Vivek Sarkar

Curriculum Vitae

Contents

1 Summary **2**

2 Education **3**

3 Professional Experience **3**

3.1 2017-present: College of Computing, Georgia Institute of Technology 3

3.2 2007-present: Department of Computer Science, Rice University 5

3.3 1987-2007: International Business Machines Corporation 8

4 Professional Awards **11**

5 Research Awards **12**

6 Industry Gifts **15**

7 Graduate Student and Other Mentoring **17**

8 Professional Service **19**

8.1 Conference Committees 20

8.2 Advisory/Award/Review/Steering Committees 25

9 Keynote Talks, Invited Talks, Panels (Selected) **27**

10 Teaching **33**

11 Publications **37**

11.1 Refereed Conference and Journal Publications 37

11.2 Refereed Workshop Publications 52

11.3 Books, Book Chapters, and Edited Volumes 59

12 Patents **60**

13 Software Artifacts (Selected) **61**

14 Personal Information **61**

1 Summary

During my twenty years in industry (IBM) and fourteen years in academia (Rice, Georgia Tech), I have led multiple research organizations while also working closely with industry partners (AMD, IBM, Intel, Microsoft, NVIDIA) and program managers in federal agencies (DARPA, DOE, NSA, NSF). My administrative experience at IBM includes serving as Senior Manager of the Programming Technologies department in IBM Research during 2000–2007, a department that grew from fifteen to forty PhD-level research staff members as core personnel during that period. Multiple accomplishments from that period in my department are still hailed as top accomplishments for IBM Research. More recently, I served as Chair of the Department of Computer Science at Rice during 2013–2016, a period during which I oversaw doubling of the number of CS undergraduate majors with an increase in the fraction of women from 23% to 35% at the same time, reversal in the downward funding trend for the department prior to my becoming chair, and growth in our on-campus MS CS program to become the second-largest master's program on campus. After moving to Georgia Tech in 2017, my administrative experiences include serving as co-director of the Center for Research into Novel Computing Hierarchies (CRNCH) and as Chair of the School of Computer Science since August 2020. As CRNCH co-director, I worked with the founding co-director to expand the scope of post-Moore computing research in the center to include new thrusts related to software. As school chair, I rebuilt our community during the pandemic and led a successful Spring 2021 faculty recruiting season with five new faculty hires, including our school's first female hire in many years (the next "juniormost" female faculty member in our school obtained her PhD in 2007).

My research has led to wide-ranging technical and leadership accomplishments related to programming model, compiler, runtime, and verification systems for parallel computing. I led the parallel programming models, tools, and productivity thrust in the DARPA HPCS IBM PERCS project during 2002–2007; this included the creation of the X10 language which received a most influential paper award in 2015. During 1998–2001, I led the creation of the Jikes Research Virtual Machine, a unique JVM design and open-source implementation that has been used by over a hundred universities worldwide for research and teaching, and received the 2012 ACM SIGPLAN Programming Languages Software Award. Since 2007, I have obtained sponsored research grants of over $20M (excluding subcontracts for partners) for my Habanero Extreme Scale Software Research Laboratory at Rice and Georgia Tech, which have resulted in a number of research contributions as well as influences on industry standards for parallelism. At IBM, my contributions were recognized by major awards, induction into the IBM Academy of Technology in 1995, and my appointment in relationship management roles with senior customer executives. I am an ACM Fellow, an IEEE Fellow, and recipient of the 2020 ACM-IEEE CS Ken Kennedy Award.

My external service includes serving as a member of the US Department of Energy's Advanced Scientific Computing Advisory Committee (ASCAC) since 2009, and on CRA's Board of Directors since 2015. I helped create the new CRA-Industry standing committee (as founding co-chair) with the goal of bringing together industry partners in computing research to convene discussions and actions on topics of mutual strategic interest with CRA's academic and government constituents.

2 Education

1987 PhD, Computer Systems Laboratory, Dept. of Electrical Engineering, Stanford University.
Thesis: "Partitioning and Scheduling Parallel Programs for Execution on Multiprocessors".
Advisor: Professor John L. Hennessy.

1982 MS, Computer Science, University of Wisconsin-Madison.

1981 BTech, Electrical Engineering, Indian Institute of Technology, Kanpur, India.
Project title: "Command Language Processor Generator".
Advisor: Professor Kesav V. Nori.

3 Professional Experience

3.1 2017-present: College of Computing, Georgia Institute of Technology

After 10 years at Rice University, my move to Georgia Tech in August 2017 was motivated in part by the opportunity to collaborate on large-scale research projects related to post-Moore computing with world-class experts in computer architecture for high performance computing. These opportunities were realized by the successful transition of my lab from Rice, and its expansion into software-hardware codesign research topics as evidenced by my group's recent publications in hardware-focused venues as well as the broader set of research grants that I have been able to obtain with my colleagues.

2020-present: Chair, School of Computer Science

The School of Computer Science (SCS) consists of approximately 40 tenure-track/tenured faculty members who lead Georgia Tech's College of Computing in seven foundational areas of Computer Science – Architecture, Databases, Networking, Programming Languages & Software Engineering, Security, Systems and Theory. While starting as school chair in August 2020 posed multiple challenges due to the COVID-19 pandemic, it has been rewarding to see how our faculty and students have stepped up to the challenge of maintaining continuity in research and education during this period. My top priority during this pandemic year has been to focus on the wellbeing of our school's community of faculty, staff, and students. This involved ensuring that our entire school community always had access to the latest information and resources related to COVID-19 testing and vaccination, and also scheduling regular drop-in office hours and end-semester faculty retreats. Upon learning of the significant challenges being faced by our graduate students, I encouraged them to schedule regular virtual social events and also mentored them in organizing as a formal student organization at Georgia Tech. At the same time, I found opportunities to broaden the perspective of school faculty to look beyond their specific research area and grow into thought leaders in integrating CS research areas in novel ways motivated by solving society's hardest and most urgent CS-related challenges.

2018-2020: Co-director, Center for Research into Novel Computing Hierarchies

The premise underlying the Center for Research into Novel Computing Hierarchies (CRNCH), an Interdisciplinary Research Center at Georgia Tech, is that we are at the forefront of a new computing revolution as Moore's Law comes to an end, and that is it is time to rethink all foundations of computing technology. The center was created by my colleague, Tom Conte, an expert in computer Architecture who is also co-chair of IEEE's Rebooting Computing Initiative. After my arrival at Georgia Tech, Tom recommended to the dean that I be appointed co-director to lead the software directions in the center. I appreciated being welcomed into this role, and one of my priorities has been to further enhance collaborations among CS and ECE faculty whose research overlaps with the scope of the center. My PI role in two DARPA projects is also highly synergistic with advancing the goals of CRNCH.

2018-2019: Principal Investigator, DDARING project in DARPA Software Defined Hardware program

The DDARING project was funded in 2018 as a software-focused TA2 effort within the Software Defined Hardware (SDH) program, which is part of DARPA's recent Electronics Resurgence Initiative (ERI). As PI, my responsibility includes coordinating the work of all the co-PIs, Tom Conte and Rich Vuduc from Georgia Tech, Deming Chen and Wen-Mei Hwu from UIUC, Scott Mahlke from U.Michigan, and Viktor Prasanna from USC, and their team members, towards a unified vision of the project goals. In addition, my group contributes directly to the programming model and compiler aspects of this project. The DDARING project aims to develop a novel programming system to accelerate data-intensive workflows with the productivity levels that analysts have come to expect from modern problem-solving environments such as using Python with Jupyter notebooks Our approach encapsulates a unification of software compilation and hardware reconfiguration techniques. It includes a new Python-based programming model, called Intrepydd, that is suitable for ahead-of-time compilation on current and future hardware platforms and accelerators, thereby enabling performance programming in a lightweight Python syntax. It also includes extensions for matrix, tensor, graph, and deep learning computations that expose opportunities for dynamic data-aware code optimizations and hardware reconfiguration. Intrepydd is not intended for writing complete/main programs. Instead, code generated from Intrepydd programs can be compiled so as to be integrated within a Python application or a C++ application. The DDARING system also includes a static code optimizer, a dynamic code optimizer, an auto-tuner for selecting code-and-configuration pairs, and a knowledge base that learns and stores information on the mapping of workflow steps to compiler/library generated kernels, and the mapping of kernels to available hardware configurations.

2017-present: Professor, School of Computer Science, Stephen Fleming Chair for Telecommunications, College of Computing, Georgia Institute of Technology

I lead the Habanero Extreme Scale Software Research Laboratory, which moved from Rice University to Georgia Tech in 2017 and currently includes eleven PhD students, one postdoctoral researcher, and four research scientists. After the move, I continued our group's prevailing research thrusts related to software foundations for extreme-scale computing, but also expanded our research into software-hardware codesign topics. In the area of software foundations, my group has continued research on advancing the state of the art in programming models, optimizing compilers, runtime systems, and debugging for homogeneous and heterogeneous parallel systems. Our group's expansion into software-hardware codesign research topics is evidenced by our recent publications in hardware-focused venues such as ISCA'18, MICRO'18, FCCM'19 and MICRO'19, and algorithms-focused venues such as NeurIPS'18.

With respect to internal impact in Georgia Tech, I became Co-Director of the Center for Research into Novel Computing Hierarchies (CRNCH) in 2018. In addition, I served on the Search Committee for the new Dean of the College of Computing; in this role, I paid special attention to recruiting top diversity candidates to apply, one of whom emerged as a finalist. Since 2018, I have served as our school's representative on the college-level Reappointment, Promotion, and Tenure (RPT) committee. In Spring 2018, I served on our school's Faculty Recruiting Committee (FRC), and played an active role in encouraging five faculty candidates to accept Georgia Tech's offer that year; I now chair our school's FRC for the 2019-20 academic year. I was also elected by my faculty colleagues to be Chair of the SCS School Advisory Committee (SAC), and I served on the committee that selects candidates in the College of Computing to nominate for recognition as Regents Professors. As listed in Section 10, my teaching responsibilities include the undergraduate compilers class (CS 4240) and advanced graduate-level classes in compilers (CS 6245) and programming languages (CS 6390).

3.2 2007-present: Department of Computer Science, Rice University

After 20 years at IBM, my move to Rice University was motivated in part by the mainstream adoption of parallel computing, with the end of increases in uniprocessor clock frequency. I believed that academia offered unique opportunities for a broader impact of my parallel computing research through partnerships with multiple industry partners, as well as the creation of teaching material for the new era of mainstream parallel computing. These opportunities were realized by the creation of the Habanero Extreme Scale Software Research Laboratory which has served as a wellspring for advances in the foundations of parallel computer software during the last decade. Our lab's research has also enabled technology transfer to industry partners and open standards, and provided the context for introducing a new undergraduate course (COMP 322) at Rice University and a related online specialization on Coursera.

2017-present: Adjunct Professor, Department of Computer Science, Rice University

This appointment has enabled the transition of my lab from Rice to Georgia Tech to occur smoothly with minimal disruption on students and research staff members in my lab.

2014-2019: Principal Investigator, Pliny project in DARPA MUSE program

The core Pliny project was funded by the DARPA program on Mining and Understanding Software Enclaves (MUSE) for four years at $11.1M in 2014. In 2018, we obtained funding for an extension effort until December 2019 to focus on technology transfer activities with GitHub as a primary transition partner. As PI, my responsibility includes coordinating the work of all the co-PIs from Rice University (Swarat Chaudhuri, Chris Jermaine), GrammaTech (David Melski), University of Wisconsin (Tom Reps, Ben Liblit), and UT Austin (Isil Dillig, Thomas Dillig), and their team members. Our proposal addressed Technical Areas 2–4 of the BAA to create a system that leverages "big code" to achieve orders-of-magnitude productivity gains for important software engineering tasks, including code search, bug-finding, code repair and code synthesis. At the heart of our approach is the creation of new statistical models that correlate program syntax with program semantics and informal specifications that users of programming tools may have in mind. Code search is enabled by extraction of software features that go beyond what is typically used in text mining. Anomaly detection is enabled by computing anomaly scores as a statistical distance between the inferred distribution of program behaviors, and the actually observed behavior of the input program. In practice, statistical anomalies correspond to software defects or non-standard coding practices, both of which provide valuable information to the software development process. Our project also explored both search-based and learning-based approaches to code synthesis. In search-based synthesis, our code search technology is used to find candidate code fragments that can be used to complete code sketches. In learning-based synthesis, our approach works by first generating a latent variable that summarizes the user input (for example, a set of types, API calls, keywords, or natural language) and returning candidate code fragments by sampling the learned distribution of programs previously ingested from the corpus, while being conditioned on the specification. Finally, the Pliny project also included the development of the PlinyCompute platform, a scalable distributed system for developing high-performance, data-intensive, distributed computing tools and libraries, such as those needed by the search, anomaly detection, and synthesis tools in the Pliny project.

2013-2016: Chair, Department of Computer Science, Rice University

The Rice CS department is unique in many ways. When I started as department chair, it was a small department with 19 faculty members, and also a relatively young department created in 1984. The report that I created for our advancement review committee in my first semester as chair made a data-driven case for significantly growing the department. During my term as chair, we continued to lead in our historical areas of strengths, which included programming languages and parallel computing, while also growing our research portfolio in strategic areas such as bioinformatics,

machine learning and data analytics. As department chair, I oversaw a significant increase in the number of CS majors, making CS the largest undergraduate major at Rice, and, more importantly, an increase in the percentage of women CS undergraduate majors from 23% to over 35% (the highest in our department's history) at the same time. I also made a strategic push to significantly grow our on-campus Professional Master of Computer Science to become the second-largest professional program at Rice.

One of my major priorities when I became CS department chair was to focus on building a sense of community among all members of the department — academic faculty, teaching faculty, research faculty, administrative staff, graduate students, and undergraduate students. I initiated a tradition of "all hands" meetings where all members of these groups were invited, updates were shared from all the units represented, and recognitions were announced. I also made a point of meeting regularly with the engineering dean, Ned Thomas, so as to understand his priorities and gain his trust as a contributor to the school's goals not just the department goals. I sought multiple sources of revenue to support our community building goals and our challenges with booming enrollments in CS; these included local companies, as well as a strategy to increase the number of Professional Masters students. The additional revenue also enabled me to subsidize summer research internships for students in underrepresented groups who were interested in majoring in CS, and to provide funding to the graduate students to form a graduate student association.

As department chair at Rice, I interacted with leaders of graduate and undergraduate student organizations, helping shape their initiatives and connecting them with potential sponsors. My rapport with graduate students followed naturally from the fact that my lab had the largest number of graduate students in the department, and I got to know many of their colleagues in other groups as a result. I supported and funded the formation of a new CS Graduate Student Association (CS-GSA) at Rice in my first semester as department chair. A side-benefit of teaching the sophomore-level COMP 322 when I was department chair was that I also got to know most CS undergraduates well. When discussing important departmental issues with leaders of undergraduate CS clubs, it was rewarding to experience the high level of trust they had in me because of our past interactions in COMP 322. I advised and supported the rapid growth of club events, especially Hack Rice which grew from about 50 to over 300 participants while I was chair. Finally, I helped identify resources to support the activities of CSters, a club for undergraduate and graduate women CS students at Rice, which enabled a large number of them to attend the Grace Hopper Conference every year.

2007-2017: Professor of Computer Science, Professor of Electrical and Computer Engineering (joint appointment), E.D. Butcher Chair in Engineering, Rice University

Responsible for creating and leading the Habanero Extreme Scale Software research laboratory that aspires to unify elements of high-end computing, multicore, and embedded software stacks so as to produce portable software that can run unchanged on a range of homogeneous and heterogeneous extreme scale platforms (http://habanero.rice.edu). The research agenda for the Habanero project includes the creation of new parallel language, compiler, runtime, and verification technologies with

a common execution model foundation.

PI or co-PI of multiple research projects funded by DOE, DARPA, DOD, MARCO, and NSF. Co-PI and Associate Director of NSF Expeditions Center for Domain-Specific Computing started in 2009 (PI is Prof. Jason Cong at UCLA). Recipient of funding from multiple industry sponsors and partners including AMD, BHP Billiton, Halliburton Services, IBM Research, IBM Toronto laboratory, Intel, Microsoft, Samsung and Shell. Recipient of equipment donation from AMD, IBM, NVIDIA, Samsung, Texas Instruments, Samsung, and Sun Microsystems.

Soon after joining Rice, I worked with the vice provost for information technology to make the case to IBM to donate new resources for Rice's research computing center. I was able to communicate a win-win strategy for IBM to invest in Rice so as to increase their connections with the Texas Medical Center. This led to a $7.6M Shared University Research (SUR) equipment award in 2010 for a POWER7 system dubbed BlueBioU. At that time, the compute power of BlueBioU exceeded the combined total of Rice's existing supercomputers, and it was used very productively for eight years thereafter.

As listed in Section 10, my teaching responsibilities include the creation of a new undergraduate class on the fundamentals of parallel programming (COMP 322), and advanced graduate-level classes which included COMP 422 (Introduction to Parallel Computing), COMP 515 (Advanced Compilation for Vector and Parallel Processors, formerly taught by Prof. Ken Kennedy), and COMP 635 (Seminar on Heterogeneous Processors).

3.3 1987-2007: International Business Machines Corporation

I started my career at IBM Research in the glory days of industry research labs, and had Fran Allen and Alfred Spector as mentors who helped me grow as a leader. My career took shape in a climate that placed a high priority on recruiting, retaining, and nurturing technical and leadership talent. Each advance at IBM broadened my horizons, from a focus on individual research contributions, to my group's research agenda as a first-line manager, my department's technical strategy as a second-line manager, then to goals for the Programming Languages and Software Engineering technical area which included 100+ researchers at IBM locations worldwide, and finally, to the scale of IBM combined with its partners and customers. During 2000-07, as Senior Manager for the Programming Technologies Department, I led its growth from about fifteen to forty PhD-level CS research staff members (with qualifications comparable to those of CS faculty members), and an operating budget of over $12M/year. I was responsible for raising a significant fraction of my department's budget from IBM product groups and external sources. For example, collaborating with colleagues in hardware systems, I raised funding for R&D in the DARPA HPCS program and, as one of six key personnel, was instrumental in the success of IBM's proposals for funding in Phase 1 ($3M), Phase 2 ($53M) and Phase 3 ($244M). I also focused on developing technical and management talent in my organization. It has been very satisfying to see some of my department members grow into management roles and achieve recognition for their technical achievements as ACM Fellows and Distinguished Members. These include David Bacon, Evelyn Duesterwald, Stephen Fink, David

Grove, Michael Hind, and Frank Tip, among others.

