

MILIND M. CHABBI

Curriculum Vitae

Programming Systems Group, Uber Technologies
Palo Alto, CA

Email: Chabbi.Milind@gmail.com

Phone: (520) 247-9385

OBJECTIVE

Advance the horizons of human understanding by inventions and discoveries. Deliver useful automation tools for the community.

DOMAINS OF EXPERTISE

Compilers, Static and Dynamic Program Analysis, High Performance Parallel Computing (HPC), Performance Analysis Techniques (Measurement, Analysis, and Visualization), Performance Optimization, Memory-Hierarchy Optimizations, Vectorization, Synchronization Algorithms, Multicore Accelerated Architectures, Architecture-aware Algorithms, Generative AI for developer tools, and Optimize Deep-Learning Workloads.

SOFT SKILLS

Technical leadership and vision, Incubation, Initiating and fructifying research agendas, Mentoring, Multitasking, Cross-team collaboration and alignment, Technical writing, Presentation, and University and Inter-company engagements.

CURRENT POSITIONS

Senior Staff Research Scientist and Tech Lead, Uber Programming Systems Group.

EXPERIENCE

- **Uber Technologies** (Sunnyvale) Sr. Staff Research Scientist [June/2018 -]

Tech lead for a team of ten Ph.D.s working on program analysis and GenAI for efficiency, reliability and developer velocity.

Data center efficiency via code optimization

- Designed and architected data-center-wide code-level performance profiling for microservices.
- Championed and incorporated profile-guided compiler optimizations into Uber service fleet.
- Envisioned and delivered developer tools to analyze long tail latencies in microservice RPCs.
- Guided datacenter hardware selection by characterizing how enterprise micro-services exercise different data-center hardware.
- Initiated foundational efforts in a data-driven code optimization.
- Reduced data center tax by improving throughput (efficiency) and reducing resource (e.g., memory footprint) consumption of services.

Impact: reduced data center total cost of ownership by >15%; enabled hundreds of engineers to swiftly reduce latencies in hundreds of critical endpoints.

Language and runtime

- Envisioned, architected and up-leveled the technical trajectory for large-scale concurrency bug detection and repair for Golang programs.
Found over 7000 data races, repaired over 5500 data races in 80-million lines of Go codebase hosting 3500 distinct microservices driving over 900 unique developers.
Impact: 150 developer years and 50 Million USD saved.
- Envisioned and architected hardware performance counter-based profiler for Golang.
Impact: improved profiling accuracy by over 2× compared with the default Go profiler.

Mobile

- Developed link-time compiler optimizations for reducing Uber app size.
Impact: reduced Uber iOS rider app size by 23%; improved first-time customer bookings by 12%.
- Developed schedulers for mapping streaming applications on to memory-constrained hardware.
Impact: reduced Uber android driver app out-of-memory crashes by 50%.

- **Baidu Research** (Sunnyvale) Senior Research Scientist [Oct/2017 - June/2018]
 - Developed a communication optimization algorithm to scale Natural-Language Processing (NLP) applications.
Impact: 7× speedup in training language models, 35% improvement in prediction accuracy, and 2.5× larger training set enablement.
 - Developed an approximation technique to reduce computational demands of NLP and image-processing deep-learning workloads.
Impact: 10% training time reduction with no accuracy loss on common image and speech recognition models.
- **Hewlett Packard Labs** (Palo Alto) Senior Research Scientist [Aug/2015 - Sep/2017]

System software, hardware-software co-design, performance analysis techniques—measurement, modeling, and simulation—for many-core NUMA systems with peta-byte non-volatile memory and fast interconnection fabric, and application optimization under HPE’s **memory-driven computing** initiative and the US DoE **exascale** project. Mentored eight interns.
Impact: 100× speedup of a bioinformatics pipeline on the new technology that led to 10s of millions of USD HPE server purchases by a customer and garnered handful additional customers.
- **Lawrence Livermore National Lab** Computational Intern [Jun/2013 - Sep/2013]

Performance tuning and feature enhancements of *ROSE* source-to-source compiler.
Impact: speedup compilation time by 9×.
- **Wave Equations (Houston)** HPC Expert [Jan/2012 - Aug/2015]

Championed, prototyped, and demonstrated the effectiveness of GPU computing for solving finite difference wave solvers for the Oil and Gas Industry. Extensively utilized loop transformation, vectorization, and memory-hierarchy optimizations to achieve dramatic speedups.
Impact: 40 × speedup of reverse-time migration seismic MPI application.
- **Microsoft** (Redmond) Software Engineer 2 [Aug/2007 - Jul/2010]

Performance enhancement and versioning resilience of JIT compiler for the Common Language Runtime (CLR).
Impact: reduced a critical startup time of CLR by 10×.
- **Microsoft** (Redmond) Intern [May/2006 - Aug/2006]

Design and development of *Python 2.5* language features in the Python compiler (IronPython).
Impact: improved feature parity with CPython standard implementation.
- **Motorola** (Bangalore) Software Engineer [Jul/2002 - Jul/2005]
 - Development and testing of internet browser shipped on various Motorola phones.
 - Porting HTTP and WAP stacks from legacy P2K to Symbian Operating System.
 - Direct interaction with wireless service providers such as AT&T, T-Mobile, Vodaphone, and Verizon in various world-wide geographies (North America, South America, and Europe) to ensure product requirements, quality, and customer satisfaction.

