.
- Our technique assigns logical intervals to the nodes in a hierarchy which helps us in efficiently detecting the hierarchical overlaps between any two lock requests.
- DomLock reduces the locking cost of parallel operations in multi-threaded environments by avoiding unnecessary path traversals and by acquiring lock only on the dominator node.
- We implemented DomLock in STMBench7 benchmark suite and obtained on an average 42% performance improvement over the existing locking in STMBench7.

Multi-granularity Locking in Hierarchies with Synergistic Hierarchical and Fine-Grained Locks
(Mentored MTech Project)

- We propose HiFi, a novel locking protocol which allows synergistic co-existence of fine-grained and hierarchical locks.
- We propose a novel indexing technique for hierarchies which uniquely identifies every node as an interval value and effectively captures hierarchical dependencies between nodes.
- Using real-world XML hierarchies and synthetic datasets, we illustrate that HiFi considerably improves the parallel performance of the underlying application.
- I mentored this project and the project received **the best MTech project award** at IIT Madras. (Ganesh K, CSE, MTech, 2018)

Partitioning of 3-dimensional data-set in parallel Sep to Nov 2016

- This work won **first rank** in HiPC'16 Student Parallel Programming Challenge (team of 3).
- We designed a technique which partitions the data into connected components in parallel.
- The key idea in this technique is a concurrent lock free data structure to represent a graph.
- We implemented our techniques for NVIDIA GPUs using CUDA and Intel Xeon Phi coprocessor using OpenMP.
- Using our optimizations, we partitioned the data-set of 8 million points in just 7 seconds on NVIDIA GPUs and 12 seconds on Intel Xeon Phi processor.

COURSE PROJECTS

Indexing large graph database to speed up the sub-graph mining (Team of 2) Jul to Nov 2016

- For a given graph database and a query graph, we need to find out all graphs from the database which are subgraphs of the given query graph.
- As the sub-graph isomorphism test is an NP-Complete problem, checking every graph in the database with query graph is a costly operation in terms of computation time.
- Our technique prunes out the graphs which cannot be subgraphs of the query graph.
- We used a technique for frequent subgraph mining to extract features from database and we designed the feature selection technique to select the best possible feature set.
- In our experiments on real-world molecular datasets, our selected features pruned out on an average 96% of the graphs without actually testing the subgraph isomorphism with query graph.
- Our technique performed the best in the course project contest for minimizing querying time for 200 queries over a database containing 70K graphs.

POSITIONS OF RESPONSIBILITY

- Member of Artifact Evaluation Committee for the 23rd ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming (PPoPP'18).
- Member of Artifact Evaluation Committee for the 38th ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI'17).
- Teaching Assistant: Program Analysis, Computer Architecture, Concurrent Programming, Computer Programming Lab, National Workshop on Programming with Intel Xeon Phi Co-processor.

TALKS

- Presented the paper titled *Toggle: Contention-Aware Task Scheduler for Concurrent Hierarchical Operations* in EuroPar 2019, Gottingen Germany, Aug 26 - Aug 30, 2019.
- Presented the paper titled *NumLock: Towards Optimal Multi-Granularity Locking in Hierarchies* in ICPP 2018, Eugene, USA, Aug 12 - Aug 14, 2018.
- Presented the paper titled *Multi-granularity Locking in Hierarchies with Synergistic Hierarchical and Fine-Grained Locks* in EuroPar 2018, Turin, Italy, Aug 27 - Aug 31, 2018.
- Presented a poster in 23^{nd} IEEE International Conference on High Performance Computing, Dec 20 - Dec 22, 2016, in Hyderabad, India.
- Presented a poster in Google PhD Student Summit on Compiler & Programming Technology, 5-7 December, 2016, Google (Munich), Germany, Dec 4 - Dec 8, 2016.
- Presented the paper titled *DomLock: A New Multi-Granularity Locking Technique for Hierarchies* in PPoPP'16, Barcelona, Spain, March 12 - March 16, 2016.
- Presented the poster titled *Locking for Hierarchical Data Structures* in POPL'15, TIFR, Mumbai, India, January 12-18, 2015.