My responsibilities at IBM went beyond my line management commitments in multiple ways. After my first three years in the PTRAN research group, Fran Allen recommended to our product division that I lead a product group at IBM's Santa Teresa Laboratory to build the new high-level ASTI compiler optimizer from scratch. Six years in IBM's Software Group, 1991-96, immersed me in a real-world environment during an especially challenging time for IBM as a company. In addition to executing successfully as product team lead and becoming a first-line manager, I became part of the inner circle of a strategy group created by Steve Mills, Senior Vice President, who depended on us to spearhead cultural changes needed for IBM to become more agile and competitive in a new era of fast-paced software development. In 1995, about 8 years after earning my doctorate, I was inducted into IBM's Academy of Technology. CEO Lou Gerstner viewed the Academy as his 300-person CTO and, at annual meetings with us, he spoke very candidly about IBM's strategic challenges and what kept him awake at night.

These product and business related experiences stood me in good stead when I returned to IBM Research in 1998 and was asked to take on IBM Research's Relationship Manager role for IBM's Rational Software division in addition to my line management responsibilities. I learned the importance of seeing challenges and priorities from the perspectives of my strategic partners, and was able to gain the trust of business unit executives that I would help them succeed instead of just pushing the funding needs of the Research division. When I learned of their plans to release Eclipse as an open-source platform in 2001, I proposed an initiative in which Research could help increase university adoption of Eclipse for research and teaching. In response to my proposal, I was given a budget of $2M/year for three years which I used to lead the creation of an Eclipse Innovation Grants program and worldwide community in academia. During this period (2002-2004), Eclipse evolved from being completely unknown into a ubiquitous platform for software development on most college campuses. I also served as IBM Research's Relationship Manager for the Aerospace and Defense industry sector, where I would assist sales teams, working on multi-million dollar contracts, by meeting with customers to discuss their future strategic needs and to identify areas of overlap with IBM's solutions and future research directions. I learned to become an advocate for all the CS research being undertaken at IBM, not just the research in my area.

2002-2007: Team Lead, PERCS Programming Models, Tools, and Productivity, IBM Research

Responsible for creating and leading the programming models, tools, and productivity research agenda in the DARPA-funded IBM PERCS project, and representing this area in all DARPA milestone reviews until June 2007. I was one of six key personnel for the project from the start, and contributed to the creation of IBM's proposals and their selected for funding in Phase 1 ($3M), Phase 2 ($53M) and Phase 3 ($244M) of the DARPA HPCS program. Two major outcomes of this effort have been the creation of the X10 language for future parallel systems which I co-led with Vijay Saraswat during 2004–2007, and the creation of the Eclipse Parallel Tools Platform

open source project with LANL and other partners. In Phase 2, my role also included funding and coordinating work performed by multiple university partners – UC Berkeley (Ras Bodik), MIT (Saman Amarasinghe), U. Delaware (Guang Gao), UT Austin (Jim Browne), U. Illinois (Ralph Johnson), Purdue University (Jan Vitek). I received two IBM Research Division awards and one IBM Outstanding Technical Achievement award at IBM for my work on PERCS.

2000-2007: Senior Manager, Programming Technologies Department, IBM Research

Responsible for initiating and overseeing research projects carried out by a department with approximately forty permanent researchers, additional temporary researchers, and an annual budget of over $12 million. These research projects spanned the areas of Programming Models (X10, XJ, Collage), Tools (Eclipse Parallel Tools Platform, Advanced Refactorings in Eclipse, Scalable And Flexible Error detection, Security analysis, Scripting analysis), and Optimized Execution Environments (Jikes RVM, Metronome, Progressive Deployment Systems). During this period, I personally led the open-source release of the Jikes Research Virtual Machine (RVM), and initiated multiple new projects including PERCS, X10, XJ, and DOMO/WALA. I received an IBM Outstanding Technical Achievement Award at IBM for my work on Jikes RVM, which included the leadership of the design and implementation of the Jikes RVM dynamic optimizing compiler. Jikes RVM has been used by over a hundred universities worldwide for research and teaching, and received the 2012 ACM SIGPLAN Programming Languages Software Award. As IBM's Research Relationship Manager for Aerospace and Defense, I interacted with multiple customers in that industry to discuss their future strategic needs and synergies with IBM's solutions to help them achieve their strategic goals. My Senior Manager responsibilities also included defining and coordinating IBM Research strategy around the open source Eclipse framework for programming tools (www.eclipse.org) and the Eclipse Innovation Grants program.

1998-2000: Manager, Dynamic Compilation Group, IBM Research

Conducted research in and led the group in the design and implementation of a new dynamic optimizing compiler for the Java language, as part of the Jalapeno project which resulted in the Jikes Research Virtual Machine. The design built on my personal research in the areas of Linear Scan register allocation, load/store elimination of heap references using Array SSA form, BURS-based register-sensitive instruction selection, and Array Bounds Check elimination on Demand (ABCD). The dynamic optimizing compiler also provided the foundation for research on adaptive optimization.

1996-1998: Visiting Associate Professor, MIT EECS department (on leave from IBM)

Restructured and taught MIT's main compilers class (6.035) in Fall 1996 and 1997. Co-invented the Array SSA form and Linear Scan register allocation analyses and optimizations. Founding member of the RAW project. Supervised PhD, Masters, and undergraduate students.

1995-2007: Member, IBM Academy of Technology

The IBM Academy is the top technical advisory group in the company. Academy membership is permanent and is a recognition of sustained technical contributions. Academy membership is limited to 300 people, less than 1% of IBM employees, who are recognized as the top technical leaders from IBM in areas that span software, hardware, manufacturing, applications, and services.

1994-1996: Manager, Application Development Technology Institute (ADTI), IBM Software Group

Managed the ADTI technology-transfer department with the mission of working with multiple research groups to invent, harvest, and transfer high priority technologies to IBM's programming language products. I personally led the design and implementation of automatic and OpenMP parallelization in IBM's XL Fortran product compilers and runtime systems for SMP systems, which built on the ASTI optimizer.

1991-1993: Project Leader, ASTI Optimizer, IBM Santa Teresa Laboratory

Conducted research in and led a 10-person team in the design and implementation of the ASTI optimizer, IBMs first high-level product optimizer for performing loop and data transformations for cache locality, efficient scalarization of Fortran 90 array language, SMP parallelization, and HPF compilation. The ASTI optimizer is still shipped as a key component of IBM's XL Fortran product compilers. I received an Outstanding Innovation Award at IBM for my work on ASTI.

1987-1990: Research Staff Member, PTRAN project, IBM Research

Conducted research on automatic partitioning of program dependence graphs into parallel tasks, and on new representations and uses of the control dependence relation, leading to a new foundation for selecting useful parallelism in program dependence graphs. Also contributed to the design and implementation of the PTRAN research system led by Fran Allen (recipient of 2006 ACM Turing Award).

4 Professional Awards

2020	Recipient, ACM - IEEE CS Ken Kennedy Award
2020	IEEE Fellow
2019	"Thank a Teacher" appreciation award, CS 6390, Fall 2019, Georgia Tech
2019	Best paper award, 15th International Workshop on OpenMP (IWOMP 2019).
2018	"Thank a Teacher" appreciation award, CS 6245, Fall 2018, Georgia Tech
2015	Most Influential OOPSLA Paper Award for OOPSLA 2005 paper on X10
2012	ACM SIGPLAN Programming Languages Software Award for Jikes RVM
2010	Best paper award, Runtime Verification conference (RV 2010)

2010 First Decade High Impact Paper, CASCON conference
2010 Outstanding Faculty Associate, Wiess College, Rice University
2008 ACM Fellow
2008 IBM Outstanding Technical Achievement Award, PERCS project (Phase 3)
2006 ACM Distinguished Scientist
2006 IBM Research Division Award, PERCS project (Phase 2)
2006 Honorable Mention, HPC Challenge Class 2 Award, Supercomputing 2006
2005 IBM Outstanding Technical Achievement Award, Jikes Research Virtual Machine
2003 IBM Research Division Award, PERCS project (Phase 1)
2001 IBM Research Division Award, Jikes Research Virtual Machine
1998 Third Invention Achievement Award, IBM
1998 Best paper, LCPC workshop (one of three best papers)
1998 Second Invention Achievement Award, IBM
1997 IBM Outstanding Innovation Award, ASTI optimizer
1997 Best paper, International Conference on Parallel Processing
1996 First Invention Achievement Award, IBM
1994 Best IBM paper, CASCON '94
1985 First place, ACM International Collegiate Programming Contest
1981 Best BTech project, IIT Kanpur

5 Research Awards

- DARPA PAPPA award, November 2019 - May 2021, "AutoMPHC: Automating Massively Parallel Heterogeneous Computing". PI: Vivek Sarkar, co-PIs: Taesoo Kim, Sukarno Mertoguno, Jun Shirako, Alexey Tumanov (all at GT). Total Award Amount: $1,000,000 (for Phases 1 and 2).

- DARPA SDH award, July 2018 - December 2019 (Phase 1), "DDARING: Dynamic Data-Aware Reconfiguration, INtegration and Generation". PI: Vivek Sarkar, co-PIs: David Bader, Tom Conte, Rich Vuduc (Georgia Tech), Deming Cheng, Wen-mei Hwu (U. Illinois), Scott Mahlke (U. Michigan), Viktor Prasanna (USC). Total Award Amount: $1,893,566 (shared among four institutions for Phase 1).

- Co-PI in NSF SPX collaborative research award for "Scalable Heterogeneous Migrating Threads for Post-Moore Computing", PI: Peter Kogge (Notre Dame), July 2018 - June 2021. Award Amount: $150,000 per year for 3 years.

- PI for DOE subcontract from Sandia National Laboratory for "Infrastructure Development for Analyzing Resilience with Asynchronous Many Task Parallelism", May 2018 – present. Total Award Amount: $50,000 in 2018 and $100,000 in 2019.

- PI for DOE subcontract from Oak Ridge National Laboratory for "CAASCADE: Understanding HPC Applications for Evidence-based Co-design", December 2017 – July 2018. Total Award Amount: $76,487.
- PI for DOE subcontract from Oak Ridge National Laboratory for "Instrumentation and Analysis of Memory Accesses in Accelerator Code", May 2018 – August 2018. Total Award Amount: $38,394.
- PI for subcontract from Brookhaven National Laboratory for DOE ECP project, "Scaling OpenMP via LLVM for Exascale Performance and Profitability (SOLLVE)". September 2016 – August 2019. Award Amount: $103,666 per year for 3 years.
- Co-PI in NSF XPS collaborative research award for "Parallel and Distributed Circuit Programming for Structured Prediction", PI: Jason Eisner (JHU), August 2016 – December 2019. Award Amount: $137,081 per year for 3 years.
- PI for DOE subcontract from Los Alamos National Laboratory for "Extending OpenSHMEM to Support an Asynchronous Execution Model", February 2016 – present. PI: Vivek Sarkar. Award Amount: $250,000/year.
- Industry contract from Shell International Exploration & Production, Inc. for "Scalable Parallel Pore-scale Multi-phase Flow Simulation", September 2015 – August 2018. PI: Vivek Sarkar. Total Award Amount: $300,000.
- DARPA MUSE award, September 2014 – November 2018, "Pliny: An End-to-End Framework for Big Code Analytics". PI: Vivek Sarkar, co-PIs: Swarat Chaudhuri, Chris Jermaine (Rice University), Isil Dillig, Thomas Dillig (UT Austin), Ben Liblit, Tom Reps (U. Wisconsin – Madison), Dave Melski (GrammaTech). Total Award Amount: $11,142,361 (shared among four institutions)
- Co-PI in NSF InTrans award for "Accelerator-Rich Architectures with Applications to Healthcare", PI: Jason Cong (UCLA), January 2014 – December 2016. Total Award Amount: $225,000
- Industry contract from Samsung Electronics Co., Ltd for "JavaScript: Ahead-of-Time Compilation with LLVM", August 2013 – March 2015. PI: Vivek Sarkar. Total Award Amount: $300,000.
- NSF award CCF-1302570, April 2013 – April 2017, "Collaborative Research: A Static and Dynamic Verification Framework for Parallel Programming", PI: Ganesh Gopalakrishnan (U. Utah), co-PIs: Eric Mercer (BYU), Vivek Sarkar (Rice). Amount awarded to Vivek Sarkar for first two years: $200,000 (additional $200,000 for Years 3 and 4 is pending).
- Department of Defense contract for "Advanced Computing Runtime", April 2013 – April 2015, PI: Vivek Sarkar. Total amount: $1,994,965 ($1,049,292 for Prof. Sarkar's group at Rice, and $945,673 for a subcontract to Intel, co-PI: Tim Mattson.)

- Industry contract from Halliburton Corporation for "Optimization of Lattice Boltzmann (LB) Simulations on Multicore CUP and Manycore GPU Processors", April 2013 – December 2013. PI: Vivek Sarkar. Total amount: $250,000.
- Industry contract from Samsung Advanced Institutes of Technology for "Rice Tizen Analysis for Security (RTAS)", October 2012 – October 2013. PI: Vivek Sarkar. Total amount: $259,999
- DOE X-stack contract DE-SC0008883 for "DEGAS : Dynamic, Exascale Global Address Space", September 2012 – August 2016. Overall PI: Kathy Yelick (LBL). PI for Rice subcontract: Vivek Sarkar. Co-PI for Rice subcontract: John M Mellor-Crummey. Award amount for Rice: $450,000/year.
- Intel DOE Traleika Glacier X-Stack subcontract, September 2012 – August 2016, PI: Vivek Sarkar. Amount awarded to Vivek Sarkar (only co-PI from Rice): $1,000,000.
- DOE X-stack contract DE-SC0008882 for "Domain Specific Language Support for Exascale (D-TEC)", September 2012 – August 2016. Overall PI: Dan Quinlan (LLNL). PI for Rice subcontract: John M Mellor-Crummey. Co-PI for Rice subcontract: Vivek Sarkar. Award amount for Rice: $575,000/year.
- Intel UHPC subcontract for "Scalable Runtimes and CnC Implementations for Extreme Scale Systems", October 2011 – June 2012, PI: Vivek Sarkar. Amount awarded to Vivek Sarkar: $180,000.
- DOE/LLNS contract for "Data Abstractions for Portable HPC Performance", October 2011 – September 2013, PI: Vivek Sarkar. Total amount: $200,000.
- DOE/LLNS contract for "Program Analyses for Rose", September 2011 – June 2013, PI: Vivek Sarkar. Total amount: $150,000 (Phase 1), $100,000 (Phase 2).
- DOE/LLNS contract for "ROSE Extensions for Java", July 2011 – September 2014, PI: Vivek Sarkar. Total amount: $150,000 (Phase 1), $165,548 (Phase 3).
- NSF award CCF-0964520, June 2010 – May 2013, "Collaborative Research: Chorus: Dynamic Isolation in Shared- Memory Parallelism", PI: Swarat Chaudhuri (PSU), co-PI: Vivek Sarkar (Rice). Amount awarded to Vivek Sarkar: $653,918.
- NSF Expeditions award CCF-0926127, September 2009 – August 2014, "Customizable Domain-Specific Computing". This research is conducted by the multi-institute Center for Domain-Specific Computing (Center Director: Jason Cong, UCLA, Center Associate Director: Vivek Sarkar, Rice.) Total amount: $9,999,997. Amount awarded to Vivek Sarkar: $1,610,000.
- MARCO award for Multiscale Systems Center, September 2009 – August 2012, "Automated Modeling and Management of Energy in Managed Runtime Systems", PI: Jan Rabaey (UC Berkeley). Amount awarded to Vivek Sarkar (Rice): $184,134.

- Intel contract, December 2009 – December 2010, "Scheduling Policies and Patterns for Intel's Concurrent Collections Parallel Programming Model", PI: Vivek Sarkar (Rice). Total amount: $75,000.
- NSF award CCF-0938018, September 2009 – August 2012, "Collaborative Research: Programming Models and Storage System for High Performance Computation with Many- Core Processors", PI: Vivek Sarkar (Rice), co-PIs: Jack Dennis (MIT), Guang Gao (U. Delaware). Amount awarded to Vivek Sarkar: $300,000.
- DARPA AACE award, AFRL Contract FA8650-09-C-7915, March 2009 – September 2013, "Platform-Aware Compilation Environment (PACE)", PI: Keith Cooper (Rice), co-PIs: John Mellor-Crummey (Rice), Krishna Palem (Rice), Vivek Sarkar (Rice), Linda Torczon (Rice). Total amount: $16,000,000.
- NSF award, CCF-0833166, September 2008 – August 2011, "Collaborative Research: Programming Models, Compilers, and Runtimes for High-End Computing on Manycore Processors", PI: Vivek Sarkar (Rice), co-PI: Guang Gao (U. Delaware). Amount awarded to Vivek Sarkar: $405,999.
- AMD contract, September 2007 – September 2008, "Integration and Evaluation of Lightweight Profiling (LWP) in Java Runtime Environments for Multicore Processors", PI: Vivek Sarkar (Rice). Total amount: $125,000.
- DARPA HPCS program, IBM PERCS project, 2003 – 2010. Contributed to IBM's proposals being selected for funding in Phase 1 ($3M), Phase 2 ($53M) and Phase 3 ($244M) of the DARPA HPCS program, as one of six key personnel in the proposals. *This was the only external funding that I applied for while at IBM during 1987–2007.*

6 Industry Gifts

- IBM CAS Fellowship for research on "XL Compiler and Runtime Extensions for GPU Platforms (continuation)" by PhD student, Prithayan Barua and research scientist, Jun Shirako, 2018, $22,000.
- IBM CAS Fellowship for research on "Extending OMR with Automatic CPU/GPU Parallelization (continuation)" by PhD student Ankush Mandal and research scientist, Akihiro Hayashi, 2018, $22,000.
- IBM CAS Fellowship for research on "XL Compiler and Runtime Extensions for GPU Platforms (continuation)" by PhD student, Prasanth Chatarasi and research scientist, Jun Shirako, 2017, $28,000.
- IBM CAS Fellowship for research on "Extending OMR with Automatic CPU/GPU Parallelization" by PhD student Max Grossman and research scientist, Akihiro Hayashi, 2017, $28,000.