EDUCATION

- Ph.D. in Computer Science, Rice University, Houston, TX. [Aug/2010 - Jul/2015]
G.P.A. 4.01/4.0.
Thesis title: Software Support for Efficient Use of Modern Computer Architectures.
Advisor: Prof. John Mellor-Crummey.
- Master of Science in Computer Science, The University of Arizona, Tucson, AZ. [Aug/2005 - May/2007]
G.P.A. 4.0/4.0.
Thesis title: Efficient Dynamic Taint Analysis of Programs Using Multicore Machines.
Advisors: Prof. Saumya Debray and Prof. Greg Andrews.
- Bachelor of Engineering in Computer Science, Sri Jayachamarajendra College of Engineering (SJCE), Mysore, India. [Oct/1998 - May/2002]
G.P.A. 3.8/4.0.

PROGRAMMING SKILLS

Languages: C, C++, Python, Java, C#, Go, CUDA, Fortran, and assembly (x86, PowerPC, and ARM).

Tools/Frameworks: OpenMP (Open Multi-Processing), Intel TBB (Thread Building Blocks), Cilk Plus, MPI (Message Passing Interface), LLVM, Tensorflow, SPIN model checking.

PUBLICATIONS

1. Feiyang Jin, Zhizhou Zhang, Rajkishore Barik, Gautam Korlam, and **Milind Chabbi**. *GenAI-based Datarace Fix for Real-World Golang Programs*. In the proceedings of the Conference on Machine Learning for Systems (MLSys '23).
2. Angelina Lee, and **Milind Chabbi**, *The Tail of Errors in Microservices* (Under review).
3. Georgian-Vlad Saioc, Dmitriy Shirchenko, and **Milind Chabbi**. *Unveiling and Vanquishing Goroutine Leaks in Enterprise Microservices with Dynamic Analysis*. In the proceedings of the international symposium on Code Generation and Optimization (CGO '24). Edinburg, UK. March, 2024.
4. Qidong Zhao, **Milind Chabbi**, and Xu Liu, *EasyView: Bringing Performance Profiles into Integrated Development Environments*. In the proceedings of the international symposium on Code Generation and Optimization (CGO '24). Edinburg, UK. March, 2024.
5. Vivek Shahare, **Milind Chabbi**, Nikhil Hegde. *Protecting Locks Against Unbalanced Unlock*. In the proceedings of the Symposium on Parallel Architectures and Algorithms (SPAA '23). Orlando, FL.
6. Muhammad Aditya Sasongko, **Milind Chabbi**, Paul Kelly, and Didem Unat. *Precise Event Sampling on AMD and Intel: In-Depth Quantitative and Qualitative Comparison*. In the Transactions on Parallel and Distributed Systems. May, 2023.
7. Muhammad Aditya Sasongko, **Milind Chabbi**, Paul Kelly, and Didem Unat. *Precise Event Sampling-based Data Locality Tools for AMD Multicore Architectures*. In the journal of Concurrency and Computation: Practice and Experience. April, 2023.
8. Bolun Li, Pengfei Su, **Milind Chabbi**, Shuyin Jiao, and Xu Liu. *DJXPerf: Identifying Memory Inefficiencies via Object-centric Profiling for Java*. In the proceedings of the international symposium on Code Generation and Optimization (CGO '23). Montreal, Canada. Feb, 2023.
9. **Milind Chabbi** and Murali Krishana Ramanathan. *A Study of Real-World Data Races in Golang*. In the proceedings of the 43rd ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI '22). San Diego, CA.
10. Chris Zhang, Murali Krishana Ramanathan, Prithvi Raj, Abhishek Parwal, Tim Sherwood, and **Milind Chabbi**. *Critical Path Analysis of Large-Scale Microservice Architecture*. In the proceedings of the USENIX Annual Technical Conference (ATC '22). July, 2022.
11. Bolun Li, Hao Xu, Pengfei Su, **Milind Chabbi**, Shuyin Jiao, and Xu Liu. *OJXPerf: Pinpointing Object Duplications in Java Programs*. In the proceedings of the 44th International Conference on Software Engineering (ICSE'22). Pittsburg, PA. May, 2022.
12. Muhammad Aditya Sasongko, **Milind Chabbi**, Mandana Bagheri Marzijarani, and Didem Unat. *ReuseTracker: Fast Yet Accurate Multicore Reuse Distance Analyzer*. In the ACM Transactions on Architecture and Code Optimization (TACO), Volume 19, Issue 1. March, 2022.
13. Muhammad Aditya Sasongko, **Milind Chabbi**, and Didem Unat. *Low-Overhead Reuse Distance Profiling Tool for Multicore*. In European Conference on Parallel Processing (Euro-Par 2021).
14. Chris Zhang, **Milind Chabbi**, Adam Welc, and Tim Sherwood. *Optimistic Concurrency Control for Real-world Go Programs*. In the proceedings of the USENIX Annual Technical Conference (ATC '21). July, 2021.
15. **Milind Chabbi**, Jin Lin, and Raj Barik. *An Experience with Code-Size Optimization for Production iOS Mobile Applications*. In the proceedings of the international symposium on Code Generation and Optimization (CGO '21). Virtual conference. Feb, 2021.
16. Qidong Zhao, Xu Liu, and **Milind Chabbi**. *DRCCTPROF: A Fine-grained Call Path Profiler for ARM-based Clusters*. In proceedings of the international conference for High Performance Computing, Networking, Storage and Analysis (SC'20). Atlanta, GA. **(Performance Track Best Paper, all tracks Best Paper nominee)**. Nov, 2020.

