Saurabh Kalikar

Contact Information

624, Tamiraparni hostel
IIT Madras, Chennai 600 036

+91-9043756015
kalikar.saurabh@gmail.com

Research Interests

Parallel computing, Thread synchronization

Education

Ph.D., Computer Science and Engineering Jan 2014 to Present

IIT Madras, India.

- Thesis Topic: *Lock synchronization in parallel programs*
- Advisor: Rupesh Nasre

B.Tech., Computer Science and Engineering 2008 to 2012

Government College of Engineering, Amravati, Maharashtra, India.

Publications

Saurabh Kalikar, Rupesh Nasre. "DomLock: A New Multi-Granularity Locking Technique for Hierarchies". In *Proceedings of the 21st ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming,* PPoPP 2016 held at Barcelona, Spain.
Acceptance rate: 19.2%.

- **Distinguished paper award at PPoPP'16.**
- Artifacts successfully evaluated. (http://pace.cse.iitm.ac.in/tools.php)
- The extended version of DomLock is invited as an article in *ACM Transactions on Parallel Computing.*

Professional Experience

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

- Won the HiPC'16 Student Parallel Programming Challenge-Intel Track (Team of 3).
- Invited to attend Google's 4th PhD Student Summit on Compiler and Programming Technology, 5-7 December, 2016, Munich, Germany.
- Distinguished paper award at PPoPP 2016.
- Travel grants
  - 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).
- 4th rank in HiPC 2015 Student Parallel Programming Challenge-Intel Track.
- First rank in coding contest in national level technical festival, Amravati (Feb 2011).

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
- Parallel Programming: Pthreads, CUDA, OpenMP
- Compiler Framework: LLVM

Projects

DomLock: A new multi-granularity locking technique for hierarchies (Ph.D. research) On going

- Proposed a new multi-granularity locking technique for hierarchical data structures.
- A novel technique of assigning logical intervals to the nodes in a hierarchy is presented.
- DomLock reduces the locking cost of parallel operations in multi-threaded environments by acquiring lock only on dominator node.
- We implemented DomLock in well-known STMBench7 benchmark suite and obtained on an average 42% performance improvement over the existing lockings in STMBench7.
- I presented this work at PPoPP 2016, Barcelona, Spain in March 2016.

Partitioning of 3-dimensional data-set in parallel — Sep 2015 to Nov 2015

- This work is done as part of HiPC'16 Student Parallel Programming Challenge (team of 3) and got **first rank**.
- We designed a technique which partitions the data into connected components using OpenMP which can run on Intel Xeon Phi (KNL) coprocessor.
- The key idea in this technique is to design a concurrent lock free data structure to represent a graph.
- Using our parallel technique executed on 240 cores, we partitioned the data-set of 8 million points in just 12 seconds.

Course Projects

Indexing large graph database to speed up the sub-graph mining queries (Indexing and Searching in Large Datasets-Team of 2) — Jun 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 NPC problem, checking every graph with query graph is a costly operation.
- Our technique prunes out the graphs which cannot be subgraphs of the query graph.
- We use a technique for frequent sub-graph mining to extract features from database and we designed the feature selection technique to select best possible feature set.
- In our experiments on real datasets, our selected features pruned out on an average 96% of the graphs without actually testing the subgraph isomorphism with query graph.
- Our technique **won** the course project contest for minimizing querying time for 200 queries over 70K graphs database.

Parallel K-Means clustering for large data sets (Concurrent Programming-Team of 3) Jan to May 2014

- Implemented a parallel algorithm for clustering the data points and assigning the cluster IDs to every point in the input data set.
- Intel Cilk framework was used for parallelization.
- Parallel K-Means performed 3× faster than existing K-means++ clustering algorithm.

Parallel *grep* (Program Analysis-Team of 2) — Jan to May 2014

- Implemented a parallel and scalable version of Linux *grep* command.
- Pthread library was used to create and manage multiple parallel threads.
- On 16 core machine, we obtained 10× speedup over sequential execution.

Positions of Responsibility

Teaching Assistant

- First year B.Tech. Introduction to Programming lab. — Jul to Dec 2016
- National Workshop on Programming with Intel Xeon Phi Co-processor, held at IIT Madras. — Jul 17-18, 2015
- Second year B.Tech. Computer Programming lab. — Jul to Dec 2015
- Program Analysis, a graduate level course. — Jan to May 2014