- IBM CAS Fellowship for research on "XL Compiler and Runtime Extensions for GPU Platforms (continuation)" by PhD student, Prasanth Chatarasi and research scientist, Jun Shirako, 2016, $28,000.
- IBM CAS Fellowship for research on "GPU Enablement of Java Applications (continuation)" by PhD student Max Grossman and research scientist, Akihiro Hayashi, 2016, $28,000.
- IBM CAS Fellowship for research on "XL Compiler and Runtime Extensions for GPU Platforms (continuation)" by PhD student, Deepak Majeti and research scientist, Jun Shirako, 2015, $28,000.
- IBM CAS Fellowship for research on "GPU Enablement of Java Applications (continuation)" by Postdoctoral Researcher, Akihiro Hayashi, 2015, $28,000.
- IBM CAS Fellowship for research on "XL Compiler and Runtime Extensions for GPU Platforms" by PhD student, Deepak Majeti and research scientist, Jun Shirako, 2014, $28,000.
- IBM CAS Fellowship for research on "GPU Enablement of Java Applications" by Postdoctoral Researcher, Akihiro Hayashi, 2014, $28,000.
- IBM CAS Fellowship for research on "OpenMP and PGAS Enhancements for Manycore Processor (continuation)" by PhD student, Sanjay Chatterjee, 2013, $28,000.
- IBM CAS Fellowship for research on "OpenMP and PGAS Enhancements for Manycore Processor (continuation)" by PhD student, Sanjay Chatterjee, 2012, $28,000.
- IBM CAS Fellowship for research on "OpenMP and PGAS Enhancements for Manycore Processors" by PhD student, Sanjay Chatterjee, 2011, $28,000.
- Intel grant, December 2010, for research on "CnC-Python", $75,000.
- Co-PI for IBM SUR donation of 18-node POWER7 BlueBioU system at Rice, valued at $7.6M (March 2010).
- IBM X10 Innovation grant, March 2010, $20,000.
- IBM CAS Fellowship for research on "Scalable and Efficient Scheduling of OpenMP Tasks (continuation)" by PhD student, Yi Guo, 2010, $28,000.
- IBM Open Collaboration Research grant for Multicore Software, January 2009, $200,000.
- IBM CAS Fellowship for research on "Scalable and Efficient Scheduling of OpenMP Tasks" by PhD student, Yi Guo, 2009, $28,000.
- BHP Billiton gift for Numerical Optimization of Java Codes in Seismic Applications, June 2008, $100,000.
- IBM Open Collaboration Research grant for Multicore Software, January 2008, $200,000.

7 Graduate Student and Other Mentoring

I consider the mentoring of students and junior colleagues to be among the most satisfying experiences of my career. I currently supervise eleven PhD students at Georgia Tech since 2017, after supervising sixteen PhD students at Rice University during 2007-16. (All graduate student mentoring prior to 2007 was performed in a co-advisor role when I was at IBM Research.) I regularly seek out opportunities for my graduated PhD students and post-doctoral researchers to serve on program committees for major conferences and workshops, and to advise them on career transitions. While it remains a major challenge, I always make it a high priority to increase diversity in my research group. Members of my group at Rice originated from more than ten different countries, and the 16 PhD students that I supervised at Rice during 2007-2017 included one African-American student (Mack Joyner) and one female student (Alina Sbirlea). I currently have one female student in my group at Georgia Tech (Sana Damani).

I also take special pride in my record of mentorship of junior researchers in both industry and academia. At Georgia Tech, in my role as SAC chair, I have focused on meeting regularly with junior faculty to mentor them on different aspects of their career, and to reassure them during periods of uncertainty in our school and college, such as the recent transition of both our school chair and college dean. I also had the privilege of recruiting and mentoring a large number of researchers during my career at IBM. Many of my mentees from that period have now become world-class leaders in the Programming Languages research community; notable examples include Evelyn Duesterwald (appointed as a first-line manager in my department in 2006, became an ACM Distinguished Scientist in 2010), Stephen Fink (transferred to my department in 1999, appointed as team lead for WALA project in 2002, became ACM Senior Member in 2008, and ACM Distinguished Scientist in 2011), David Grove (recruited by me in 1998 after his PhD, appointed as team lead for Jikes RVM after the open source release in 2001, became ACM Senior Member in 2006, ACM Distinguished Scientist in 2010, and ACM Fellow in 2012), Michael Hind (transferred to my department in 1998, became my successor as first-line manager in 2000, and second-line manager in 2007, and an ACM Distinguished Scientist in 2009), Frank Tip (transferred to my department in 2000, appointed as first-line manager in 2003, became ACM Distinguished Scientist in 2008), and Eran Yahav (recruited by me in 2004 after his PhD, joined Technion as a tenure-track faculty member in 2010). While at IBM, I would meet with my mentees regularly (often on a biweekly basis) to discuss all aspects of their career at IBM. Even after I left IBM in 2007, I remained involved in assisting with their award nominations and advising them on their future career transitions.

The graduate students whom I have mentored include the following:

- Supervised Sriraj Paul (Rice PhD, 2018) on his PhD research on mapping high level parallel programming models to asynchronous many-task runtime systems.
- Supervised Kumud Bhandari (Rice PhD, 2018) on his PhD research on programming techniques for systems with non-volatile byte-addressable random access memory (NVRAM).

- Supervised Nick Vrvilo (Rice PhD, 2017) on his PhD research on data and task abstractions for extreme-scale runtime systems.
- Supervised Rishi Surendran (Rice PhD, 2017) on his PhD research on debugging, repair, and synthesis of task-parallel programs.
- Supervised Max Grossman (Rice PhD, 2017) on his PhD research on programming systems for heterogeneous supercomputers.
- Supervised Alina Sbirlea (Rice PhD, 2015) on her PhD research related to high-level execution models for multicore architectures.
- Supervised Sagnak Tasirlar (Rice PhD, 2015) on his PhD research related to optimized event-driven runtime systems.
- Supervised Shams Imam (Rice PhD, 2015) on his PhD research related to cooperative scheduling of parallel tasks with synchronization constraints.
- Supervised Drago Sbrlea (PhD, 2015) on his PhD research related to memory and communication optimizations for macro-dataflow programs.
- Supervised Deepak Majeti (PhD, 2015) on his PhD research related to portable programming models for heterogeneous platforms.
- Supervised Kamal Sharma (Rice PhD, 2014) on his PhD research related to locality transformations of computation and data.
- Supervised Sanjay Chatterjee (Rice PhD, 2013) on his PhD research related to runtime systems for extreme scale platforms.
- Supervised Raghavan Raman (Rice PhD, 2012) on his PhD research related to dynamic data race detection in parallel programs.
- Supervised Yi Guo (Rice PhD, 2010) on his PhD research related to runtime systems for multicore processors.
- Supervised Rajkishore Barik (Rice PhD, 2009) on his PhD research related to register allocation and analysis and optimization of explicitly parallel programs.
- Supervised Mackale Joyner (Rice PhD, 2008) on his PhD research related to optimization of array accesses in high-productivity languages (Chapel, Fortress, X10).
- Co-supervised Igor Peshansky (NYU PhD, 2003) on his PhD research on optimistic program optimization.
- Co-supervised Massimiliano Poletto (MIT PhD, 1999) on his PhD research related to Linear Scan register allocation.

- Co-supervised Kourosh Gharachorloo (Stanford PhD, 1995) on his graduate research related to implementation of Sisal programs, prior to his PhD research.

- In addition to the above, I co-supervised and mentored the following graduate students while they were at IBM, and have co-authored papers with each of them: Radhika Thekkath (U.Washington PhD, 1995), Ras Bodik (U.Pittsburgh PhD, 1999), Chandra Krintz (UCSB PhD, 2001), Jan-Willem Maessen (MIT PhD, 2002), Matthew Arnold (Rutgers PhD, 2002), Keunwoo Lee (U.Washington PhD, 2006), Alexey Loginov (U.Wisconsin PhD, 2006), John Whaley (Stanford PhD, 2006), Matthew Harren (PhD, UC Berkeley, 2007), Vincent Cave (MS, INRIA, 2007), Kartik Agaram (PhD, UT Austin, 2008), Yuan Zhang (PhD, U.Delaware, 2008), Shivali Agarwal (PhD,TIFR, India, 2009).

- I also collaborated with the following graduate students while I was on sabbatical at MIT during 1996–1998: Rajeev Barua, Jacob Harris, Kathleen Knobe, Walter Lee, Devabhaktuni Srikrishna, Michael Taylor, Eliot Waingold.

8 Professional Service

I believe that service is leadership and leadership is service. Over the years, I have found myself assuming an increasing number of leadership roles across the broad CS community. While these include multiple Program Chair and General Chair roles for major CS conferences, the roles that I have found especially rewarding are those where I have had to work on leading and motivating academic colleagues towards a common vision and set of goals. One example of this was as co-PI of the DARPA HPCS project at IBM, where my role in Phase 2 (2003-06) was akin to that of a program manager and included funding and coordinating work performed by six university partners towards a common vision of improving productivity in developing high performance applications. It was satisfying to see groundbreaking research emerge from this initiative, including the early sketch-based work on program synthesis led by Ras Bodik and Armando Solar-Lezama at UC Berkeley, as well as the StreamIt domain-specific programming environment led by Saman Amarasinghe's research group at MIT.

After moving to academia, I was asked by DARPA to lead a new study on Exascale Software during 2008-2009. I formed a team of 24 world-class leaders across multiple areas of CS in academia, government and industry. We produced a report in 2009 that was considered to be highly influential by program managers in DARPA, DOE, and NSF. I have also served as a member of DOE's ASCAC advisory committee since 2009, and have chaired subcommittees related to future directions for high performance computing. Finally, by serving on CRA's board since 2015, I have gained a deeper appreciation of the challenges facing CS departments in North America.

A common theme in my CS community leadership is a collaborative approach, that includes gaining the trust of my peers that I will work towards consensus but also use my best judgment to make timely decisions when a consensus is not emerging.

8.1 Conference Committees

Program Chair and General Chair roles

- General Co-chair, IEEE International Conference on Rebooting Computing (ICRC), November 2019.
- Workshop Co-chair, 32nd Workshop on Languages and Compilers for Parallel Computing (LCPC), October 2019.
- General Chair, ACM Federated Computing Research Conference (FCRC), June 2019.
- Co-chair, 2018 CRA Conference at Snowbird, July 2018.
- Program Co-chair, Twenty-third International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS), March 2018.
- General Chair, 2017 ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming (PPoPP), February 2017.
- Workshop Co-chair, Second Workshop on Runtime Systems for Extreme Scale Programming Models and Architectures (RESPA), co-located with SC16, November 2016.
- Program Co-chair, 2016 Virtual Execution Environments (VEE) conference, co-located with ASPLOS 2016, April 2016.
- Workshop Co-chair, First Workshop on Runtime Systems for Extreme Scale Programming Models and Architectures (RESPA), co-located with SC15, November 2015.
- Plenary Speaker Chair, ACM Federated Computing Research Conference (FCRC), June 2015.
- Program Co-chair, 29th International Conference on Supercomputing (ICS), June 2015.
- Program Vice-chair (Software), 29th IEEE International Parallel and Distributed Processing Symposium (IPDPS), May, 2015.
- General Chair, Eighteenth International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS), March 2013.
- Program Co-chair, Systems Software area, SC11, November 2011.
- Program Chair, PACT 2011: Twentieth International Conference on Parallel Architectures and Compilation Techniques (PACT), October 2011.
- General Co-chair, Partitioned Global Address Space (PGAS) conference, October 2011.
- Program Co-chair, Fourth workshop on Programming Language Approaches to Concurrency and Communication-cEntric Software (PLACES), April 2011.

- General Co-chair, 7th International Conference on Distributed Computing and Internet Technologies (ICDCIT), February 2011.
- Workshop Co-chair, 23rd Workshop on Languages and Compilers for Parallel Computing (LCPC), October 2010.
- Vice Chair, Programming Models, Languages and Compilers track, International Conference on Parallel Processing (ICPP), September 2009.
- Program Chair, 2009 ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming (PPoPP), February 2009.
- Conference Co-chair, IBM Academy Conference on Software Engineering for Tomorrow (SWEFT), October 2006.
- General Chair, ACM SIGPLAN '05 Conference on Programming Language Design and Implementation (PLDI), June 2005.
- Program Co-chair, PACT 2003: Twelfth International Conference on Parallel Architectures and Compilation Techniques, September 2003.
- General Chair, PLAN-X: First Workshop on Programming Language Technologies for XML, October 2002 (co-located with ACM PLI 2002).
- Conference Chair, IBM Academy Conf. on Best Practices in Multi-Site Software Development, May 2002.
- Program Chair, ACM SIGPLAN '94 Conference on Programming Language Design and Implementation (PLDI), June 1994.

Program Committees

- Program Committee member, IEEE International Parallel and Distributed Processing Symposium (IPDPS), May 2019.
- Program Committee member, IEEE International Conference on Rebooting Computing (ICRC), November 2018.
- Program Committee member, 31st Workshop on Languages and Compilers for Parallel Computing (LCPC), October 2018.
- External Program Committee member, ACM SIGPLAN '17 Conference on Programming Language Design and Implementation (PLDI), June 2017.
- Program Committee Member, IA^3 Workshop on Irregular Applications: Architectures and Algorithms, co-located with SC14, November 2014.

- Program Committee Member, International Conference on Principles and Practice of Programming on the Java platform: virtual machines, languages, and tools (PPPJ), September 2014.
- External Review Committee member, OOPSLA conference, October 2012.
- Program Committee member, Third International Conference on Runtime Verification (RV), September 2012.
- Program Committee member, 25th Workshop on Languages and Compilers for Parallel Computing (LCPC), September 2012.
- Program Committee member, 2012 IEEE International Parallel & Distributed Processing Symposium (Software Track), May 2012.
- Program Committee member, CC 2012: Twenty-first International Conference on Compiler Construction, March 2012.
- Program Committee member, IEEE International Parallel and Distributed Processing Symposium (IPDPS), April 2011.
- Program Committee member, 2nd Workshop on Determinism and Correctness in Parallel Programming (WoDet), March 2011.
- External Review Committee member, Sixteenth International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS 2011), March 2011.
- External Review Committee member, 2011 ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming (PPoPP), February 2011.
- Program Committee member, PACT 2010: Nineteenth International Conference on Parallel Architectures and Compilation Techniques (PACT), September 2010.
- Program Committee member, TOOLS EUROPE 2010: International Conference on Objects, Models, Components, Patterns, June 2010.
- Program Committee member, ACM SIGPLAN '10 Conference on Programming Language Design and Implementation (PLDI), June 2010.
- Program Committee member, 2nd USENIX Workshop on Hot Topics in Parallelism (HotPar '10), June 2010.
- Program Committee member, ACM International Conference on Computing Frontiers (CF), May 2010.
- Program Committee member, 2010 International Symposium on Code Generation and Optimization (CGO), April 2010.

- Program Committee member, CC 2010: Nineteenth International Conference on Compiler Construction, March 2010.
- Program Committee member,3rd workshop on Programming Language Approaches to Concurrency and communication-cEntric Software (PLACES 2010)
- Program Committee member, 6th International Conference on Distributed Computing and Internet Technologies (ICDCIT), February 2010.
- Program Committee member, Supercomputing 2009 (SC09), November 2009.
- Program Committee member, 22nd Workshop on Languages and Compilers for Parallel Computing (LCPC), October 2009.
- Program Committee Member, Second International Workshop on Parallel Programming Models and Systems Software for High-End Computing (P2S2), co-located with ICPP 2009, September 2009.
- Program Committee Member, IFIP Working Conference on Domain Specific Languages (DSL WC), July 2009.
- External Review Committee member, ACM SIGPLAN '09 Conference on Programming Language Design and Implementation (PLDI), June 2009.
- Program Committee Member, ACM SIGOPS Operating System Review Special Issue on the Interaction among the OS, Compilers, and Multicore Processors, April 2009.
- Program Committee Member, DAMP 2009 Workshop (Declarative Aspects of Multicore Programming), co-located with POPL, January 2009.
- Member of Editorial Board, Encyclopedia of Parallel Computing, Springer, 2009.
- Program Committee Member, First International Workshop on Parallel Programming Models and Systems Software for High-End Computing (P2S2), co-located with ICPP 2008, September 2008.
- Program Committee Member, 20th ACM Symposium on Parallelism in Algorithms and Architectures (SPAA), June 2008.
- Program Committee Member, 2008 Workshop on Software and Hardware Challenges of Manycore Platforms (SHCMP'08), co-located with ISCA'08, June 2008.
- Program Committee Member, Workshop on Software Tools for Manycore Systems (STMS'08), June 2008.
- Program Committee Member, International Workshop on Multi-Core Computing System (MuCoCoS'08), March 2008.

- Program Committee Member, Workshop on Parallel Programming on Accelerator-Based Systems (PPABS), co-located with PPoPP 2008, February 2008.
- Program Committee Member, ICDCIT 2007: Fourth International Conference on Distributed Computing and Internet Technology, December 2007.
- Program Committee Member, PACT 2007: Sixteenth International Conference on Parallel Architectures and Compilation Techniques, September 2007.
- Program Committee Member, Workshop on Programming Models for Ubiquitous Parallelism (PMUP), September 2006 (co-located with PACT 2006).
- Program Committee Member, 2006 International Conference on High Performance Computing and Communications (HPCC), September 2006.
- Program Committee Member, IBM Academy Conference on Software Engineering for Tomorrow (SWEFT), November 2005.
- Guest Editor, IBM Systems Journal special issue on Open Source Software, Volume 44, Number 2, June 2005.
- Program Committee Member, ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming 2005 (PPoPP), June 2005.
- Program Committee Member, PODC Workshop on Concurrency and Synchronization in Java Programs, July 2004.
- Program Committee Member, CC 2003: Twelfth International Conference on Compiler Construction, April 2003.
- Program Committee Member, ACM International Conference on Supercomputing, June 2001.
- Program Committee Member, ACM Java Grande Conference, June 2000.
- Program Committee Member, Second Workshop on Java for High-Performance Computing, June 2000.
- Program Committee Member, International Conference on Parallel Processing (ICPP), August 1996.
- Program Committee Member, 28th Annual IEEE/ACM International Symposium on Microarchitecture (MICRO), November 1995.
- Program Committee Member, PACT 1994: Parallel Architectures and Compilation Techniques, August 1994.
- Program Committee Member, Twenty-first ACM Symposium on Principles of Programming Languages (POPL), January 1994.

- Program Committee Member, ACM SIGPLAN '93 Conference on Programming Language Design and Implementation (PLDI), June 1993.
- Program Committee Member, ACM SIGPLAN '91 Conference on Programming Language Design and Implementation (PLDI), June 1991.

8.2 Advisory/Award/Review/Steering Committees

2019: External Review Committee Member, Department of Computer Science, UVA

I served on the 2019 external review committee for the Department of Computer Science at U.Virginia with Sandhya Dwarkadas (U.Rochester), Kevin Jeffay (UNC-Chapel Hill), Don Fussell (UT Austin), and Xiaodong Zhang (OSU). Our committee visited U.Virginia for a day and a half in May 2019, during which we had multiple meetings with department faculty, students and staff, and also met with the department chair, Kevin Skadron, and the dean of UVA Engineering, Craig Benson. Though no one was formally designated to be the chair of the committee, I took the lead on coordinating our out-brief presentation at the end of the site visit in May, and the creation of our 12-page report which was submitted in July.