17. Jialiang Tan, Shuyin Jiao, **Milind Chabbi**, Xu Liu. *What Every Scientific Programmer Should Know About Compiler Optimizations?*. In proceedings of the International Conference on Supercomputing (ICS'20). Barcelona, Spain. June, 2020.
18. **Milind Chabbi**, Abdelhalim Amer, and Xu Liu. *Efficient Abortable-locking Protocol for Multi-level NUMA Systems: Design and Correctness*. In the journal of ACM Transactions on Parallel Computing (TOPC), Volume 7, Issue 3 July 2020.
19. Matthias Becker, Stefanie Warnat-Herresthal, Kathrin Klee, Jonas Schulte-Schrepping, Pawel Biernat, Patrick Günther, Kevin Babler, **Milind Chabbi**, Rocky Craig, Hartmut Schultze, Sharad Singhal, Thomas Ulas, and Joachim Schultze. *Accelerated Genomics Data Processing using Memory-Driven Computing*. [Under review.]
20. Raj Barik, Manu Sridharan, Murali Krishna Ramanathan, **Milind Chabbi**. *Optimization of Swift Protocols*. In proceedings of the international conference for Object-Oriented Programming, Systems, Languages, and Applications (OOPSLA'19). Athens, Greece. Oct, 2019.
21. Pengfei Su, **Milind Chabbi**, Xu Liu. *Pinpointing Performance Inefficiencies via Lightweight Variance Profiling*. In proceedings of the international conference for High Performance Computing, Networking, Storage and Analysis (SC'19). Denver, CO. Nov, 2019.
22. Muhammad Aditya Sasongko, **Milind Chabbi**, Pirah Noor Soomro, and Didem Unat. *ComDetective: A Fast and Accurate Communication Detection Tool for Threads*. In proceedings of the international conference for High Performance Computing, Networking, Storage and Analysis (SC'19). Denver, CO. **(Performance Track Best Paper, all tracks Best Paper nominee)**. Nov, 2019.
23. Matthias Becker, **Milind Chabbi**, Stefanie Warnat-Herresthal, Umesh Worlikar, Shobhit Agrawal, Jaydeep Bhat, Jonas Schulte-Schrepping, Kevin Babler, Patrick Günther, Hartmut Schultze, Thomas Ulas, Sharad Singhal, and Joachim L. Schultze. *Memory-driven Computing Accelerates Genomic Data Processing*. In proceedings of the 6th International Workshop on High Performance Computing on Bioinformatics (HPCB 2019) in conjunction with the IEEE International Conference on Bioinformatics and Biomedicine (BIBM 2019). San Diego, CA, USA.
24. Pengfei Su, Qingsen Wang, **Milind Chabbi**, Xu Liu. *Pinpointing Performance Inefficiencies in Java*. In proceedings of the ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering (ESEC/FSE '19).
25. Mostofa Patwary, **Milind Chabbi**, Heewoo Jun, Jiaji Huang, Gregory Diamos, and Kenneth Church. *Language Modeling at Scale*. In the proceedings of the International Parallel and Distributed Processing Symposium (IPDPS '19). Rio de Janeiro, Brazil.
26. Pengfei Su, Shasha Wen, Hailong Yang, **Milind Chabbi**, and Xu Liu. *Load Redundancy: A Software Inefficiency Indicator*. In the proceedings of the 41st ACM/IEEE International Conference on Software Engineering (ICSE'19). Montreal, Canada. **(Distinguished paper award)**.
27. Qingsen Wang, Xu Liu, and **Milind Chabbi**. *Featherlight Reuse-distance Measurement*. In proceedings of the 25th IEEE International Symposium on High-Performance Computer Architecture (HPCA '19). Washington D.C.
28. Qingsen Wang, Pengfei Su, **Milind Chabbi** and Xu Liu. *Lightweight Hardware Transactional Memory Profiling*. In proceedings of the ACM SIGPLAN symposium on Principles and Practice of Parallel Programming, 2019 (PPoPP '19). Washington D.C. **(Best paper award)**.
29. Abdelhalim Amer, Huiwei Lu, Pavan Balaji, **Milind Chabbi**, Yanjiw Wei, and Satoshi Matsuoka. *Lock Contention Management in Multithreaded MPI*. In ACM Transactions on Parallel Computing (TOPC).
30. **Milind Chabbi**, Shasha Wen, and Xu Liu. *Featherlight On-the-fly False Sharing Detection*. In proceedings of the ACM SIGPLAN symposium on Principles and Practice of Parallel Programming, 2018 (PPoPP '18). Vienna, Austria. **(Best paper award)**.
31. Du Shen, Xu Liu, **Milind Chabbi**. *An Evaluation of Vectorization and Cache Reuse Tradeoffs on Modern CPUs*. In proceedings of the 9th International Workshop on Programming Models and Applications for Multicores and Manycores, 2018 (PMAM '18). Vienna, Austria.

32. Shasha Wen, John Byrne, Xu Liu, and **Milind Chabbi**. *Watching for Software Inefficiencies with Witch*. In proceedings of the 23rd ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS '18). (**SIGPLAN Research Highlight**).
33. Adarsh Yoga and **Milind Chabbi**. *Path-synchronous Performance Monitoring in HPC Interconnection Networks with Source-code Attribution*. In proceedings of the 8th IEEE workshop in Performance Modeling, Benchmarking and Simulation of High Performance Computer Systems (PMBS '17) held as part of ACM/IEEE Supercomputing 2017 (SC '17), Denver, CO.
34. Shasha Wen, **Milind Chabbi**, and Xu Liu. *RedSpy: Exploring Value Locality in Software*. In proceedings of the 22nd ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS '17) (**ASPLOS Highlight Paper**).
35. **Milind Chabbi**, Halim Amer, Shasha Wen, and Xu Liu. *An Efficient Abortable-Locking Protocol for Multi-level NUMA Systems*. In proceedings of the ACM SIGPLAN symposium on Principles and Practice of Parallel Programming, 2017 (PPoPP '17). Austin, TX.
36. **Milind Chabbi** and John Mellor-Crummey. *Contention-Conscious, Locality-Preserving Locks*. In proceedings of the ACM SIGPLAN symposium on Principles and Practice of Parallel Programming, 2016 (PPoPP '16). Barcelona, Spain.
37. Tianzheng Wang, **Milind Chabbi**, and Hideaki Kimura. *Be My Guest—MCS Lock Now Welcomes Guests*. In proceedings of the ACM SIGPLAN symposium on Principles and Practice of Parallel Programming, 2016 (PPoPP '16). Barcelona, Spain.
38. Shasha Wen, Xu Liu, and **Milind Chabbi**. *Runtime Value Numbering: A Profiling Technique to Pinpoint Redundant Computations*. In proceedings of the ACM SIGARCH international conference on Parallel Architectures and Compilation Techniques, 2015 (PACT '15). San Francisco, CA.
39. Ashwin Aji, Lokendra Panwar, Wu-chun Feng, Pavan Balaji, James Dinan, Rajeev Thakur, Feng Ji, Xiaosong Ma, **Milind Chabbi**, Karthik Murthy, John Mellor-Crummey, and Keith Bisset. *MPI-ACC: Accelerator-Aware MPI for Scientific Applications*. In IEEE Transactions on Parallel and Distributed Systems (TPDS). Vol. PP, issue 99, pages 1-14, 2015.
40. **Milind Chabbi**, Michael Fagan, and John Mellor-Crummey. *High Performance Locks for Multi-level NUMA Systems*. In proceedings of the ACM SIGPLAN symposium on Principles and Practice of Parallel Programming, 2015 (PPoPP '15). San Francisco, CA.
41. **Milind Chabbi**, Wim Lavrijsen, Wibe de Jong, Koushik Sen, John Mellor-Crummey, and Costin Iancu. *Barrier Elision for Production Parallel Programs*. In proceedings of the ACM SIGPLAN symposium on Principles and Practice of Parallel Programming, 2015 (PPoPP '15). San Francisco, CA.
42. **Milind Chabbi**, Xu Liu, and John Mellor-Crummey. *Call Paths for Pin Tools*. In proceedings of the international symposium on Code Generation and Optimization, 2014 (CGO '14). Orlando, FL.
43. **Milind Chabbi**, Karthik Murthy, Michael Fagan, and John Mellor-Crummey. *Effective Sampling-Driven Performance Tools for GPU-Accelerated Supercomputers*. In proceedings of the international conference for High Performance Computing, Networking, Storage and Analysis (SC'13). Denver, CO.
44. Ashwin Aji, Lokendra Panwar, Wu-chun Feng, Pavan Balaji, James Dinan, Rajeev Thakur, Feng Ji, Xiaosong Ma, **Milind Chabbi**, Karthik Murthy, John Mellor-Crummey, and Keith Bisset. *On the Efficacy of GPU-Integrated MPI for Scientific Applications*. In proceedings of the symposium on High-Performance Parallel and Distributed Computing (HPDC '13). New York City, NY.
45. Sanjay Chatterjee, Zoran Budimlić, Vincent Cavé, **Milind Chabbi**, Max Grossman, Saĝnak Taşırlar, Yonghong Yan, and Vivek Sarkar. *Integrating Asynchronous Task Parallelism with MPI*. In proceedings of the International Parallel and Distributed Processing Symposium (IPDPS '13). Boston, MA.
46. **Milind Chabbi**, and John Mellor-Crummey. *DeadSpy: A Tool to Pinpoint Program Inefficiencies*. In proceedings of the international symposium on Code Generation and Optimization, 2012 (CGO '12). San Jose, CA. (**Best presentation award**)
47. **Milind Chabbi**, John Mellor-Crummey, and Keith Cooper. *Efficiently Exploring Compiler Optimization Sequences with Pairwise Pruning*. In proceedings of the 1st international workshop on Adaptive Self-Tuning Computing Systems for the Exaflop Era (EXADAPT '11). San Jose, CA.