2019, 2015, 2013: CRD and CS Area Review Committee Member, Lawrence Berkeley National Laboratory

I have served as a committee member on three strategy-related reviews for the Computing Sciences area and Computational Research Division at LBNL — a review of the CS area strategy and CRD capabilities in February 2019, a review of the strategic plan of the CS area in September 2015, and a review of CRD in January 2013. In addition to contributing to the review reports, I engaged in mentoring discussions with postdocs and junior research staff to both encourage them in their career directions at the lab and to take back their concerns to share at ASCAC meetings.

2015-present: Member, CRA Board of Directors

CRA's mission is to strengthen research and advanced education in computing, and includes membership of Computer Science departments in the USA and Canada. By serving on CRA's board since 2015, I have gained a deeper appreciation of the challenges facing CS departments in North America. My most significant CRA activity thus far has been to serve as co-chair of the committee for the highly successful CRA Snowbird 2018 conference; the other committee members include Lorenzo Aviso, Carla Brodley, Kim Hazelwood (co-chair), Chris Johnson, Mario Nascimento, and Jaime Teevan. I am proud of what we achieved as a team in building the conference program. Notable highlights include selection of "Diversity of Leadership" and "Ubiquity of AI" as major conference themes, in addition to focusing on specific topics that are most relevant to CS departments, such as Faculty Growth, Department Rankings, Faculty Recruiting, Schools/Colleges of Computing, Industrial Research in CS, and Development of Teaching Faculty.

2015-2017: Member, ACM/IEEE CS Ken Kennedy Award Selection Committee

The ACM/IEEE CS Ken Kennedy Award is given annually for outstanding contributions to programmability or productivity in high-performance computing together with significant community service or mentoring contributions. It was established in memory of Ken Kennedy, the founder of Rice University's computer science department and one of the world's foremost experts on high-performance computing. It was my privilege to have Ken as a mentor, and to serve on this committee for three years. In 2016, I served as chair of this committee, and as a voting member of the IEEE CS Awards Committee.

2015-present: Member, ORNL Computing and Computational Sciences Directorate Advisory Committee

The advisory committee for Oak Ridge National Lab's Computing and Computational Sciences Directorate (CCSD) is responsible for reviewing current and future activities in the directorate, with respect to different focus areas of strategic importance each year.

2015-2016: Member, UCSB Computer Engineering Advisory Board

The UCSB Computer Engineering Advisory Board (CEAB) advises the ECE and CS departments at UCSB on future directions for their Computer Engineering Program.

2013-2015: Member, ASPLOS Steering Committee

After serving as General Chair of the ASPLOS 2013 conference (Architectural Support for Programming Languages and Operating Systems) in Houston, I became a member of the ASPLOS Steering Committee, and chaired the committee in 2015.

2009-present: Member, Advanced Scientific Computing Advisory Committee (ASCAC), US Department of Energy

The Advanced Scientific Computing Advisory Committee (ASCAC) provides advice to the Department of Energy on a variety of scientific and technical issues related to its Advanced Scientific Computing Research program. ASCAC's recommendations include advice on long-range plans, priorities, and strategies to address more effectively the scientific aspects of advanced scientific computing including the relationship of advanced scientific computing to other scientific disciplines, and maintaining appropriate balance among elements of the program. The Committee, which formally reports to the Director, Office of Science, primarily includes representatives of universities, national laboratories, and industries involved in advanced computing research. I have participated in multiple ASCAC subcommittees including one on exascale computing in 2010. I chaired an ASCAC subcommittee formed in 2012 in response to a charge by Dr. William F. Brinkman, Director of the Office of Science, Department of Energy, asking ASCAC to examine the potential synergies between

the challenges of data-intensive science and exascale. I also chaired a subcommittee formed in 2017 to examine future technologies for high performance computing in the post-Moore era.

2005-2007: Consultant, US Army Science Board

Helped review the US Army's future technology programs, and participated in briefings in Fort Leavenworth on new challenges faced by terrorist tactics. Hosted a meeting at IBM related to software modernization in future Army initiatives.

9 Keynote Talks, Invited Talks, Panels (Selected)

- Keynote speaker, IEEE International Conference on High Performance Computing, Data and Analytics (HiPC), "Data Flow Execution Models – a Third Opinion", December 2019.
- Invited speaker in U.Illinois Computer Science Distinguished Lecture Series, "Data Flow Execution Models – a Third Opinion", October, 2019
- Keynote speaker, 11th Workshop on Programming Language Approaches to Concurrency and Communication-cEntric Software (PLACES), co-located with ETAPS, April 2019
- Panelist, "Exascale Runtime Systems", SC18, Dallas, Texas, November, 2018.
- Invited speaker, 31st International Workshop on Languages and Compilers for Parallel Computing (LCPC), " Programming Model and Compiler Extensions for Unifying Asynchronous Tasks, Futures, and Events", October 2018.
- Invited speaker, University of Vienna, Austria, "Software Challenges for Extreme Heterogeneity", August 2018.
- Invited speaker, TU Vienna, Austria, "Why the End-Game for Moore's Law will be driven by a Compiler Renaissance", August 2018.
- Invited speaker, Indian Institute of Science, Bangalore, India, "Software Challenges for Extreme Heterogeneity", March 2018.
- Invited speaker, Indian Institute of Technology, Bombay, India, "Software Challenges for Extreme Heterogeneity", March 2018.
- Keynote speaker, Compiler, Architecture and Tools Conference (CATC), Haifa, Israel, "Software Challenges for Extreme Heterogeneity", December 2017.
- Invited speaker, Technion, Haifa, Israel, "Unstructured Parallelism Considered Harmful – Improving Software Reliability and Performance through Structured Parallelism", December 2017.

- Keynote speaker, Workshop on Memory Centric Programming for HPC (MCHPC'17), co-located with SC17, "Compiler and Runtime Challenges for Memory Centric Programming", November 2017.
- Keynote speaker, 30th International Workshop on Languages and Compilers for Parallel Computing (LCPC), "Software Challenges for Extreme Heterogeneity", October 2017.
- Invited speaker, Second Workshop of Mission-Critical Big Data Analytics (MCBDA 2017), "Programming Model Challenges for Extreme Scale Computing and Analytics", May 2017.
- Keynote speaker, 26th International Conference on Compiler Construction (CC), "Why the End-Game for Moore's Law will be driven by a Compiler Renaissance", February 2017.
- Computer Engineering Seminar Series, University of Southern California (USC), "Software Challenges for Extreme Scale Systems, or how to play the End-Game for Moore's Law", October 2016.
- Distinguished Speaker Colloquium Series, U. Virginia, Department of Computer Science, "Software Challenges for Extreme Scale Systems, or how to play the End-Game for Moore's Law", September 2016.
- Keynote speaker, Partitioned Global Address Space (PGAS) conference, "The Role of Global Address/Name Spaces in Extreme Scale Computing and Analytics", September 2015.
- Invited speaker, ORNL Smoky Mountains Computational Sciences and Engineering Conference, "Programming Model Challenges for Extreme Scale Computing and Analytics", September 2015.
- Invited speaker, INRIA Paris, "Structured Parallel Programming Primitives and their use in Compilers, Runtimes and Debuggers for Parallel Systems", June 2015.
- Keynote speaker, International Conference on Principles and Practice of Programming on the Java platform: virtual machines, languages, and tools (PPPJ), "Beyond Multicores: Parallel Computing on GPUs and clusters with Java 8", September 2014.
- Distinguished lecturer, "Hybrid Programming Challenges for Extreme Scale Software", Department of Computer Science and Engineering, Texas A&M University, September 2014.
- Invited speaker, "Test-Driven Detection and Repair of Data Races in Structured Parallel Programs", Indian Institute of Science (IISc), Bangalore, India, August 2014.
- Invited speaker, "Test-Driven Detection and Repair of Data Races in Structured Parallel Programs", Tata Institute of Fundamental Research (TIFR), Mumbai, India, August 2014.
- Invited speaker, "Leveraging Structured Parallelism for Improved Programmability and Scalability", U. Alberta, July 2014.

- Invited speaker, "Software Challenges for Extreme Scale Systems", Princeton University, May 2014.
- Keynote speaker, "Hybrid Programming Challenges for Extreme Scale Software", 12th Annual Workshop on Charm++ and Application, University of Illinois at Urbana-Champaign, April 2014.
- Invited speaker, "Why Runtime Systems and System Software Must Work Together for Future HPC Applications", 2014 DOE Salishan Conference on High Speed Computing, Gleneden Beach, Oregon, April 2014.
- Invited speaker, "Software Challenges for Extreme Scale Systems", Intel Labs, Santa Clara, April 2014.
- Keynote speaker, "Hybrid Programming Challenges for Extreme Scale Software", First OpenSHMEM Workshop: Experiences, Implementations and Tools, Annapolis, Maryland, March 2014.
- Panelist, "Exascale Runtime Systems", SC13, Denver, Colorado, November, 2013.
- Keynote speaker, "The Role of OpenMP in Extreme Scale Software", International Workshop on OpenMP (IWOMP) 2013, Canberra, Australia, September 2013.
- Invited lecture, "Analysis and Transformation of Programs with Explicit Parallelism", Seoul National University, August 2013.
- Invited colloquium speaker, "Software Challenges for Extreme Scale Systems", U. Utah School of Computing, July 2013.
- Invited lecture, "Analysis and Transformation of Programs with Explicit Parallelism", ENS Lyon, June 2013.
- Invited Lecture, "Synergistic Challenges in Data-Intensive Science and Extreme Scale Computing", ISI Kolkata and TIFR Mumbai, May 2013.
- Organizer and presenter, SC12 Birds of a Feather Session on "Open Community Runtime (OCR)", November 2012.
- Invited colloquium speaker, Harvey Mudd College, "Determinacy and Data Races in Task-Parallel Programs", September 2012.
- Distinguished Lecture, "Programming Challenges for Multicore Processors", UC Riverside, February 2012.
- Invited lecture, "Towards a Portable Execution Model for Extreme Scale Multicore Systems", CSIRO, Perth, Australia, December 2011.
- Panelist, PGAS 2011 conference, October 2011.

- Panelist, DFM 2011 workshop, October 2011.
- Keynote speaker, APPT Conference, "Software Challenges for Extreme Scale Systems", Shanghai, China, September 2011.
- Panelist, Supercomputing 2010 (SC10), "Advanced HPC Execution Models: Innovation or Disruption", November 2010.
- Panelist, IEEE International Parallel and Distributed Processing Symposium (IPDPS) 2010, "Unconventional Wisdom in Multicore Computing", April 2010.
- Invited speaker, ACM SIGPLAN/SIGBED Conference on Languages, Compilers and Tools for Embedded Systems (LCTES), "Towards a Unified Execution Model for Mainstream and Embedded Multicore Systems", April 2010.
- Invited speaker, 4th Joint Symposium on Radiotherapy Research, "High Performance Computing in Biomedical Research", April 2010.
- UT San Antonio Distinguished Lecture Series in Computer Science, "Multicore Programming Models and their Implementation Challenges", March 2010.
- U. Michigan Distinguished Lecture Series in Computer Science and Engineering, "Multicore Programming Models and their Implementation Challenges", January 2010.
- Panelist, HPCA 2010 and PPoPP 2010 conferences, "Extreme Scale Computing: Challenges and Opportunities", January 2010.
- Panelist, Disruptive Technologies Panel, Supercomputing 2009 (SC09) conference, "Software Challenges in Extreme Scale Systems", November 2009.
- Invited speaker, 3rd workshop on Virtual Machines and Intermediate Languages (VMIL) co-located with OOPSLA 2009, "Virtual Machine and Intermediate Language Challenges for Parallelism", October 2009.
- Keynote speaker, 2009 European Joint Conferences on Theory and Practice of Software (ETAPS), "Challenges in Code Optimization of Parallel Programs", March 2009.
- ISTec Distinguished Lecturer, Colorado State University, "Multicore Programming Models and their Implementation Challenges" and "A Quick Tour of Modern Multicore Programming Models", December 2008.
- Panelist, "Can Developing Applications for Massively Parallel Systems with Heterogeneous Processors Be Made Easy(er)?", Supercomputing 2008.
- Invited talk, Supercomputing 2008 Workshop on Bridging Multicore's Programmability Gap, "Multicore Programming Models and their Implementation Challenges".

- Keynote speaker, U. Washington and Microsoft Research 2008 Summer Institute on The Concurrency Challenge, "Multicore Programming Models and their Implementation Challenges", August 2008.
- Invited speaker, IBM Austin Research Lab Distinguished Seminar Series, "Multicore Programming Models and their Implementation Challenges", July 2008.
- Invited speaker, GCOE Ambient SoC Symposium, Waseda University, Japan, "Programming Challenges for Multicore Parallel Systems", July 2008.
- Invited talk, 2008 Workshop on Exploiting Concurrency Efficiently and Correctly (co-located with CAV 2008), "Static and Dynamic Analysis of Parallel Programs", July 2008.
- Invited colloquium speaker, U. Maryland ECE department, "Programming Challenges for Multicore Parallel Systems", May 2008.
- Keynote speaker, 2008 International Symposium on Code Generation and Optimization (CGO)", "Code Optimization of Parallel Programs", April 2008.
- Panelist, IEEE International Parallel and Distributed Processing Symposium (IPDPS), "How to avoid making the same Mistakes all over again — what the parallel-processing Community has (failed) to offer the multi/many-core Generation?", April 2008.
- Invited colloquium speaker, UT Austin CS Department, "Programming Challenges for Multicore Parallel Systems", April 2008.
- Invited talk, Workshop on Architectures and Compilers for Multithreading, "Compiler Challenges for Multicore Parallel Systems", December 2007, IIT Kanpur, India.
- Keynote speaker, High Performance Computation Conference (HPCC), "Programming Challenges for Petascale and Multicore Parallel Systems", September 2007.
- Plenary speaker, Third International Conference on Distributed Computing and Internet Technology, "The Role of Programming Languages in Future Data-Centric and Net-Centric Applications", December 2006.
- Panelist, "Wish List: Architectural Support and Tool Infrastructure for Improving Software Dependability", Workshop on Architectural and System Support for Improving Software Dependability (ASID) co-located with ASPLOS, October 2006.
- Panelist, "Opportunities and Challenges in Partitioned Global Address Space Languages", Los Alamos Computer Science Institute (LACSI) Symposium, October 2006.
- Keynote speaker, Workshop on Software Challenges for Multicore Architectures, "X10: A High-Productivity Approach to Programming Multi-Core Systems", September 2006.

- Panelist,"High Productivity Languages for HPC: Compiler Challenges", LCPC 2005 workshop, October 2005.
- Invited speaker, Ninth Annual Workshop on High Performance Embedded Computing (HPEC), "X10 Programming: Towards High Productivity High Performance Systems in the post-Moore's Law Era", September 2005.
- Panelist, panel discussion on "Will Software Save Moore's Law?", HPEC 2005 workshop, September 2005.
- Invited speaker, AHPCRC DARPA PGAS Programming Model Conference, "X10: An Object-Oriented Approach to PGAS Programming", September 2005.
- Invited speaker, Workshop in programming models for HPCS ultra-scale applications (PMUA 2005), "X10 — a New Programming Model for Productive Scalable Parallel Programming", June 2005.
- Invited speaker, Seventh Workshop on Languages, Compilers, and Run-time Support for Scalable Systems (LCR), "X10: Addressing Language, Compiler, and Runtime Challenges for Scalable Systems in 2010", October 2004.
- Invited speaker, 2nd International Conference on the Principles and Practice of Programming in Java (PPPJ), "Java Innovation in Industry and Academia: Current Synergy and Future Challenges", June 2003.
- Keynote speaker, High Performance Computing (HiPC) conference, "Scalable High-Performance Java Virtual Machines", December 2001.
- Invited speaker, Workshop on Cutting Edge Computing, "End-to-end Adaptive Optimization: Towards Autonomic Virtual Machines", December 2001.
- Keynote speaker, 2nd Workshop on Java for High-Performance Computing (held in conjunction with ACM ICS 2000 conference), "The Evolution of Optimization and Parallelization technologies for Java, or why Java for High-Performance Computing is not an oxymoron", June 2000.
- Panelist, panel discussion on Dynamic Compilation at ACM Dynamo 2000 workshop held in conjunction with ACM POPL 2000 conference, Jan 2000.
- Invited speaker, ISCOPE conference, "Evolution of Optimization and Parallelization technologies from Fortran to Java — why High-Performance Object-Oriented Computing need not be an oxymoron", December 1999.
- Invited speaker, Workshop on Scheduling Algorithms for Parallel/Distributed Computing, "Static Scheduling with Communication Weights — Theory and Practice", June 1999.

- Panelist, Workshop on Challenges for Parallel Processing (co-located with ICPP), "Issues in Distributed Memory Systems", August 1995.
- Invited speaker, IXth Distinguished Lecture Series, University Video Communications, "Don't Waste Those Cycles: An In-Depth Look at Scheduling Instructions in Basic Blocks and Loops", August 1994 (with Barbara Simons).
- Lecturer, ACM Lectureship Series, "Program Optimization — a Quantitative Approach", "Compiling for Parallelism", and "A General Framework for Iteration-Reordering Loop Transformations", 1992-93.
- Panelist, Third Workshop on Compilers for Parallel Computers, "How good are parallelizing compilers in practice?", July 1992.
- Panelist, International Workshop on Multithreaded Computers (co-located with Supercomputing '91 conference), "Programming, Compilation, and Resource Management Issues for Multithreading", November 1991.
- Invited speaker, ACM International Conference on Supercomputing, "Compiling for Parallel Computers", June 1989.
- Panelist, ACM International Conference on Supercomputing, "Delivering Supercomputer Performance to the User", June 1989.
- Keynote speaker, Fourth Parallel Processing Circus, "The Programming Problem for General-purpose Multiprocessors", December 1988.
- Panelist, International Conference on Supercomputing, "Future of Parallel Programming", May 1988.

10 Teaching

My experiences in teaching and curriculum development include teaching three different courses at Georgia Tech (CS 4240, CS 6245, CS 6390), four different courses at Rice University during 2007-2017 (COMP 211, COMP 322, COMP 422, COMP 515), one course at MIT during in 1996 and 1997 (6.035), and one half-semester course at Harvey-Mudd College in 2012. The experience of independently teaching 6.035 at MIT in Fall 96 and Fall 97 (while on leave from IBM) was a stimulating experience for me; I recall having to rise to multiple challenges, such as ensuring that I verbally read out everything that I wrote on the board to make my lectures fully accessible to a visually impaired student (who, to my great satisfaction, would then ask questions about what I wrote). Finally, I have developed or co-developed material for several short courses and tutorials that have been given at leading international conferences, summer schools, and institutions over the last twenty-five years.