POSTERS

- Chris Zhang, Murali Krishana Ramanathan, Prithvi Raj, Abhishek Parwal, Tim Sherwood, and **Milind Chabbi**. *Critical Path Analysis of Large-Scale Microservice Architecture*. At the USENIX Annual Technical Conference (ATC '22). July, 2022.
- Muhammad Aditya Sasongko, **Milind Chabbi**, and Didem Unat. *Novel Reuse-distance Tool for Multicore*. In proceedings of the ACM SIGARCH international conference on Parallel Architectures and Compilation Techniques, 2020 (PACT '20). Virtual event.
- Haris Volos, Kimberly Keeton, Yupu Zhang, **Milind Chabbi**, Se Kwon Lee, Mark Lillibridge, Yuvraj Patel, and Wei Zhang. *Software Challenges for Persistent Fabric-attached Memory*. USENIX Symposium on Operating Systems Design and Implementation (OSDI'18). Carlsbad, CA. October, 2018.
- Haris Volos, Kimberly Keeton, Yupu Zhang, **Milind Chabbi**, Se Kwon Lee, Mark Lillibridge, Yuvraj Patel, and Wei Zhang. *Memory-Oriented Distributed Computing at Rack Scale*. ACM Symposium on Cloud Computing 2018. Carlsbad, CA. October, 2018.
- **Milind Chabbi**, Xu Liu, and John Mellor-Crummey. *Performance Challenges for Emerging HPC Systems*. 7th Oil and Gas High Performance Computing Workshop. Houston, TX. March, 2014.
- **Milind Chabbi**, Karthik Murthy, Michael Fagan, and John Mellor-Crummey. *HPCToolkit: A Tool for Performance Analysis on Heterogeneous Supercomputers*. 6th Oil and Gas High Performance Computing Workshop. Houston, TX. February, 2013.
- **Milind Chabbi**, Karthik Murthy, Michael Fagan, and John Mellor-Crummey. *Next Generation Performance Tools for Heterogeneous Architectures*. 5th Oil and Gas High Performance Computing Workshop. Houston, TX. March, 2012.

TUTORIALS and PANELS

- Xu Liu, **Milind Chabbi**, Derek Bruening, Pengfei Su, and Qidong Zhao. *Enabling Easy Fine-Grained Binary Analysis on ARM and X86 Architectures*. Tutorial held at the International Symposium on Code Generation and Optimization (CGO'22). Virtual conference.
- **Milind Chabbi**, Shasha Wen, and Xu Liu. *CCTLib: Pinpointing Software Inefficiencies with Fine-grained Program Monitoring*. Tutorial held at the International Symposium on Code Generation and Optimization (CGO'17). Austin, TX. February, 2017.
- Harish Patil, Charles Yount, **Milind Chabbi**, and Cristiano Pereira. *Using PinPlay for Reproducible Analysis and Replay Debugging*. Invited speaker at PLDI'15 tutorial. Portland, OR. June, 2015.
- **Milind Chabbi**. *Critically Missing Pieces on Accelerators: A Performance Tools Perspective*. Invited talk at Birds of a Feather session at SC'13. Denver, CO. November, 2013.
- David Goodwin, Guido Juckeland, Allen Malony, **Milind Chabbi**, Stan Tomov. *Using the CUDA Profiling API and Related Third Party Tools*. Panelist at Nvidia GTC. San Jose, CA. March, 2013.

INVITED TALKS

- **Milind Chabbi**. *Demystifying Concurrency Bugs in Golang*. Invited talk University of Edinburgh, UK. March, 2024.
- **Milind Chabbi**. *Demystifying Concurrency Bugs in Golang*. Invited talk Royal Holloway University of London, UK. March, 2024.
- **Milind Chabbi**. *Demystifying Concurrency Bugs in Golang*. Invited talk Aarhus University, Denmark. March, 2024.
- **Milind Chabbi**. *Demystifying Concurrency Bugs in Golang*. Invited talk at Systems Group, North Carolina State University. April, 2023.
- **Milind Chabbi**. *Demystifying Concurrency Bugs in Golang*. Invited talk at Reliable and Secure Systems group, Purdue. Feb, 2023.
- **Milind Chabbi**. *A Study of Real-World Data Races in Golang*. Invited talk at ACM Baltimore Chapter Seminar. Aug, 2022.