I believe that development of new pedagogic material goes hand-in-hand with conducting world-class research, and one of my hopes when moving to academia in 2007 was that I would be able to contribute to the creation of teaching material for the new era of mainstream parallel computing. This goal was realized by my creating a new undergraduate class, COMP 322 (Fundamentals of Parallel Programming) targeted at sophomores, in response to curriculum guidelines from ACM and other organizations to introduce parallel programming at the undergraduate level. COMP 322 is now a required course for all CS undergraduate majors at Rice. Thus far, COMP 322 has been taught at Rice in every Spring semester since 2011, and its material has been used in courses at other institutions including Brigham Young University, Harvey-Mudd College and Washington University. Programming assignments in COMP 322 make extensive use of the Habanero-Java library developed in my research group. I have created 80+ short, 5-7 minute videos to enable a "blended" experience for students: in-class discussions and worksheets combined with asynchronous learning through the videos. I have also written over 200 pages of lecture handout material, and have initiated discussions with publishers about textbook publication opportunities. A three-course online specialization based on COMP 322 was launched by Rice Online on Coursera in July 2017.

Overall, I consider teaching students to be a privilege, and among the most rewarding experiences of my career. In all my years in academia, I have never "bought out" of a teaching commitment even though I had the research funds to do so.

Courses

- Three-course online specialization on Parallel, Concurrent and Distributed Programming, Coursera, July 2017 – present.
- CS 6245 (Compiling for Parallelism), Georgia Tech, Fall 2020.
- CS 4240 (Compilers), Georgia Tech, Spring 2020.
- CS 6390 (Programming Languages), Georgia Tech, Fall 2019.
- CS 4240 (Compilers), Georgia Tech, Spring 2019.
- CS 6245 (Compiling for Parallelism), Georgia Tech, Fall 2018.
- CS 4240 (Compilers), Georgia Tech, Spring 2018.
- COMP 322 (Fundamentals of Parallel Programming), Rice University, Spring 2017.
- COMP 322 (Fundamentals of Parallel Programming), Rice University, Spring 2016.
- COMP 515 (Advanced Compilation for Vector and Parallel Processors), Rice University, Fall 2015.
- COMP 322 (Fundamentals of Parallel Programming), Rice University, Spring 2015.
- COMP 322 (Fundamentals of Parallel Programming), Rice University, Spring 2014.

- COMP 515 (Advanced Compilation for Vector and Parallel Processors), Rice University, Fall 2013.
- COMP 322 (Fundamentals of Parallel Programming), Rice University, Spring 2013.
- Co-instructor for CS 181E (Fundamentals of Parallel Programming), Harvey Mudd College, Fall 2012.
- COMP 322 (Fundamentals of Parallel Programming), Rice University, Spring 2012.
- Three-day course on "Introduction to Parallel Programming", CSIRO, Perth, Australia, December 2011.
- COMP 515 (Advanced Compilation for Vector and Parallel Processors), Rice University, Fall 2011.
- COMP 322 (Fundamentals of Parallel Programming), Rice University, Spring 2011.
- Invited short course on "Multicore Programming Models and their Compilation Challenges", ACACES 2010 Sixth International Summer School on Advanced Computer Architecture and Compilation for High-Performance and Embedded Systems, July 2010.
- COMP 211 (Principles of Program Design), Rice University, Spring 2010.
- COMP 322 (Fundamentals of Parallel Programming), Rice University, Fall 2009.
- COMP 515 (Advanced Compilation for Vector and Parallel Processors), Rice University, Spring 2009.
- Independent Study courses: COMP 590 (3 CAAM PhD students), COMP 490 (1 CS Undergraduate student), Fall 2008.
- Invited short course on "Multicore Programming Models", Second International School on Trends in Concurrency, June 2008.
- COMP 422 (Introduction to Parallel Computing), Rice University, Spring 2008.
- COMP 635 (Seminar on Heterogeneous Processors), Rice University, Fall 2007.
- 6.035 (Computer Language Engineering), MIT, Fall '96 & Fall '97.
  Sole instructor for the main compilers class at MIT. Restructured curriculum to use Java as the language foundation instead of CLU.
- Code Optimization in Modern Compilers, Western Institute for Computer Science (WICS), Stanford University, August '94, August '95, August '96.
  This one-week overview class on code optimization was designed for industry professionals and overseas students (taught jointly with Prof. Krishna Palem from the NYU Courant Institute).

- Compiling for Parallelism, Indian Institute of Science, March 1990.
  This intensive two-week course was sponsored by the United Nations Development Programme.
- CS 302, U. Wisconsin-Madison, Fall '81 & Spring '82.
  Sole lecturer for introductory programming course in Pascal and Fortran,

Conference Tutorials

- "HClib: A Task-based Parallel Programming Model" (with Vivek Kumar), given at HiPC 2018.
- "Habanero-Java: Multicore Programming for the Masses" (with Shams Imam), given at PPoPP 2014.
- "Compiler Challenges for Task-Parallel Languages", given at PLDI 2011.
- "The Concurrent Collections Parallel Programming Model - Foundations and Implementation Challenges" (with Kathleen Knobe), given at PLDI 2009.
- "Analysis and Optimization of Parallel Programs" (with Sam Midkiff), given at PLDI 2008.
- "Programming using the Partitioned Global Address Space (PGAS) Model" (with Tarek El-Ghazawi), given at Supercomputing 2007.
- "X10: Concurrent Object-Oriented Programming for Modern Architectures" (with Vijay Saraswat), given at OOPSLA 2006 and PPoPP 2007.
- "Programming Internet-Scale Distributed Applications in the 21st Century: BPEL and Beyond" (with John Field), given at ECOOP 2004 and OOPSLA 2003 conferences.
- "Optimized Compilation of Java Programs", given at ACM PLDI 2000 and ACM Java Grande 2001 conferences.
- "Static and Dynamic Optimized Compilation of Java Programs", given at the workshop on Java for High-Performance Computing, held in conjunction with ICS '99 (with Manish Gupta).
- "Code Optimization in Modern Compilers" given at the ASPLOS '96, HPCA '96, POPL '96, HiPC '95 conferences (with Krishna Palem).
- "Advanced Optimizations for Memory Hierarchies" given at the PLDI '93 and CASCON '94 conferences.
- "Instruction Scheduling" given at the SPDP '93 conference (with Barbara Simons). This tutorial was extended into an invited video lecture in University Video Communications' Distinguished Lecture Series.

11 Publications

I believe that my CS research accomplishments and my standing in the field contribute an essential dimension of experience and credibility to my leadership roles in general. As a researcher, I am fascinated by the challenge of advancing the foundations of high-level programming systems (which include languages, compilers, and runtimes) so as to fully exploit the latest advances in computing systems. While there has been a long tradition of research in this area since the dawn of computing, the rapid evolution of hardware has continuously fueled a need for new software technologies as old approaches quickly become obsolete, especially as we approach the end of Moore's Law. This has motivated the expansion of my group's research into software-hardware codesign research topics as evidenced by our recent publications in hardware-focused venues and algorithms-focused venues.

11.1 Refereed Conference and Journal Publications

1. Linear Promises: Towards Safer Concurrent Programming. O Rau, C Voss, V Sarkar. *35th European Conference on Object-Oriented Programming (ECOOP 2021).*

2. Task-graph scheduling extensions for efficient synchronization and communication. S Bak, O Hernandez, M Gates, P Luszczek, V Sarkar. *Proceedings of the 2021 ACM International Conference on Supercomputing.*

3. ARBALEST: Dynamic Detection of Data Mapping Issues in Heterogeneous OpenMP Applications. L Yu, J Protze, O Hernandez, V Sarkar. *2021 IEEE International Parallel and Distributed Processing Symposium (IPDPS).*

4. An ownership policy and deadlock detector for promises. C Voss, V Sarkar. *Proceedings of the 26th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming (PPoPP 2021).*

5. Towards Chapel-based Exascale Tree Search Algorithms: dealing with multiple GPU accelerators. T Carneiro, N Melab, A Hayashi, V Sarkar. *The 2020 International Conference on High Performance Computing & Simulation (HPCS 2020).*

6. Intrepydd: Performance, Productivity, and Portability for Data Science Application Kernels. Tong Zhou, Jun Shirako, Anirudh Jain, Sriseshan Srikanth, Thomas Conte, Richard Vuduc, Vivek Sarkar. *Proceedings of the 2020 ACM SIGPLAN International Symposium on New Ideas, New Paradigms, and Reflections on Programming and Software, Onward! 2020, November 2020.*

7. Vyasa: A High-Performance Vectorizing Compiler for Tensor Convolutions on the Xilinx AI Engine. Prasanth Chatarasi, Stephen Neuendorffer, Samuel Bayliss, Kees Vissers, and Vivek Sarkar *Proceedings of the 24th IEEE High Performance Extreme Computing Conference (HPEC20), September 2020.*

8. OmpMemOpt: Optimized Memory Movement for Heterogeneous Computing. Prithayan Barua, Jisheng Zhao, Vivek Sarkar. *27th International European Conference on Parallel and Distributed Computing (EuroPar), August 2020.*

9. MAESTRO: A Data-Centric Approach to Understand Reuse, Performance, and Hardware Cost of DNN Mappings. Hyoukjun Kwon, Prasanth Chatarasi, Vivek Sarkar, Tushar Krishna, Michael Pellauer, Angshuman Parashar. *IEEE Micro Vol. 40, no. 3, May-June 2020.*

10. Experimental Insights from the Rogues Gallery. Jeffrey S Young, Jason Riedy, Thomas M Conte, Vivek Sarkar, Prasanth Chatarasi, Sriseshan Srikanth. *2019 IEEE International Conference on Rebooting Computing (ICRC), November 2019.*

11. Understanding Reuse, Performance, and Hardware Cost of DNN Dataflows: A Data-Centric Approach. Hyoukjun Kwon, Prasanth Chatarasi, Michael Pellauer, Angshuman Parashar, Vivek Sarkar, Tushar Krishna. *The 52nd Annual IEEE/ACM International Symposium on Microarchitecture (MICRO), October 2019.* (Selected as IEEE Micro Top Pick for 2019.)

12. Enabling Resilience in Asynchronous Many-Task Programming Models Authors Sri Raj Paul, Akihiro Hayashi, Nicole Slattengren, Hemanth Kolla, Matthew Whitlock, Seonmyeong Bak, Keita Teranishi, Jackson Mayo, Vivek Sarkar. *26th International European Conference on Parallel and Distributed Computing (Euro-Par), August 2019.*

13. Optimized Execution of Parallel Loops via User-Defined Scheduling Policies Seonmyeong Bak, Yanfei Guo, Pavan Balaji, Vivek Sarkar. *The 48th International Conference on Parallel Processing (ICPP), August 2019.*

14. T2S-Tensor: Productively Generating High-Performance Spatial Hardware for Dense Tensor Computations. Nitish Kumar Srivastava, Hongbo Rong, Prithayan Barua, Guanyu Feng, Huanqi Cao, Zhiru Zhang, Vivek Sarkar, Wenguang Chen, Paul Petersen, Geoff Lowney, Christopher Hughes, Timothy Mattson, Pradeep Dubey. *27th IEEE International Symposium On Field-Programmable Custom Computing Machines (FCCM), April 2019.*

15. Transitive Joins: A Sound and Efficient Online Deadlock-Avoidance Policy. Caleb Voss, Tiago Cogumbreiro, Vivek Sarkar. *ACM Conference onPrinciples and Practice of Parallel Programming (PPoPP), February 2019.*

16. Valence: Variable Length Calling Context Encoding. Tong Zhou, Michael R. Jantz, Prasad A. Kulkarni, Kshitij A. Doshi, Vivek Sarkar. *28th International Conference on Compiler Construction (CC), February 2019.*

17. Topkapi: Parallel and Fast Sketches for Finding Top-K Frequent Elements. Ankush Mandal, He Jiang, Anshumali Shrivastava, Vivek Sarkar. *Advances in Neural Information Processing Systems 31 (NeurIPS), December 2018.*

18. Detecting MPI usage anomalies via partial program symbolic execution. Fangke Ye, Jisheng Zhao, Vivek Sarkar. *The International Conference for High Performance Computing, Networking, Storage, and Analysis (SC18), November 2018.*

19. Cost-driven thread coarsening for GPU kernels. Prithayan Barua, Jun Shirako, Vivek Sarkar. *27th International Conference on Parallel Architectures and Compilation Techniques (PACT), November 2018.*

20. In-Register Parameter Caching for Dynamic Neural Nets with Virtual Persistent Processor Specialization. Farzad Khorasani, Hodjat Asghari Esfeden, Nael Abu-Ghazaleh, Vivek Sarkar. *The 51st Annual IEEE/ACM International Symposium on Microarchitecture (MICRO), October 2018.*

21. Using Dynamic Compilation to Achieve Ninja Performance for CNN Training on Many-Core Processors. Ankush Mandal, Raj Barik, Vivek Sarkar. *25th International European Conference on Parallel and Distributed Computing (Euro-Par), August 2018.*

22. GT-Race: graph traversal based data race detection for asynchronous many-task parallelism. Lechen Yu, Vivek Sarkar. *25th International European Conference on Parallel and Distributed Computing (Euro-Par), August 2018.*

23. RegMutex: Inter-Warp GPU Register Time-Sharing. Farzad Khorasani, Hodjat Asghari Esfeden, Amin Farmahini-Farahani, Nuwan Jayasena, Vivek Sarkar. *2018 ACM/IEEE 45th Annual International Symposium on Computer Architecture (ISCA), June 2018.*

24. Parallel Sparse Flow-Sensitive Points-to Analysis. Jisheng Zhao, Michael G. Burke, Vivek Sarkar. *Proceedings of the 2018 International Conference on Compiler Construction (CC 2018), February 2018.*

25. Modeling the Conflicting Demands of Multi-Level Parallelism and Temporal/Spatial Locality in Affine Scheduling. Oleksandr Zinenko, Chandan Reddy, Sven Verdoolaege, Jun Shirako, Tobias Grosser, Vivek Sarkar, Albert Cohen. *Proceedings of the 2018 International Conference on Compiler Construction (CC 2018), February 2018.*

26. Deadlock Avoidance in Parallel Programs with Futures: Why parallel tasks should not wait for strangers. Tiago Cogumbreiro, Rishi Surendran, Francisco Martins, Vivek Sarkar, Vasco T. Vasconcelos, and Max Grossman. *ACM SIGPLAN International Conference on Object-Oriented Programming, Systems, Languages, and Applications (OOPSLA). ACM, 2017.*

27. DAMMP: A Distributed Actor Model for Mobile Platforms. Arghya Chatterjee, Srdjan Milakovic, Bing Xue, Zoran Budimlic, Vivek Sarkar. *14th International Conference on Managed Languages and Runtimes (ManLang'17), September 2017.*

28. A Marshalled Data Format for Pointers in Relocatable Data Blocks. Nick Vrvilo, Lechen Yu and Vivek Sarkar. *Proceedings of the 2017 ACM SIGPLAN International Symposium on Memory Management (ISMM). June 2017.*

29. Performance Evaluation of OpenMP's Target Construct on GPUs. Akihiro Hayashi, Jun Shirako, Ettore Tiotto, Robert Ho, Vivek Sarkar. *International Journal of High Performance Computing and Networking (IJHPCN), June 2017.*

30. Formalization of Habanero Phasers using Coq. Tiago Cogumbreiro, Jun Shirako, and Vivek Sarkar. *Journal of Logical and Algebraic Methods in Programming (JLAMP), March 2017.*

31. Optimized Two-Level Parallelization for GPU Accelerators using the Polyhedral Model. Jun Shirako, Akihiro Hayashi, Vivek Sarkar. *Proceedings of the 2017 International Conference on Compiler Construction (CC 2017), February 2017.*

32. PIPES: A Language and Compiler for Task-Based Programming on Distributed-Memory Clusters. *The Conference on High Performance Computing, Networking, Storage and Analysis (SC16), November 2016.*

33. Automatic Parallelization of Pure Method Calls via Conditional Future Synthesis. Rishi Surendran, Vivek Sarkar. *2016 ACM SIGPLAN International Conference on Object-Oriented Programming, Systems, Languages, and Applications (OOPSLA 2016), November 2016.*

34. The Open Community Runtime: A Runtime System for Extreme Scale Computing. Timothy G. Mattson et al. *2016 IEEE High Performance Extreme Computing Conference (HPEC 16), October 2016.*

35. Dynamic Determinacy Race Detection for Task Parallelism with Futures. Rishi Surendran and Vivek Sarkar. *16th International Conference on Runtime Verification (RV'16), September 2016.*

36. Declarative Tuning for Locality in Parallel Programs. Sanjay Chatterjee, Nick Vrvilo, Zoran Budimlic, Kathleen Knobe, Vivek Sarkar. *The 45th International Conference on Parallel Processing (ICPP), August 2016.*

37. A Distributed Selectors Runtime System for Java Applications. Arghya Chatterjee, BrankoGvoka, Bing Xue, Zoran Budimlic, Shams Imam, Vivek Sarkar. *13th International Conference on the Principles and Practice of Programming on the Java Platform: virtual machines, languages, and tools (PPPJ'16), August 2016.*

38. SWAT: A Programmable, In-Memory, Distributed, High-Performance Computing Platform. Max Grossman, Vivek Sarkar. *International ACM Symposium on High-Performance Parallel and Distributed Computing (HPDC), May 2016.*

39. Efficient Checkpointing of Multi-Threaded Applications as a Tool for Debugging, Performance Tuning, and Resilience. Max Grossman, Vivek Sarkar. *IEEE International Parallel and Distributed Processing Symposium (IPDPS). May 2016.*

40. Formalization of phase ordering. Tiago Cogumbreiro, Jun Shirako, Vivek Sarkar. *Programming Language Approaches to Concurrency- and Communication-cEntric Software (PLACES 2016), April 2016.*

41. Automatic Data Layout Generation and Kernel Mapping for CPU+GPU Architectures. Deepak Majeti, Kuldeep Meel, Raj Barik and Vivek Sarkar. *25th International Conference on Compiler Construction (CC 2016), March 2016.*

42. Polyhedral Optimizations of Explicitly Parallel Programs. Prasanth Chatarasi, Jun Shirako, and Vivek Sarkar. *24th International Conference on Parallel Architectures and Compilation Techniques (PACT), October 2015.*

43. Compiling and Optimizing Java 8 Programs for GPU execution. Kazuaki Ishizaki, Akihiro Hayashi, Gita Koblents, Vivek Sarkar. *24th International Conference on Parallel Architectures and Compilation Techniques (PACT), October 2015.*

44. Heterogeneous Work-stealing across CPU and DSP cores. Vivek Kumar, Alina Sbirlea, Ajay Jayaraj, Zoran Budimlic, Deepak Majeti, Vivek Sarkar. *19th IEEE High Performance Extreme Computing conference (HPEC'15). September 2015.*

45. HJlib-CL: Reducing the Gap Between the JVM and Accelerators. Max Grossman, Shams Imam, Vivek Sarkar. *12th International Conference on the Principles and Practice of Programming on the Java Platform (PPPJ'15), September 2015.*

46. Machine-Learning-based Performance Heuristics for Runtime CPU/GPU Selection. Akihiro Hayashi, Kazuaki Ishizaki, Gita Koblents, Vivek Sarkar. *12th International Conference on the Principles and Practice of Programming on the Java Platform: virtual machines, languages, and tools (PPPJ'15), September 2015.*

47. A Composable Deadlock-free Approach to Object-based Isolation. Shams Imam, Jisheng Zhao and Vivek Sarkar. *21st International European Conference on Parallel and Distributed Computing (Euro-Par'15), August 2015.*

48. Elastic Tasks: Unifying Task Parallelism and SPMD Parallelism with an Adaptive Runtime. Alina Sbirlea, Kunal Agrawal and Vivek Sarkar. *21st International European Conference on Parallel and Distributed Computing (Euro-Par'15), August 2015.*

49. Load Balancing Prioritized Tasks via Work-Stealing. Shams Imam and Vivek Sarkar. *21st International European Conference on Parallel and Distributed Computing (Euro-Par'15), August 2015.*

50. Data Layout Optimization for Portable Performance. Kamal Sharma, Ian Karlin, Jeff Keasler, James McGraw, Vivek Sarkar. *21st International European Conference on Parallel and Distributed Computing (Euro-Par'15), August 2015.*

51. The Eureka Programming Model for Speculative Task Parallelism. Shams Imam, Vivek Sarkar. *29th European Conference on Object-Oriented Programming (ECOOP), July 2015.*

52. HadoopCL2: Motivating the Design of a Distributed, Heterogeneous Programming System With Machine-Learning Applications. Max Grossman, Mauricio Breternitz, Vivek Sarkar. *IEEE Transactions on Parallel and Distributed Systems (TPDS). 2015.*

53. JPF Verification of Habanero Java Programs using Gradual Type Permission Regions. Peter Anderson, Nick Vrvilo, Eric Mercer, and Vivek Sarkar. 2015. *SIGSOFT Softw. Eng. Notes 40, 1 (February 2015), 1-5.*

54. Oil and Water Can Mix: An Integration of Polyhedral and AST-based Transformations. Jun Shirako, Louis-Noel Pouchet, Vivek Sarkar. *The Conference on High Performance Computing, Networking, Storage and Analysis (SC'14), November 2014.*

55. HabaneroUPC++: a Compiler-free PGAS Library. Vivek Kumar, Yili Zheng, Vincent Cave, Zoran Budimlic, Vivek Sarkar. *8th International Conference on Partitioned Global Address Space Programming Models (PGAS14), October 2014.*

56. Habanero-Java Library: a Java 8 Framework for Multicore Programming. Shams Imam, Vivek Sarkar. *11th International Conference on the Principles and Practice of Programming on the Java platform: virtual machines, languages, and tools (PPPJ'14), September 2014.*

57. Bounded Memory Scheduling of Dynamic Task Graphs. Dragos Sbirlea, Zoran Budimlic, Vivek Sarkar. *International Conference on Parallel Architectures and Compilation Techniques (PACT), August 2014.*

58. Cooperative Scheduling of Parallel Tasks with General Synchronization Patterns. Shams Imam, Vivek Sarkar. *28th European Conference on Object-Oriented Programming (ECOOP), July 2014.*

59. Test-Driven Repair of Data Races in Structured Parallel Programs. Rishi Surendran, Raghavan Raman, Swarat Chaudhuri, John Mellor-Crummey, Vivek Sarkar. *35th ACM Conference on Programming Language Design and Implementation (PLDI), June 2014.*

60. Inter-iteration Scalar Replacement Using Array SSA Form. Rishi Surendran, Rajkishore Barik, Jisheng Zhao, Vivek Sarkar. *The 23rd International Conference on Compiler Construction (CC 2014), April 2014.*

61. A Decoupled non-SSA Global Register Allocation using Bipartite Liveness Graphs. Rajkishore Barik, Jisheng Zhao and Vivek Sarkar. *ACM Transactions on Architecture and Code Optimization (TACO), Volume 10 Issue 4, December 2013.*

62. Automatic Detection of Inter-application Permission Leaks in Android Applications. Dragos Sbirlea, Michael G. Burke, Salvatore Guarnieri, Marco Pistoia, Vivek Sarkar. *IBM Journal of Research and Development (Volume:57 , Issue: 6), November–December 2013.*

63. Isolation for Nested Task Parallelism. Jisheng Zhao, Roberto Lublinerman, Zoran Budimlic, Swarat Chaudhuri, Vivek Sarkar. *The 29th International Conference on the Object-Oriented Programming, System, Languages and Application (OOPSLA), October 2013.*

64. Interprocedural Strength Reduction of Critical Sections in Explicitly-Parallel Programs. Raj Barik, Jisheng Zhao, Vivek Sarkar. *Proceedings of the 22nd International Conference on Parallel Architectures and Compilation Techniques (PACT), September 2013.*

65. Accelerating Habanero-Java Programs with OpenCL Generation. Akihiro Hayashi, Max Grossman, Jisheng Zhao, Jun Shirako, Vivek Sarkar. *10th International Conference on the Principles and Practice of Programming in Java (PPPJ), September 2013.*

66. A Transformation Framework for Optimizing Task-Parallel Programs. Krishna Nandivada, Jun Shirako, Jisheng Zhao, Vivek Sarkar. *ACM Transactions on Programming Languages and Systems (TOPLAS), Volume 35 Issue 1, April 2013.*

67. Integrating Asynchronous Task Parallelism with MPI. Sanjay Chatterjee, Sagnak Tasirlar, Zoran Budimlic, Vincent Cave, Milind Chabbi, Max Grossman, Vivek Sarkar. *27th IEEE International Parallel and Distributed Processing Symposium (IPDPS), April 2013.*

68. Integrating Task Parallelism with Actors. Shams Imam, Vivek Sarkar. *SPLASH Conference on Object-Oriented Programming, Systems, Languages, and Applications (OOPSLA), October 2012.*

69. Folding of Tagged Single Assignment Values for Memory-Efficient Parallelism. Dragos Sbirlea, Kathleen Knobe, Vivek Sarkar. *International European Conference on Parallel and Distributed Computing (Euro-Par), August 2012.*

70. A Practical Approach to DOACROSS Parallelization. Priya Unnikrishnan, Jun Shirako, Kit Barton, Sanjay Chatterjee, Raul Silvera, Vivek Sarkar. *International European Conference on Parallel and Distributed Computing (Euro-Par), August 2012.*

71. Design, Verification and Applications of a New Read-Write Lock Algorithm. Jun Shirako, Nick Vrvilo, Eric G. Mercer, Vivek Sarkar. *24th ACM Symposium on Parallelism in Algorithms and Architectures (SPAA), June 2012.*

72. Scalable and Precise Dynamic Datarace Detection for Structured Parallelism. Raghavan Raman, Jisheng Zhao, Vivek Sarkar, Martin Vechev, Eran Yahav. *Proceedings of 2012 ACM Conference on Programming Language Design and Implementation (PLDI), June 2012.*

73. Mapping a Data-Flow Programming Model onto Heterogeneous Platforms. Alina Sbirlea, Yi Zou, Zoran Budimlic, Jason Cong, Vivek Sarkar. *Conference on Languages, Compilers, Tools and Theory for Embedded Systems (LCTES), June 2012.*

74. Practical Permissions for Race-Free Parallelism. Edwin Westbrook, Jisheng Zhao, Zoran Budimlic, Vivek Sarkar. *26th European Conference on Object-Oriented Programming (ECOOP), June 2012.*

75. Analytical Bounds for Optimal Tile Size Selection. Jun Shirako, Kamal Sharma, Naznin Fauzia, Louis-Noel Pouchet, J. Ramanujam, P. Sadayappan, Vivek Sarkar. *Proceedings of the 2012 International Conference on Compiler Construction (CC 2012), April 2012.*

76. Delegated Isolation. Roberto Lublinerman, Jisheng Zhao, Zoran Budimlic, Swarat Chaudhuri, Vivek Sarkar. *Proceedings of OOPSLA 2011, October 2011.*

77. Interfacing Chapel with Traditional HPC Programming Languages. Adrian Prantl, Thomas Epperly, Shams Imam, Vivek Sarkar. *PGAS11 Conference, October 2011.*

78. Permission Regions for Race-Free Parallelism. Edwin Westbrook, Jisheng Zhao, Zoran Budimlic, Vivek Sarkar. *Proceedings of the 2nd International Conference on Runtime Verification (RV '11), September 2011.*

79. Data-Driven Tasks and their Implementation. Sagnak Tasirlar, Vivek Sarkar. *Proceedings of the International Conference on Parallel Processing (ICPP), September 2011.*

80. Habanero-Java: the New Adventures of Old X10. Vincent Cave, Jisheng Zhao, Jun Shirako, Vivek Sarkar. *9th International Conference on the Principles and Practice of Programming in Java (PPPJ), August 2011.*

81. DrHJ — a lightweight pedagogic IDE for Habanero Java. Jarred Payne, Vincent Cave, Raghavan Raman, Mathias Ricken, Robert Cartwright, Vivek Sarkar. *Tool Demonstration paper, 9th International Conference on the Principles and Practice of Programming in Java (PPPJ), August 2011.*

82. Hardware and Software Tradeoffs for Task Synchronization on Manycore Architectures. Yonghong Yan, Sanjay Chatterjee, Daniel Orozco, Elkin Garcia, Zoran Budimlic, Jun Shirako, Robert Pavel, Guang R. Gao, and Vivek Sarkar. *Proceedings of Euro-Par 2011, August 2011.*

83. Optimizations for Distributed-Memory X10 Programs. Rajkishore Barik, Jisheng Zhao, David Grove, Igor Peshansky, Zoran Budimli, Vivek Sarkar. *25th IEEE International Parallel and Distributed Processing Symposium (IPDPS), April 2011.*

84. Subregion Analysis and Bounds Check Elimination for High Level Arrays. Mackale Joyner, Zoran Budimli, Vivek Sarkar. *Proceedings of the 2011 International Conference on Compiler Construction (CC 2011), April 2011.*

85. Customizable Domain-Specific Computing. Jason Cong, Vivek Sarkar, Glenn Reinman, and Alex Bui. *IEEE Design & Test, 2:28, pp. 6–15, March 2011.*

86. Efficient Selection of Vector Instructions using Dynamic Programming. Rajkishore Barik, Jisheng Zhao, Vivek Sarkar. *MICRO-43, December 2010.*

87. Efficient Data Race Detection for Async-Finish Parallelism. Raghavan Raman, Jisheng Zhao, Vivek Sarkar, Martin Vechev, Eran Yahav. *Proceedings of the 1st International Conference on Runtime Verification (RV 10), November 2010. Recipient of Best Paper Award.*

88. Reducing Task Creation and Termination Overhead in Explicitly Parallel Programs. Jisheng Zhao, Jun Shirako, Krishna V. Nandivada, Vivek Sarkar. *Proceedings of the Nineteenth International Conference on Parallel Architectures and Compilation Techniques (PACT), September 2010.*

89. Verifying Determinism of Structured Parallel Programs. Martin Vechev, Eran Yahav, Raghavan Raman, Jisheng Zhao, Vivek Sarkar. *Proceedings of the 2010 Static Analysis Symposium (SAS '10), September 2010.*

90. Concurrent Collections. Zoran Budimlić, Michael Burke, Vincent Cavé, Kathleen Knobe, Geoff Lowney, Ryan Newton, Jens Palsberg, David Peixotto, Vivek Sarkar, Frank Schlimbach, Sağnak Taşirlar. *Scientific Programming, 18:3–4, pp. 203–217, August 2010.*

91. A Study of a Software Cache Implementation of the OpenMP Memory Model for Multicore and Manycore Architectures. Chen Chen, Joseph B. Manzano, Ge Gan, Guang R. Gao, and Vivek Sarkar. *Proceedings of Euro-Par 2010, August 2010.*

92. SLAW: a Scalable Locality-aware Adaptive Work-stealing Scheduler. Yi Guo, Jisheng Zhao, Vincent Cav, Vivek Sarkar. *24th IEEE International Parallel and Distributed Processing Symposium (IPDPS), April 2010.*

93. Hierarchical Phasers for Scalable Synchronization and Reduction. Jun Shirako, Vivek Sarkar. *24th IEEE International Parallel and Distributed Processing Symposium (IPDPS), April 2010.*

94. Software Challenges in Extreme Scale Systems. V. Sarkar, W. Harrod, A.E. Snavely. *SciDAC Review Special Issue on Advanced Computing: The Roadmap to Exascale, pp. 60-65, 2010 (revised version of reference [99] below).*

95. Programmability Issues. Barbara Chapman, Jess Labarta, Vivek Sarkar, Mitsuhisa Sato. *International Journal of High Performance Computing Applications (IJHPCA), November 2009, volume 23, number 4, pp 328-331.*

96. Interprocedural Load Elimination for Dynamic Optimization of Parallel Programs. Rajkishore Barik, Vivek Sarkar. *The Eighteenth International Conference on Parallel Architectures and Compilation Techniques (PACT), September 2009.*

97. JCUDA: a Programmer-Friendly Interface for Accelerating Java Programs with CUDA. Yonghong Yan, Max Grossman, Vivek Sarkar. *Proceedings of Euro-Par 2009, August 2009.*

98. Chunking Parallel Loops in the Presence of Synchronization. Jun Shirako, Jisheng Zhao, V. Krishna Nandivada, Vivek Sarkar. *Proceedings of the 2009 ACM International Conference on Supercomputing (ICS), June 2009.*

99. Software Challenges in Extreme Scale Systems. V. Sarkar, W. Harrod, A.E. Snavely. *2009 Conference on Scientific Discovery through Advanced Computing Program (SciDAC), June 2009.*

100. Work-First and Help-First Scheduling Policies for Terminally Strict Parallel Programs. Yi Guo, Rajkishore Barik, Raghavan Raman, Vivek Sarkar. *23rd IEEE International Parallel and Distributed Processing Symposium (IPDPS), May 2009.*

101. Phaser Accumulators: a New Reduction Construct for Dynamic Parallelism. Jun Shirako, David Peixotto, Vivek Sarkar, William Scherer. *23rd IEEE International Parallel and Distributed Processing Symposium (IPDPS), May 2009.*

102. Phasers: a Unied Deadlock-Free Construct for Collective and Point-to-point Synchronization. Jun Shirako, David Peixotto, Vivek Sarkar, William Scherer. *Proceedings of the 2008 ACM International Conference on Supercomputing (ICS), June 2008.*

103. Type Inference for Locality Analysis of Distributed Data Structures. Satish Chandra, Vijay Saraswat, Vivek Sarkar, Ras Bodik. *Proceedings of the 2008 ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming (PPoPP), February 2008.*

104. Optimizing Array Accesses in High Productivity Languages. Mackale Joyner, Zoran Budimlić, Vivek Sarkar. *Proceedings of the 2007 High Performance Computation Conference (HPCC), September 2007.*

105. Deadlock-Free Scheduling of X10 Computations with Bounded Resources. Shivali Agarwal, Rajkishore Barik, Dan Bonachea, Vivek Sarkar, Rudrapatna Shyamasundar, Katherine Yelick. *Proceedings of the 19th Annual ACM Symposium on Parallelism in Algorithms and Architectures (SPAA '07), June 2007.*

106. May-Happen-in-Parallel Analysis of X10 programs. Shivali Agarwal, Rajkishore Barik, Vivek Sarkar, Rudrapatna Shyamasundar. *Proceedings of the 2007 ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming (PPoPP), March 2007.*

107. Extended Linear Scan: an Alternate Foundation for Global Register Allocation. Vivek Sarkar, Rajkishore Barik. *Proceedings of the 2007 International Conference on Compiler Construction (CC 2007), March 2007.*

108. The Role of Programming Languages in Future Data-Centric and Net-Centric Applications. Vivek Sarkar, Mukund Raghavachari. *Proceedings of the Third International Conference on Distributed Computing and Internet Technology (ICDCIT), December 2006.*

109. Enhanced Bitwidth-Aware Register Allocation. Rajkishore Barik, Vivek Sarkar. *Proceedings of the 2006 International Conference on Compiler Construction (CC 2006), March 2006.*

110. X10: An Object-Oriented Approach to Non-Uniform Cluster Computing. Philippe Charles, Christopher Donawa, Kemal Ebcioglu, Christian Grothoff, Allan Kielstra, Christoph von Praun, Vijay Saraswat, Vivek Sarkar. *Proceedings of the ACM 2005 OOPSLA conference, Onward! track, October 2005.*

111. Immutability Specification and its Applications. Igor Pechtchanski, Vivek Sarkar. *Concurrency and Computation Practice & Experience (CCPE), 17(5:6), April 2005.*

112. The Jikes Research Virtual Machine project: building an open-source research community. B. Alpern, S. Augart, S.M. Blackburn, M. Butrico, A. Cocchi, P. Cheng, J. Dolby, S. Fink, D. Grove, M. Hind, K.S. McKinley, M. Mergen, J.E.B. Moss, T. Ngo, V. Sarkar, M. Trapp. *IBM Systems Journal special issue on Open Source Software, Volume 44, Number 2, June 2005.*

113. XJ : Facilitating XML Processing in Java. Matthew Harren, Mukund Raghavachari, Oded Shmueli, Michael Burke, Rajesh Bordawekar, Igor Pechtchanski, Vivek Sarkar. *Proceedings of the Fourteenth International World Wide Web Conference (WWW 2005), May 2005.*

114. Decentralizing Execution of Composite Web Services. Mangala Gowri Nanda, Satish Chandra, Vivek Sarkar. *Proceedings of the ACM 2004 OOPSLA conference, October 2004.*

115. Runtime Register Allocation. Kemal Ebcioglu, Vivek Sarkar, Kartik Agaram. *Proceedings of the $P = AC^2$ conference, October 2004.*

116. Integrating Database and Programming Language Constraints. Mukund Raghavachari, Oded Shmueli, Vivek Sarkar, Michael Burke, Rajesh Bordawekar. *Proceedings of the 9th International Conference on Data Base Programming Languages (DBPL), September 2003.*

117. Immutability Specification and its Applications. Igor Pechtchanski, Vivek Sarkar. *Proceedings of the ACM 2002 Java Grande/ISCOPE Conference, October 2002.*

118. Efficient and precise datarace detection for multithreaded object-oriented programs. Jong-Deok Choi, Keunwoo Lee, Alexey Loginov, Robert O'Callahan, Vivek Sarkar, Manu Sridharan *Proceedings of the ACM SIGPLAN 2002 Conference on Programming language design and implementation PLDI '02.*

119. Unconstrained Static Scheduling with Communication Weights. Vivek Sarkar. *Journal of Scheduling, Volume 5, Issue 5, 2002.*

120. High-Performance Scalable Java Virtual Machines. Vivek Sarkar, Julian Dolby. *Proceedings of the HiPC 2001 conference, Dec 2001, pp. 151-166.*