- **Milind Chabbi** and Alan Lo. *Evolving P4 Runtime from switch to DPU*. Invited talk at P4 Workshop. Virtual. May, 2022.
- **Milind Chabbi**. *A Study of Real-World Data Races in Golang*. Invited talk at School of Computer Science, Georgia Tech. Atlanta, GA. May, 2022.
- **Milind Chabbi**. *An Experience with Code-Size Optimization for Production iOS Mobile Applications*. Invited talk at Indian Institute of Technology Dharwad. Dec, 2021.
- **Milind Chabbi**. *Optimistic Concurrency for Real-World Go Programs*. Invited talk at Google Systems Research Group. Nov, 2021.
- **Milind Chabbi**. *Optimistic Concurrency for Real-World Go Programs*. Invited talk at Amazon Web Services AI/ML group. Aug, 2021.
- **Milind Chabbi**. *Performance: the Role of Languages, Compilers, and Architectures*. Invited talk at Uber Performance and Efficiency Summit, Palo Alto CA. April, 2019.
- **Milind Chabbi**. *Microscopic Performance Analysis*. Invited talk at hardware infrastructure meet-up, San Francisco, CA. Jan, 2019.
- **Milind Chabbi**. *A New Perspective on Software Inefficiency*. Invited talk at Nvidia Research, Santa Clara, CA. March, 2017.
- **Milind Chabbi**. *Software and Tools for HPE The Machine Project*. Invited speaker at Scalable Tools Workshop 2016. Lake Tahoe, CA. August, 2016.
- **Milind Chabbi**. *Software Support for Efficient Use of Modern Parallel Systems*. Invited talk at Microsoft Research, Redmond, WA. April, 2015.
- **Milind Chabbi**. *Software Support for Efficient Use of Modern Parallel Systems*. Invited talk at Lawrence Berkeley National Laboratory. Berkeley, CA. February, 2015.
- **Milind Chabbi**. *Software Support for Efficient Use of Modern Parallel Systems*. Invited talk at Hewlett Packard Laboratory. February, 2015.
- **Milind Chabbi**. *Software Support for Efficient Use of Modern Parallel Systems*. Invited talk at PepperData Inc.. Sunnyvale, CA. February, 2015.
- **Milind Chabbi**. *Software Support for Efficient Use of Modern Parallel Systems*. Invited talk at ARM Research and Development. Austin, TX. January, 2015.
- **Milind Chabbi**. *Software Support for Efficient Use of Modern Parallel Systems*. Invited talk at Intel Programming Systems Laboratory. Santa Clara, CA. January, 2015.
- **Milind Chabbi**. *Performance Analysis and Optimizations of NWChem*. Invited talk at the Dynamic Exascale Global Address Space Programming Environments (DEGAS) group. May, 2014.
- **Milind Chabbi**. *Call Paths to Pinpoint Performance and Correctness Problems in Parallel Programs*. Invited talk at Corvette research group. University of California Berkeley, CA. February, 2014.
- **Milind Chabbi**. *Call Paths for Pin Tools*. Talk at the graduate seminar in the dept. of computer science, Rice University. Houston, TX. February, 2014.
- **Milind Chabbi**. *HPCToolkit: Performance Tools for Tuning Application on Heterogeneous Supercomputers*. Invited talk at Dagstuhl seminar on Automatic Application Tuning for HPC Architectures. Dagstuhl, Germany. October, 2013.
- **Milind Chabbi**, Karthik Murthy, and John Mellor-Crummey. *HPCToolkit: A Tool for Performance Analysis on Heterogeneous Supercomputers*. Invited talk at Nvidia GTC. San Jose, CA. March, 2013.
- **Milind Chabbi**. *Scalable Performance Analysis on Heterogeneous Architectures with HPCToolkit*. Invited talk at CScADS 2012 workshop on Performance Tools for Extreme-scale Computing. Snowbird, Utah. June, 2012.
- **Milind Chabbi**. *An Unconventional Approach for Performance Analysis of Modern Software Systems*. Talk at the graduate seminar in the dept. of computer science, Rice University. Houston, TX. March, 2012.

USPTO PATENTS

- Concurrent Reading and Writing with Crash Recovery in Persistent Memory (2019).
- Discontiguous Storage and Contiguous Retrieval of Logically Partitioned Data (2017).
- Path-synchronous Performance Monitoring of Interconnection Networks Based-on Source-code Attribution (2017).
- Hierarchical Virtual File Systems for Accessing Data Sets (2017).
- A Persistent Ticket Lock (2016).
- A Unified Lock (2016).
- Access of an MCS Locks by Guests (2015).
- Low-level Code Rewriter Verification (2010).

AWARDS AND ACHIEVEMENTS

- *Critical Path Analysis of Large-Scale Microservice Architectures*, Uber Developer Platform Technology award. May, 2022.
- *Offline data race detection for Go*, Uber Infrastructure Technology award. July, 2021.
- DRCCCTPROF: *A Fine-grained Call Path Profiler for ARM-based Clusters*, best paper nominee at SC'20.
- *ComDetective: A Fast and Accurate Communication Detection Tool for Threads*, best paper nominee at SC'19.
- *Load Redundancy: A Software Inefficiency Indicator*, distinguished paper award at ICSE'19.
- *Lightweight Hardware Transactional Memory Profiling*, best paper award at PPOPP '19.
- *Featherlight On-the-fly False Sharing Detection*, best paper award at PPOPP '18.
- *Watching for Software Inefficiencies with Witch*, ASPLOS 2018 Highlights paper and SIGPLAN Research Highlights.
- *RedSpy: Exploring Value Locality in Software*, ASPLOS 2017 Highlights paper.
- Travel grant awards: *PPoPP '15*, *CGO '14*, *SC '13*, and *CGO '12*.
- *Oil & Gas HPC Workshop Graduate Fellowship* 2013-2014 and 2014-2015.
- *Best Presentation Award* at the Intl. Symp. on Code Generation and Optimization, CGO '2012.
- *Rice University Graduate Fellowship* for the academic year 2010-2011.
- *Texas Instruments Fellowship* for the full duration of Ph.D. at Rice University.
- *Microsoft Lighthouse Award* for elimination of cascaded recompilation of .Net assemblies.

OPEN-SOURCE CONTRIBUTIONS

- CRISP: Critical Path Analysis for Microservice Architectures. [Tool](#).
- Bringing optimistic concurrency to the masses via Source-to-Source transformation in Golang. [Prototype tool](#).
- Hardware performance counters for CPU profiling in Go. [Proposal](#) and [prototype](#).
- Accelerating Alzheimer's Research, a collaboration between HPE Labs and German Center for Neuro-Degenerative disease that resulted in speeding a petabyte bio-medical pipeline by 100x and resulted in business for HPE.
https://github.com/makaho/kallisto/tree/memory_driven_computing
- Contribution to the Linux performance events interface that sped up a use case by 10x.
<https://github.com/torvalds/linux/commit/32ff77e8cc9e66cc4fb38098f64fd54cc8f54573>
- CCTLib, a suite of performance analysis tools and a framework for call path collection.
<https://github.com/CCTLib/cctlib>
- Lightweight tools for performance analysis (extensions over Rice University's HPCToolkit).
<https://github.com/WitchTools/hpctoolkit>

- A network performance monitoring tool that tracks packets as they flow through the network components and associates performance measurements in each interconnection component to the application source code.
https://github.com/HewlettPackard/genz_tools_network_monitoring
- A high-performance lock for NUMA systems adopted in the MPICH MPI (Message Passing Interface) library.
https://github.com/pmodels/izem/blob/22f63b31fe5d3b006a8922fc63eda0f93fa2ae87/src/lock/zm_hmcs.c
- Contributions to the open-source ROSE source-to-source compiler that among others sped up the compilation time by 10x.
<https://github.com/rose-compiler/rose/search?utf8=%E2%9C%93&q=milind&type=>

PRESS

- LeakProf: Featherlight In-Production Goroutine Leak Detection.
<https://www.uber.com/en-US/blog/leakprof-featherlight-in-production-goroutine-leak-detection/>.
- Data Race Patterns in Go. [**Over 100K views**].
<https://www.uber.com/blog/data-race-patterns-in-go/>.
- Dynamic Data Race Detection in Go Code.
<https://eng.uber.com/dynamic-data-race-detection-in-go-code/>.
- CRISP: Critical Path Analysis for Microservice Architectures.
<https://eng.uber.com/crisp-critical-path-analysis-for-microservice-architectures/>.
- pprof++: A Go Profiler with Hardware Performance Monitoring.
<https://eng.uber.com/pprof-go-profiler/>
- How Uber Deals with Large iOS App Size.
<https://eng.uber.com/how-uber-deals-with-large-ios-app-size/>
- Closer to a Cure: Combating Alzheimer’s with New Compute Technology.
<https://www.youtube.com/watch?v=c3M8q1X6ydg>
<https://www.youtube.com/watch?v=tJ0KwMTz2iI>
- German research institute tests Memory-Driven Computing to fight neurodegenerative diseases.
<https://www.hpe.com/us/en/customer-case-studies/dzne-memory-driven.html>
- Take a look behind The Machine.
<https://www.labs.hpe.com/the-machine/behind-the-machine>
- A [sampling-based dependency analysis for performance analysis on GPUs](#) is adopted by Nvidia in its NSight, NVIDIA Visual Profiler and shown as one of the critical features of its CUDA 8.0 release.
<https://devblogs.nvidia.com/cuda-8-features-revealed/>

TEACHING

Instructor for the following courses:

- *Parallel Computing*, Spring 2022, 2023, 2024 at IIT Dharwad.

Teaching assistant for the following courses:

- *Parallel Computing*, Spring 2013, at Rice University.
- *Multicore Computing*, Fall 2012, at Rice University.
- *Operating Systems and Concurrent Programming*, Spring 2012, at Rice University.
- *Software Construction Methodologies*, Fall 2011, at Rice University.
- *Introduction to Computer Systems*, Spring 2011, at Rice University.
- *Operating Systems*, Spring 2007, at the University of Arizona.
- *Theory of Computation*, Fall 2006, at the University of Arizona.

ADVISING

- *Deepan Raj*, Indian Institute of Technology, Dharwad, India.
Doctoral committee member [2023-].
- *Vivek Shahare*, Indian Institute of Technology, Dharwad, India.
Doctoral committee member [2022-].
- *Qidong Zhao*, North Carolina State University:
Doctoral committee member [2021-].
- *Muhammad Aditya Sasongko*, Koc University, Turkey:
Doctoral committee member [2022].
- *Yufan Xu*, University of Utah:
Ph.D. research intern mentor at Uber Technologies [2023-].
- *Feiyang Jin*, Georgia Tech:
Ph.D. research intern mentor at Uber Technologies [2023-].
- *Georgian-Vlad Saioc*, Arhaus University:
Ph.D. research intern mentor at Uber Technologies [2022-].
- *Chris Zhang*, University of California Santa Barbara:
Doctoral committee member and Ph.D. research intern mentor at Uber Technologies [2020-2023].
- *Pengfei Su*, College of William and Mary:
Doctoral committee member and Ph.D. research intern mentor at Uber Technologies [2019-20].
- *Shasha Wen*, College of William and Mary:
Doctoral committee member and Ph.D. research intern mentor at Hewlett Packard Labs [2016-19].
- *Qingsen Wang*, College of William and Mary:
Doctoral committee member, College of William and Mary [2019].
- *Adarsh Yoga*, Rutgers University:
Ph.D. research intern mentor at Hewlett Packard Labs [2017].
- *Cindy Orozco Bohorquez*, Stanford University:
Ph.D. research intern mentor at Hewlett Packard Labs [2017].
- *Lorenzo Martinico*, The University of Edinburgh:
undergraduate research intern mentor at Hewlett Packard Labs [2017].
- *Wei Zhang*, The George Washington University:
Ph.D. research intern mentor at Hewlett Packard Labs [2017].
- *Timothy Chong*, Stanford University:
Ph.D. research intern mentor at Hewlett Packard Labs [2017].
- *Sara Karamati*, Georgia Institute of Technology:
Ph.D. research intern mentor at Hewlett Packard Labs [2017].
- *Tuan Tran*, University of California at Santa Cruz:
Ph.D. research intern mentor at Hewlett Packard Labs [2017].