121. Dynamic Optimistic Interprocedural Analysis: A Framework and an Application. Igor Pechtchanski, Vivek Sarkar. *Proceedings of the ACM 2001 OOPSLA conference, October 2001, pp. 195-210.*

122. Optimized Unrolling of Nested Loops. Vivek Sarkar. *International Journal on Parallel Programming 29(5), pps. 545-581, October 2001.*

123. Efficient Dependence Analysis for Java Arrays. Vivek Sarkar, Stephen J. Fink. *Proceedings of the Euro-Par 2001 conference, August 2001, pp. 273-277.*

124. Reducing the Overhead of Dynamic Compilation. Chandra Krintz, David Grove, Vivek Sarkar, Brad Calder. *Software Practice and Experience, 31(8), July 2001, pp 717–738.*

125. Register-sensitive selection, duplication, and sequencing of instructions. Vivek Sarkar, Mauricio J. Serrano, Barbara B. Simons. *Proceedings of the ACM 2001 International Conference on Supercomputing, June 2001, pp. 277-288.*

126. Unified Analysis of Array and Object References in Strongly Typed Languages. S.Fink, K.Knobe, V.Sarkar. *Proceedings of the 2000 Static Analysis Symposium (SAS '00), October 2000.*

127. Lightweight Object-Oriented Shared Variables for Cluster Computing in Java. J. Harris, V. Sarkar. *Journal on Parallel and Distributed Computing (JPDC), 60(10), October 2000, pp. 1238-1259.*

128. Location consistency – a new memory model and cache consistency protocol. Guang Gao, Vivek Sarkar. *IEEE Transactions on Computers, August 2000, Vol. 49, No. 8, pp 798-813.*

129. The Jalapeño Virtual Machine. B. Alpern et al. *Special issue of IBM Systems journal on Java performance, 39(1), 2000.*

130. ABCD: Eliminating Array Bounds Checks on Demand. R.Bodik, R.Gupta, V.Sarkar. *ACM SIGPLAN 2000 Conference on Programming Language Design and Implementation (PLDI), June 2000.*

131. Optimized Unrolling of Nested Loops. Vivek Sarkar. *ACM International Conference on Supercomputing (ICS), May, 2000.*

132. An Analytical Model for Loop Tiling and its Solution. Vivek Sarkar and Nimrod Megiddo. *IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS-2000), April 2000.*

133. Linear Scan Register Allocation. Massimiliano Poletto and Vivek Sarkar. *ACM Trans. Program. Lang. Syst. 21, 5 (Sep. 1999), Pages 895 - 913.*

134. Compiling for Parallel Systems. R. Gupta, S. Pande, K. Psarris, V.Sarkar. *Invited survey paper on parallelizing compilers in special issue of Parallel Computing journal, 1999.*

135. The Jalapeño Dynamic Optimizing Compiler for Java. Michael G. Burke, Jong-Deok Choi, Stephen Fink, David Grove, Michael Hind, Vivek Sarkar, Mauricio J. Serrano, V. C. Sreedhar, Harini Srinivasan, John Whaley. *Proceedings of the ACM 1999 Java Grande Conference, June, 1999.*

136. Enabling Sparse Constant Propagation of Array Elements via Array SSA Form. V.Sarkar and K.Knobe. *Proceedings of the 1998 Static Analysis Symposium (SAS '98), October 1998.*

137. Space-Time Scheduling of Instruction-Level Parallelism on a Raw Machine. W.Lee, R.Barua, M.Frank, D.Srikrishna, J.Babb, V.Sarkar, S.Amarasinghe. *Proceedings of the Eighth International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS-VIII), October 1998.*

138. Lightweight Object-Oriented Shared Variables for Distributed Applications on the Internet". J.Harris and V.Sarkar. *Proceedings of OOPSLA '98, October 1998.*

139. Array SSA form and its use in Parallelization. Kathleen Knobe and Vivek Sarkar. *Proceedings of the 25th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, San Diego, California, January 1998.*

140. On the Importance of an End-To-End View of Memory Consistency in Future Computer Systems. Guang R. Gao and Vivek Sarkar. *Proceedings of the 1997 International Symposium on High Performance Computing, Fukuoka, Japan, November 1997.*

141. Baring it All to Software: Raw Machines. E. Waingold, M. Taylor, V. Sarkar, W. Lee, V. Lee, J. Kim, M. Frank, P. Finch, S. Devabhaktuni, R. Barua, J. Babb, S. Amarasinghe, A. Agarwal. *IEEE Computer, September 1997.* Special issue on "Future Microprocessors — How to use a Billion Transistors".

142. False Sharing Elimination by Selection of Runtime Scheduling Parameters. Jyh-Herng Chow and Vivek Sarkar. *Proceedings of the 26th International Conference on Parallel Processing, Bloomington, Illinois, August 1997, pages 396-403.*

143. Optimal Weighted Loop Fusion for Parallel Programs. Nimrod Megiddo and Vivek Sarkar. *Proceedings of the Ninth Annual ACM Symposium on Parallel Algorithms and Architecture (SPAA), Newport, Rhode Island, June, 1997, pages 282-291.*

144. Automatic Selection of High Order Transformations in the IBM XL Fortran Compilers. Vivek Sarkar. *IBM Journal of Research and Development, 41(3), May 1997.*

145. Automatic Parallelization for Symmetric Shared-Memory Multiprocessors. Jyh-Herng Chow, Leonard E. Lyon, and Vivek Sarkar. *Proceedings of CASCON '96 conference, Toronto, Canada, November 1996.*

146. Anticipatory Instruction Scheduling. Vivek Sarkar and Barbara Simons. *Proceedings of the Eighth Annual ACM Symposium on Parallel Algorithms and Architectures (SPAA), June 1996.*

147. Incremental Computation of Static Single Assignment Form. Jong-Deok Choi, Vivek Sarkar, and Edith Schonberg. *Proceedings of the 1996 International Conference on Compiler Construction, Linkoping, Sweden, April 1996.*

148. Optimized code restructuring of OS/2 executables. Jyh-Herng Chow, Yong-fong Lee, Kalyan Muthukumar, Vivek Sarkar, Mauricio Serrano, Iris Garcia, John Hsu, Shauchi Ong, Honesty Young. *Proceedings of CASCON '95 conference, Toronto, Canada, November 1995.*

149. Location Consistency: Stepping Beyond the Memory Coherence Barrier. Guang R. Gao and Vivek Sarkar. *Proceedings of the 24th International Conference on Parallel Processing, Oconomowoc, Wisconsin, August 1995.*

150. Mapping Iterative Task Graphs on Distributed-memory Machines. Tao Yang, Cong Fu, Apostolos Gerasoulis, and Vivek Sarkar. *Proceedings of the 24th International Conference on Parallel Processing, Oconomowoc, Wisconsin, August 1995, Volume II, pages 151-158.*

151. Scheduling Iterative Task Computations on Message-passing Architectures. T. Yang, P. Diniz, A. Gerasoulis, V. Sarkar. *Proceedings of the Seventh SIAM Conference on Parallel Processing for Scientific Computing, San Francisco, California, February 1995.*

152. A Compiler Framework for Restructuring Data Declarations to Enhance Cache and TLB Effectiveness. David F. Bacon, Jyh-Herng Chow, Dz ching R. Ju, K. Muthukumar, and Vivek Sarkar. *Proceedings of CASCON '94 conference, Toronto, Canada, November 1994.*

153. Automatic Localization for Distributed-Memory Multiprocessors Using a Shared-Memory Compilation Framework. Vivek Sarkar and Lelia A. Vazquez. *Proceedings of the 27th Hawaii International Conference on System Sciences, Wailea, Hawaii, Volume II (Software Technology), pages 4-13, January 1994.*

154. An Optimal Asynchronous Scheduling Algorithm for Software Cache Consistency. Barbara Simons, Vivek Sarkar, Mauricio Breternitz, Michael Lai. *Proceedings of the 27th Hawaii International Conference on System Sciences, Wailea, Hawaii, Volume II (Software Technology), pages 502-511, January 1994.*

155. Data Merging for Shared-Memory Multiprocessors. Alan H. Karp and Vivek Sarkar. *Proceedings of the 26th Hawaii International Conference on System Sciences, Wailea, Hawaii, Volume I (Architecture), pages 244-256, January 1993.*

156. A General Framework for Iteration-Reordering Loop Transformations. Vivek Sarkar and Radhika Thekkath. *Proceedings of the ACM SIGPLAN '92 Conference on Programming Language Design and Implementation, San Francisco, California, pages 175-187, June 1992.*

157. Automatic Partitioning of a Program Dependence Graph into Parallel Tasks. Vivek Sarkar. *IBM Journal of Research and Development, 35(5/6), September/November 1991, pages 779-804.*

158. Optimization of Array Accesses by Collective Loop Transformations. Vivek Sarkar and Guang R. Gao. *Proceedings of the ACM 1991 International Conference on Supercomputing, Cologne, Germany, pages 194-205, June 1991.*

159. Vivek Sarkar and David Cann. POSC – a Partitioning and Optimizing Sisal Compiler. *Proceedings of the ACM 1990 International Conference on Supercomputing, Amsterdam, the Netherlands, pages 148-163, July 1990.*

160. Instruction Reordering for Fork-Join Parallelism. Vivek Sarkar. *Proceedings of the ACM SIGPLAN '90 Conference on Programming Language Design and Implementation, White Plains, New York, pages 322-336, June 1990.*

161. Compact Representations for Control Dependence. Ron Cytron, Jeanne Ferrante, and Vivek Sarkar. *Proceedings of the ACM SIGPLAN '90 Conference on Programming Language Design and Implementation, White Plains, New York, pages 337-351, June 1990.*

162. Determining Average Program Execution Times and their Variance. Vivek Sarkar. *Proceedings of the 1989 SIGPLAN Conference on Programming Language Design and Implementation, Portland, Oregon, pages 298-312, July 1989.*

163. PTRAN Methodologies for the Automatic Detection of Parallelism in FORTRAN Programs: Some Experimental Results (Extended Abstract). Frances Allen, Michael Burke, Philippe Charles, Ron Cytron, Jeanne Ferrante, Vivek Sarkar, and David Shields. *Proceedings of the 4th International Conference on Supercomputing, Santa Clara, California, Volume I, pages 89-93, May 1989.*

164. An Automatically Partitioning Compiler for Sisal. Vivek Sarkar, Steve Skedzielewski, and Patrick Miller. *Proceedings of CONPAR 88, Cambridge University Press, Great Britain, pages 376-383, 1988.*

165. Processor Scheduling Algorithms for Constraint Satisfaction Search Problems. K. S. Natarajan and Vivek Sarkar. *Proceedings of the 1988 International Conference on Parallel Processing, Volume II (Software), pages 140-149, August 1988.*

166. Automatic Discovery of Parallelism: A Tool and an Experiment. Michael Burke, Ron Cytron, Jeanne Ferrante, Wilson Hsieh, Vivek Sarkar, and David Shields. *Proceedings of the ACM*

SIGPLAN Symposium on Parallel Programming: Experience with Applications, Languages, and Systems. New Haven, Connecticut, pages 77-84, July 1988.

167. Synchronization Using Counting Semaphores. Vivek Sarkar. *Proceedings of the ACM 1988 International Conference on Supercomputing, St. Malo, France, pages 627-637, July 1988.*

168. A Framework for Determining Useful Parallelism. Frances Allen, Michael Burke, Ron Cytron, Jeanne Ferrante, Wilson Hsieh, Vivek Sarkar. *Proceedings of the ACM 1988 International Conference on Supercomputing, St. Malo, France, pages 207-215, July 1988.*

169. A simple and efficient implementation approach for single assignment languages. K. Gharachorloo, V. Sarkar, and J. L. Hennessy. *ACM Conference on Lisp and Functional Programming, Snowbird, Utah, pages 259-268, July 1988.*

170. Partitioning Parallel Programs for Macro-Dataflow. Vivek Sarkar and John Hennessy. *ACM Conference on Lisp and Functional Programming, Cambridge, Massachusetts, pages 202-211, August 1986.*

171. Compile-time Partitioning and Scheduling of Parallel Programs. Vivek Sarkar and John Hennessy. *Proceedings of the ACM SIGPLAN '86 Symposium on Compiler Construction, Palo Alto, California, pages 17-26, June 1986.*

11.2 Refereed Workshop Publications

1. Integrating Inter-Node Communication with a Resilient Asynchronous Many-Task Runtime System. Sri Raj Paul, Akihiro Hayashi, Matthew Whitlock, Seonmyeong Bak, Keita Teranishi, Jackson Mayo, Max Grossman, Vivek Sarkar. *ExaMPI20: Proceedings of the 2020 Workshop on Exascale MPI, November 2020.*

2. HOOVER: Leveraging OpenSHMEM for High Performance, Flexible Streaming Graph Applications. Max Grossman, Howard Pritchard, Vivek Sarkar, Steve Poole. *The 3rd Annual Parallel Applications Workshop, Alternatives To MPI+X, November 2020.*

3. A Study of Memory Anomalies in OpenMP Applications. Lechen Yu, Joachim Protze, Oscar Hernandez, Vivek Sarkar. *16th International Workshop on OpenMP (IWOMP), September 2020.*

4. OMPSan: Static Verification of OpenMPs Data Mapping Constructs. Prithayan Barua, Jun Shirako, Whitney Tsang, Jeeva Paudel, Wang Chen, Vivek Sarkar. (Recipient of Best Paper award.) *15th International Workshop on OpenMP (IWOMP), September 2019.*

5. GPUIterator: bridging the gap between Chapel and GPU platforms. Akihiro Hayashi, Sri Raj Paul, Vivek Sarkar. *Proceedings of the ACM SIGPLAN 6th on Chapel Implementers and Users Workshop (CHIUW), co-located with PLDI'19, June 2019.*

6. A Preliminary Study of Compiler Transformations for Graph Applications on the Emu System. Prasanth Chatarasi, Vivek Sarkar. *Proceedings of the Workshop on Memory Centric High Performance Computing (MCHPC, co-located with SC18), November 2018.*

7. A Unified Runtime for PGAS and Event-Driven Programming. Sri Raj Paul, Kun Chen, Akihiro Hayashi, Max Grossman, Vivek Sarkar, Jason Devinney, Bill Carlson. *Fourth International IEEE Workshop on Extreme Scale Programming Models and Middleware (ESPM2,co-located with SC18), November 2018.*

8. Using Polyhedral Analysis to Verify OpenMP Applications are Data Race Free. Fangke Ye, Markus Schordan, Chunhua Liao, Pei-Hung Lin, Ian Karlin, Vivek Sarkar. *2018 IEEE/ACM 2nd International Workshop on Software Correctness for HPC Applications (Correctness, co-located with SC18), November 2018.*

9. A One Year Retrospective on a MOOC in Parallel, Concurrent, and Distributed Programming in Java. Vivek Sarkar, Max Grossman, Zoran Budimlic, Shams Imam. *2018 IEEE/ACM Workshop on Education for High-Performance Computing (EduHPC, co-located with SC18), November 2018.*

10. Implementation and Evaluation of OpenSHMEM Contexts Using OFI Libfabric. *Workshop on OpenSHMEM and Related Technologies. Big Compute and Big Data Convergence, August 2018.*

11. Porting DMRG++ Scientific Application to OpenPOWER. Arghya Chatterjee, Gonzalo Alvarez, E. DAzevedo, Wael Elwasif, Oscar Hernandez, Vivek Sarkar. *International Workshop on OpenPOWER for HPC (IWOPH, co-located with ISC18), June 2018.*

12. Exploration of Supervised Machine Learning Techniques for Runtime Selection of CPU vs. GPU Execution in Java Programs. Gloria Kim, Akihiro Hayashi, Vivek Sarkar. *Fourth Workshop on Accelerator Programming Using Directives (WACCPD), November 2017 (co-located with SC17).*

13. Chapel-on-X: Exploring Tasking Runtimes for PGAS Languages. Akihiro Hayashi, Sri Raj Paul, Max Grossman, Jun Shirako, Vivek Sarkar. *Third IEEE Workshop on Extreme Scale Programming Models and Middleware (ESPM2), November 2017 (co-located with SC17).*

14. A Pluggable Framework for Composable HPC Scheduling Libraries. Max Grossman, Vivek Kumar, Nick Vrvilo, Zoran Budimlic, Vivek Sarkar. *The Seventh International Workshop on Accelerators and Hybrid Exascale Systems (AsHES). May 2017.*

15. Preparing an Online Java Parallel Computing Course. Vivek Sarkar, Max Grossman, Zoran Budimlic and Shams Imam. *7th NSF/TCPP Workshop on Parallel and Distributed Computing Education (EduPar-17). May 2017.*

16. Exploring Compiler Optimization Opportunities for the OpenMP 4.x Accelerator Model on a POWER8+GPU Platform. Akihiro Hayashi, Jun Shirako, Ettore Tiotto, Robert Ho, Vivek Sarkar. *Third Workshop on Accelerator Programming Using Directives (WACCPD, co-located with SC16), November 2016.*

17. Optimized Distributed Work-Stealing. Vivek Kumar, Karthik Murthy, Vivek Sarkar, Yili Zheng. *6th workshop on Irregular Applications: Architectures and Algorithms (IA3), ACM, November 2016.*

18. Fine-grained parallelism in probabilistic parsing with Habanero Java. Matthew Francis-Landau, Bing Xue, Jason Eisner, Vivek Sarkar. *6th workshop on Irregular Applications: Architectures and Algorithms (IA3), ACM, November 2016.*

19. OpenMP as a High-Level Specification Language for Parallelism. Max Grossman, Jun Shirako, Vivek Sarkar. *International Workshop on OpenMP (IWOMP), October 2016.*

20. An Extended Polyhedral Model for SPMD Programs and its use in Static Data Race Detection. Prasanth Chatarasi, Jun Shirako, Martin Kong, Vivek Sarkar. *The 29th International Workshop on Languages and Compilers for Parallel Computing (LCPC), September 2016.*

21. Integrating Asynchronous Task Parallelism with OpenSHMEM. Max Grossman, Vivek Kumar, Zoran Budimlic, Vivek Sarkar. *OpenSHMEM Workshop, August 2016.*

22. Polyhedral Optimizations for a Data-Flow Graph Language. Alina Sbirlea, Jun Shirako, Louis-Noel Pouchet, Vivek Sarkar. *The 28th International Workshop on Languages and Compilers for Parallel Computing (LCPC), September 2015.*

23. Heterogeneous Habanero-C (H2C): A Portable Programming Model for Heterogeneous Processors. Deepak Majeti, Vivek Sarkar. *Workshop on Programming Models, Languages and Compilers for Manycore and Heterogeneous Architectures (PLC), 25th May 2015, Hyderabad, India.* Co-located with IPDPS 2015

24. Polyhedral Transformations of Explicitly Parallel Programs. Prasanth Chatarasi, Jun Shirako, Vivek Sarkar. *5th International Workshop on Polyhedral Compilation Techniques (IMPACT 2015), January 2015.*

25. Selectors: Actors with Multiple Guarded Mailboxes. Shams Imam, Vivek Sarkar. *4th International Workshop on Programming based on Actors, Agents, and Decentralized Control (AGERE! 2014), co-located with OOPSLA 2014, October 2014.*

26. Savina - An Actor Benchmark Suite. Shams Imam, Vivek Sarkar. *4th International Workshop on Programming based on Actors, Agents, and Decentralized Control (AGERE! 2014), co-located with OOPSLA 2014, October 2014.*

27. DFGR: an Intermediate Graph Representation for Macro-Dataflow Programs. Alina Sbirlea, Louis-Noel Pouchet, Vivek Sarkar. *Fourth Workshop on Dataflow Execution Models for Extreme Scale Computing (DFM 2014), co-located with PACT 2014, August 2014.*

28. Exploiting Implicit Parallelism in Dynamic Array Programming Languages. Shams Imam, Vivek Sarkar, David Leibs, Peter B. Kessler. *ACM SIGPLAN International Workshop on Libraries, Languages and Compilers for Array Programming (ARRAY '14), co-located with PLDI'14, June 2014.*

29. A Case for Cooperative Scheduling in X10's Managed Runtime. Shams Imam, Vivek Sarkar. *The 2014 X10 Workshop (X10'14), co-located with PLDI'14, June 2014.*

30. Speculative Execution of Parallel Programs with Precise Exception Semantics on GPUs. Akihiro Hayashi, Max Grossman, Jisheng Zhao, Jun Shirako, Vivek Sarkar. *The 26th International Workshop on Languages and Compilers for Parallel Computing (LCPC), September 2013.*

31. Expressing DOACROSS Loop Dependencies in OpenMP. Jun Shirako, Priya Unnikrishnan, Sanjay Chatterjee, Kelvin Li, Vivek Sarkar. *9th International Workshop on OpenMP (IWOMP), September 2013.*

32. The Flexible Preconditions Model for Macro-Dataflow Execution. Drago Sbrlea, Alina Sbrlea, Kyle B. Wheeler, Vivek Sarkar. *The 3rd Data-Flow Workshop on Execution Models for Extreme Scale Computing (DFM), September 2013.*

33. Compiler-Driven Data Layout Transformation for Heterogeneous Platforms. Deepak Majeti, Rajkishore Barik, Jisheng Zhao, Vivek Sarkar and Max Grossman. *The International Workshop on Algorithms, Models and Tools for Parallel Computing on Heterogeneous Platforms (HeteroPar'2013), August 2013.*

34. Oil and Water can mix! Experiences with integrating Polyhedral and AST-based Transformations. Jun Shirako, Vivek Sarkar. *17th Workshop on Compilers for Parallel Programming (CPC), July 2013.*

35. HadoopCL: MapReduce on Distributed Heterogeneous Platforms Through Seamless Integration of Hadoop and OpenCL. Max Grossman, Mauricio Breternitz, Vivek Sarkar. *2013 International Workshop on High Performance Data Intensive Computing (HPDIC2013)*, co-located with IEEE IPDPS 2013, May 2013.

36. Finish Accumulators: a Deterministic Reduction Construct for Dynamic Task Parallelism. Jun Shirako, Vincent Cave, Jisheng Zhao, Vivek Sarkar. *The 4th Workshop on Determinism and Correctness in Parallel Programming (WoDet), March 2013.*

37. Habanero-Scala: Async-Finish Programming in Scala. Shams Imam, Vivek Sarkar. *The Third Scala Workshop (Scala Days 2012), April 2012.*

38. SCnC: Efficient Unification of Streaming with Dynamic Task Parallelism. Dragos Sbirlea, Jun Shirako, Ryan Newton, Vivek Sarkar. *Proceeding of the Data-Flow Execution Models for Extreme Scale Computing (DFM 2011), in conjunction with PACT 2011, October 2011.*

39. Intermediate Language Extensions for Parallelism. Jisheng Zhao, Vivek Sarkar. *5th Workshop on Virtual Machine and Intermediate Languages (VMIL'11), October 2011.*

40. Dynamic Task Parallelism with a GPU Work-Stealing Runtime System. Sanjay Chatterjee, Max Grossman, Alina Sbirlea, Vivek Sarkar. *2011 Workshop on Languages and Compilers for Parallel Computing (LCPC), September 2011.*

41. CnC-CUDA: Declarative Programming for GPUs. Max Grossman, Alina Simion, Zoran Budimli, Vivek Sarkar. *2010 Workshop on Languages and Compilers for Parallel Computing (LCPC), October 2010.*

42. Modeling and Mapping for Customizable Domain-Specific Computing. Zoran Budimli, Alex Bui, Jason Cong, Glenn Reinman, Vivek Sarkar. *Workshop on Concurrency for the Application Programmer (CAP), co-located with SPLASH 2010, October 2010.*

43. Comparing the Usability of Library vs. Language Approaches to Task Parallelism. Vincent Cav, Zoran Budimli, Vivek Sarkar. *Workshop on Evaluation and Usability of Programming Languages and Tools (PLATEAU), co-located with SPLASH 2010, October 2010.*

44. Hierarchical Place Trees: A Portable Abstraction for Task Parallelism and Data Movement. Yonghong Yan, Jisheng Zhao, Yi Guo, Vivek Sarkar. *Proceedings of the 22nd Workshop on Languages and Compilers for Parallel Computing (LCPC), October 2009.*

45. Declarative Aspects of Memory Management in the Concurrent Collections Parallel Programming Model. Zoran Budimlic, Aparna Chandramowlishwaran, Kathleen Knobe, Geo Lowney, Vivek Sarkar, Leo Treggiari. *Proceedings of DAMP 2009 Workshop (Declarative Aspects of Multicore Programming), co-located with POPL, January 2009.*

46. Multicore Implementations of the Concurrent Collections Programming Model. Zoran Budimlic, Aparna Chandramowlishwaran, Kathleen Knobe, Geo Lowney, Vivek Sarkar, Leo Treggiari. *Proceedings of the 2009 Workshop on Compilers for Parallel Computing (CPC), January 2009.*

47. Minimum Lock Assignment: A Method for Exploiting Concurrency Among Critical Sections. Yuan Zhang, Vugranam Sreedhar, Weirong Zhu, Vivek Sarkar, Guang Gao. *Proceedings of the 21st Workshop on Languages and Compilers for Parallel Computing (LCPC), July 2008.*

48. Array Optimizations for Parallel Implementations of High Productivity Languages Mackale Joyner, Zoran Budimlić, Vivek Sarkar. *13th International Workshop on High-Level Parallel Programming Models and Supportive Environments (HIPS 2008), April 2008.*

49. Language Extensions in Support of Compiler Parallelization. Jun Shirako, Hironori Kasahara, Vivek Sarkar. *Proceedings of the Twentieth Workshop on Languages and Compilers for Parallel Computing (LCPC), October 2007.*

50. Experiences with an SMP Implementation for X10 based on the Java Concurrency Utilities. Rajkishore Barik, Vincent Cave, Christopher Donawa, Allan Kielstra, Igor Peshansky, Vivek Sarkar. *Workshop on Programming Models for Ubiquitous Parallelism (PMUP), held in conjunction with PACT 2006, Sep 2006.*

51. An Experiment in Measuring the Productivity of Three Parallel Programming Languages. Kemal Ebcioglu, Vivek Sarkar, Tarek El-Ghazawi, John Urbanic. *HPCA Workshop on Productivity and Performance in High-End Computing (P-PHEC 2006), held in conjunction with HPCA 2006, Feb 2006.*

52. X10: an Experimental Language for High Productivity Programming of Scalable Systems. Kemal Ebcioglu, Vijay Saraswat, Vivek Sarkar. *HPCA Workshop on Productivity and Performance in High-End Computing (P-PHEC 2005), held in conjunction with HPCA 2005, Feb 2005.*

53. X10: Programming for Hierarchical Parallelism and NonUniform Data Access. Kemal Ebcioglu, Vijay Saraswat, Vivek Sarkar. *Third International Workshop on Language Runtimes (LaR 2004): Impact of Next Generation Processor Architectures On Virtual Machine Technologies, held in conjunction with OOPSLA 2004, Oct 2004.*

54. Application development productivity challenges for high-end computing. Vivek Sarkar, Kemal Ebcioglu, Clay Williams. *HPCA Workshop on Productivity and Performance in High-End Computing (P-PHEC 2004), held in conjunction with HPCA 2004, Feb 2004.*

55. Program analysis for safety guarantees in a Java virtual machine written in Java. Jan-Willem Maessen, Vivek Sarkar, David Grove. *Proceedings of the ACM PASTE 2001 workshop, June 2001, pp. 62-65.*

56. A Comparative Study of Static and Dynamic Heuristics for Inlining. M.Arnold, S.Fink, V.Sarkar, P.F.Sweeney. *ACM Dynamo '00 workshop, held in conjunction with POPL '00, Jan 2000.*

57. Enhanced Parallelization via Analyses and Transformations on Array SSA Form. K.Knobe, V.Sarkar. *Workshop on Compilers for Parallel Computers (CPC), Jan 2000.*

58. Dependence Analysis for Java. Craig Chambers, Igor Pechtchanski, Vivek Sarkar, Harini Srinivasan, Mauricio Serrano. *Twelfth Workshop on Languages and Compilers for Parallel Computing (LCPC), August 1999.*

59. Static Scheduling with Communication Weights — Theory and Practice. Vivek Sarkar. *Workshop on Scheduling Algorithms for Parallel/Distributed Computing — From Theory to*

Practice. Held in conjunction with the ACM 1999 International Conference on Supercomputing (ICS), June 1999.

60. Efficient and Precise Modeling of Exceptions for the Analysis of Java Programs. Jong-Deok Choi, David Grove, Michael Hind, Vivek Sarkar. *Proceedings of the 1999 ACM SIGPLAN Workshop on Program Analysis for Software Tools and Engineering (PASTE), September, 1999.*

61. Jalapeño — a Compiler-Supported Java Virtual Machine for Servers. Bowen Alpern, Anthony Cocchi, Derek Lieber, Mark Mergen, Vivek Sarkar. *ACM SIGPLAN 1999 Workshop on Compiler Support for System Software (WCSSS), May 1999.*

62. Optimized Execution of Fortran 90 Array Language on Symmetric Shared-Memory Multiprocessors. Vivek Sarkar. *Eleventh Workshop on Languages and Compilers for Parallel Computing (LCPC), August 1998.*

63. Loop Transformations for Hierarchical Parallelism and Locality. Vivek Sarkar. *Workshop on Languages, Compilers and Run-Time Systems for Scalable Computers, May 1998. Proceedings published as Springer-Verlag LNCS 1511.*

64. Analysis and Optimization of Explicitly Parallel Programs using the Parallel Program Graph Representation. Vivek Sarkar. *Tenth Workshop on Languages and Compilers for Parallel Computing, Minneapolis, Minnesota, August 1997.*

65. The Raw Compiler Project. A. Agarwal, S. Amarasinghe, R. Barua, S. Devabhaktuni, M. Frank, W. Lee, V. Sarkar, and M. Taylor. *Second SUIF Compiler Workshop, Stanford, California, August 1997.*

66. Locality Analysis for Distributed Shared-Memory Multiprocessors. Vivek Sarkar, Guang R. Gao, and Shaohua Han. *Ninth Workshop on Languages and Compilers for Parallel Computing, Santa Clara, California, August 1996. Springer-Verlag Lecture Notes in Computer Science, 1239, pages 20-40, 1997.*

67. Data Consistency in Distributed Memory Systems (Extended Abstract). Vivek Sarkar. *Proceedings of the 1995 ICPP Workshop on Challenges for Parallel Processing, Oconomowoc, Wisconsin, pages 124-132, August 1995.*

68. Beyond the Data Parallel Paradigm: Issues and Options. Guang R. Gao, Vivek Sarkar, Lelia A. Vazquez *Proceedings of the Working Conference on Massively Parallel Programming Models, Berlin, Germany, September 1993.*

69. Parallel Program Graphs and their Classification. Vivek Sarkar and Barbara Simons. *Sixth International Workshop on Languages and Compilers for Parallel Computing, Portland, Oregon, August 1993. Springer-Verlag Lecture Notes in Computer Science, 768, pages 633-655, 1993.*

70. A Concurrent Execution Semantics for Parallel Program Graphs and Program Dependence Graphs. Vivek Sarkar. *Fifth Workshop on Languages and Compilers for Parallel Computing, Yale University, August 1992. Springer-Verlag Lecture Notes in Computer Science, 1993.*

71. Collective Loop Fusion for Array Contraction. G. R. Gao, R. Olsen, V. Sarkar, and R. Thekkath. *Proceedings of the Fifth Workshop on Languages and Compilers for Parallel Computing, Yale University, August 1992. Springer-Verlag Lecture Notes in Computer Science, 1993.*

72. On Estimating and Enhancing Cache Effectiveness. Jeanne Ferrante, Vivek Sarkar, and Wendy Thrash. *Proceedings of the Fourth International Workshop on Languages and Compilers for Parallel Computing, Santa Clara, California, August 1991. Springer-Verlag Lecture Notes in Computer Science, 589, pages 328-343, 1991.*

73. Experiences Using Control Dependence in PTRAN. Ron Cytron, Jeanne Ferrante, and Vivek Sarkar. *Second Workshop on Languages and Compilers for Parallel Computing, U. Illinois, August 1989. Edited by D. Gelernter, A. Nicolau, and D. Padua, MIT Press, pages 186-212, 1990.*

11.3 Books, Book Chapters, and Edited Volumes

1. General Chair, Proceedings of the Eighteenth International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS 2013), March 2013.

2. Workshop Co-chair, 23rd Workshop on Languages and Compilers for Parallel Computing (LCPC), October 2010. Proceedings published in Lecture Notes in Computer Science, Vol. 6548.

3. Program Chair, 2009 ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming (PPoPP), February 2009.

4. Guest Editor, IBM Systems Journal special issue on Open Source Software, June 2005.

5. General Chair, Proceedings of the 2005 ACM SIGPLAN conference on Programming Language Design and Implementation (PLDI), June 2005.

6. Program Co-chair, Proceedings of the Twelfth International Conference on Parallel Architectures and Compilation Techniques (PACT), September 2003.

7. Program Chair, Proceedings of the 1994 ACM SIGPLAN conference on Programming Language Design and Implementation (PLDI), June 1994.

8. The PTRAN Parallel Programming System. Vivek Sarkar. *In Parallel Functional Programming Languages and Compilers, edited by B. Szymanski, ACM Press Frontier Series, pages 309-391, 1991.*

9. Partitioning and Scheduling Parallel Programs for Multiprocessors. Vivek Sarkar. *MIT Press, Cambridge, Massachusetts, 1989 (201 pages). In the series, Research Monographs in Parallel and Distributed Computing.*

12 Patents

1. Automatic verification of determinism for parallel programs. Martin Vechev, Eran Yahav, Raghavan Raman, and Vivek Sarkar. *Filed March 2011, issued as US patent 9069893 in June 2015.*

2. Method and apparatus for efficient and precise datarace detection for multithreaded object-oriented programs. Jong-Deok Choi, Keunwoo Lee, Robert O'Callahan, Vivek Sarkar, and Manu Sridharan. *Filed February 2009, issued as US patent 8464223 in April 2009.*

3. Method and apparatus for efficient and precise datarace detection for multithreaded object-oriented programs. Jong-Deok Choi, Keunwoo Lee, Robert O'Callahan, Vivek Sarkar, and Manu Sridharan. *Filed June 2002, issued as US patent 7516446 in April 2009.*

4. Static detection of a datarace condition for multithreaded object-oriented applications. Jong-Deok Choi, Alexey Loginov, and Vivek Sarkar. *Filed January 2002, issued as US patent 7469403 in December 2008.*

5. Retargeting optimized code by matching tree patterns in directed acyclic graphs. Vivek Sarkar, Mauricio J. Serrano, and Barbara Simons. *Filed December 1998, issued as US patent 6292938 in September 2001.*

6. Method of, System for, and Computer Program Product for performing Weighted Loop Fusion by an Optimizing Compiler. Nimrod Megiddo and Vivek Sarkar. *Filed June 1997, issued as US patent 6058266 in October 2000.*

7. System, method, and program product for loop instruction scheduling hardware lookahead. Barbara Simons and Vivek Sarkar. *Filed June 1997, issued as US patent 6044222 in March 2000.*

8. A Method of, System for, and Computer Program Product for providing Quick Fusion in WHERE constructs. Dz-Ching Ju, John Ng, and Vivek Sarkar. *Filed July 1997, issued as US patent 6041181 in March 2000.*

9. Method and system for generating compact code for the loop unrolling transformation. Khoa Nguyen and Vivek Sarkar. *Filed July 1997, issued as US patent 6035125 in March 2000.*

10. Method of, System for, and Computer Program Product for Minimizing Loop Execution Time by Optimizing Block/Tile Sizes. Nimrod Megiddo and Vivek Sarkar. *Filed July 1997, issued as US patent 5953531 in September 1999.*

11. System, method, and program product for instruction scheduling in the presence of hardware lookahead accomplished by the rescheduling of idles lots. Barbara Simons and Vivek Sarkar. *Filed June 1996, issued as US patent 5887174 in March 1999.*

12. Method and System for Efficient Identification of Private Variables in Program Loops by an Optimizing Compiler. Vivek Sarkar. *Filed December 1995, issued as US patent 5790859 in August 1998.*

13 Software Artifacts (Selected)

- Co-led software releases of Habanero-C programing system and Open Computing Runtime (OCR) since 2012.
- Led Habanero-Java (HJ) releases starting in 2009 (http://habanero.rice.edu/hj). HJ is used by multiple institutions for research and teaching. At Rice, this software is used for laboratory and programming assignments in COMP 322, and for research in the Habanero Extreme Scale Software research project.
- Co-led first open source release of X10 in December 2006 (http://x10-lang.org).
- Led first open source of Jikes Research Virtual Machine in October 2001 (http://jikesrvm.org/).
- Led development of ASTI optimizer component, which has been shipping as part of IBM's XL Fortran product compilers since 1996.
- Contributed to development of PTRAN research prototype compiler (led by Fran Allen).

14 Personal Information

- US citizen, Overseas Citizen of India, married with two children.
- Extra-curricular interests: theatre, hiking, horseback riding, classical music (violin).
- Work address: Georgia Institute of Technology. Klaus Advanced Computing Building, Room 2332, 266 Ferst Dr NW Atlanta GA 30332. Phone: +1-404-385-5989.
- Personal email: vsarkar@me.com.