PROFESSIONAL SERVICE

- Program Co-Chair for the 29th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming (PPoPP '24).
- Program Committee member for the ACM/IEEE Symposium on Code Generation and Optimization (CGO '24).
- Program Committee member for the IEEE International Conference on Cloud Engineering(IC2E '24).
- Program Committee member for the ACM SIGPLAN/SIGBED International Conference on Languages, Compilers, and Tools for Embedded Systems (LCTES 2024).
- Program Committee member for the IEEE International Conference on Conference on Cluster Computing (Cluster '23).

- Program Committee member for the ACM/IEEE Symposium on Code Generation and Optimization (CGO '23).
- Program Committee member for the IEEE Cluster Conference in Programming and Systems Software stack (Cluster '23).
- Program Committee member for the ACM/IEEE Symposium on Code Generation and Optimization (CGO '23).
- Program Committee member for the 28th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming (PPoPP '23).
- Artifact Evaluation co-chair for the 27th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming (PPoPP '22).
- Program Committee member for the International Conference on Parallel Architectures and Compilation Techniques (PACT'21).
- Program Committee member for the International Conference on Supercomputing (ICS'21).
- Program Committee member for the International Conference on Languages, Compilers, Tools and Theory of Embedded Systems (LCTES'21).
- Program Committee member for 27th International European Conference on Parallel and Distributed Computing (Euro-Par'21).
- Artifact Evaluation co-chair for the 26th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming (PPoPP '21).
- Program Committee member for PACMPL conference on Object-Oriented Programming, Systems, Languages, and Applications (OOPSLA '20).
- Program Committee member for the 26th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming (PPoPP '21).
- Program Committee member for the 25th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming (PPoPP '20).
- Program Committee member for the 32nd Workshop on Languages and Compilers for Parallel Computing (LCPC '2019).
- External Program Committee member for the 25th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming (PPoPP '20).
- Program Committee member, Performance Track, the International Conference for High Performance Computing, Networking, Storage, and Analysis (SC '19).
- External Program Committee member for the International Conference on Compiler Construction (CC '19).
- Program Committee member for the 24th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming (PPoPP '19).
- Program Committee member for the ACM/IEEE Symposium on Code Generation and Optimization (CGO '18).
- External Program Committee member for the 23rd ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming (PPoPP '18).
- External reviewer for the 2017 International Conference for High Performance Computing, Networking, Storage and Analysis (SC '17),
- Program Committee member for the ACM/IEEE Symposium on Code Generation and Optimization (CGO '17).
- Program Committee member for the 22nd ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming (PPoPP '17).
- Program Committee member 21st International Workshop on High-level Parallel Programming Models and Supportive Environments (HIPS '16).
- Program Committee member for the 6th International Workshop on Adaptive Self-tuning Computing Systems (ADAPT '16).

- External Program Committee member for the 21st ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming (PPoPP '16).
- Reviewer for the IEEE Journal of *Transactions on Computers (TC)*.
- Reviewer for the IEEE Journal of *Transactions on Parallel and Distributed Systems (TPDS)*.
- Reviewer for the Journal of *Software: Practice and Experience (SPE)*.
- Reviewer for the Journal of *Concurrency and Computation: Practice and Experience (CPE)*.
- Reviewer for the ACM Journal of *Transactions on Computer Systems (TOCS)*.
- Student consultant at the *Center for Written, Visual, and Oral Communication* at Rice University [2014-Present].
- Program committee member for the seventh international workshop on multi-/many-core computing systems (MuCoCoS '2014).
- Student coach at the *Professional Communication for the School of Engineering* at Rice University [2012-Present].
- Member of the advisory committee for the *Center for Career Development* at Rice University [2012-2013].
- Representative of the computer science department for the *Graduate Student Association (GSA)* at Rice University [2012-2014].
- Secondary reviewer for the 2012 International Conference for High Performance Computing, Networking, Storage and Analysis (SC '12),
- Secondary reviewer for the Symposium on Principles and Practice of Parallel Programming (*PPoPP '12*).
- Plenary presenter at the 5th Oil and Gas High Performance Computing Workshop, Houston, TX, March, 2012.
- Secondary reviewer for the 2011 Symposium on Principles and Practice of Parallel Programming (*PPoPP '11*).
- Secondary reviewer for the 2011 Workshop on Languages and Compilers for Parallel Computing (*LCPC '11*).