

Max Grossman

7pod.tech, github.com/agrippa, max@7pod.tech, (914)-907-1518

EDUCATION

Rice University

Doctor of Philosophy in Computer Science January 2015-January 2017

Master of Science in Computer Science August 2012-May 2013

Bachelor of Science in Computer Science August 2008-May 2012

Research Advisor: Dr. Vivek Sarkar, Habanero Multicore Software Group

Thesis Title: Productive Programming Systems for Heterogeneous Supercomputers

Graduate Fellow to the Computing Research Association's Committee on Education

Rice Center for Engineering Leadership, Presentation Coach (2015-2017)

Rice Center for Teaching Excellence, Graduate Liaison (2015-2017)

2012-2013 Ken Kennedy Institute Oil & Gas HPC Fellowship

2009 Brown Undergraduate Research Internship

NCAA DI Track & Field (400m hurdles, 4x400m relay)

EXPERIENCE

7pod Technologies

Feb 2017-Present

Co-Founder and CEO

Direct the development of strategic partnerships and HPC software at a small, agile HPC software/hardware shop focused on serving full stack high-performance computing solutions.

Rice University Dept of Computer Science

Feb 2017-Present

Postdoctoral Researcher, Habanero Multicore Software Group

Representative to the OpenSHMEM Standards Committee

Work focuses on the development of high performance, high productivity programming models and runtimes through 1) tighter integration of computation and communication runtimes, and 2) transparent acceleration of managed runtime systems.

Rice University Dept of Computer Science Sept 2008-June 2013, Jan 2015-Feb 2017

Research Assistant, Habanero Multicore Software Group Advisor: Dr. Vivek Sarkar

Research projects concentrate on parallel programming models and tools for distributed, multi-core, and heterogeneous architectures. Additional work includes optimizing various large production applications in medical imaging, genomics, geophysics, astrophysics, and machine learning across a variety of hardware and software platforms. Developed the autograding infrastructure, project source code, project descriptions, and quizzes for the Coursera specialization on "Parallel, Concurrent, and Distributed Programming in Java". Selected projects include:

1. HiPER, a pluggable programming framework and runtime for exascale systems:
https://github.com/habanero-rice/hcpp/tree/resource_workers
2. SWAT, automatic accelerator offload of user-written Apache Spark kernels:
<https://github.com/agrippa/spark-swat>
3. CHIMES, efficient checkpoint-restart as a software development tool:
<https://github.com/agrippa/chimes>
4. HYDOS0, a framework for performance prediction of data-parallel kernels on heterogeneous architectures

NSF East Asia and Pacific Summer Institutes
Fellow

June 2016-August 2016
Mentor: Professor Hironori Kasahara

Led research into performance prediction techniques for heterogeneous architectures as a member of the Kasahara Laboratory at Waseda University in Tokyo, Japan. Drove continuing international collaborations between Waseda, Meisei, and Rice University HPC researchers.

Delphix
Member of Technical Staff

June 2013-January 2015

Work at Delphix includes 1) feature and performance improvements to the OpenZFS filesystem, 2) functionality improvements to a data replication feature in the core database virtualization product, 3) design and implementation of a challenge-response system for OS/shell-level security of the core product based on IETF RFC 6287, 4) unit test and integration test design and implementation for Java, Javascript, Python applications, and 5) launching and growing Delphix's university recruiting program at Rice University.

Repsol USA
Research Intern

January 2011-December 2014
Managers: Mauricio Araya Polo & Gladys Gonzalez

Work focuses on optimizing geophysical applications. Initial development used heterogeneous processors as the primary avenue of performance improvement for a legacy geophysics application. Later work included intelligent scheduling algorithms, hardware adaptive algorithms, and heterogeneous optimization.

Oracle Labs
Intern, VM Research Group

May 2012-August 2012
Manager: Misha Dmitriev

Investigated the use of both static bytecode analysis and profiler metrics from parallel Java applications in accurately locating and diagnosing parallel bottlenecks. The results of that analysis was used to suggest code changes that are significantly beneficial for performance. Completed novel work on automatically manipulating Java bytecode based on the suggested changes.

NASA Jet Propulsion Laboratory
Intern, Space and Astrophysical Plasmas

June 2011-June 2012
Managers: Karen Willacy & Mark Allen

Used MPI to prepare a dynamical and chemical code base for more computationally demanding algorithms and larger data sets. Achieved scalable performance after applying code transformations to create an optimally parallel program. Responsible for advising and educating astrophysicists on the parallelization of their scientific models.

National Security Agency (NSA)
TS//SI Clearance with Polygraph and Psychological Exam.

May 2012-August 2012

Continued development of a large scale piece of software (>50,000,000 lines). Implemented new and useful functionality. Utilized C Programming, SVN Repository Management, modern wireless telecommunications protocols, Wiki Management, and the Linux command line.

BOOKS

J. Cheng, M. Grossman, T. McKercher. Professional CUDA C Programming. Wiley Publishing. 2014.

INVITED TALKS Lecturer for the Accelerator Day at the Ken Kennedy Institute's HPC Bootcamp
<http://bootcamp.rice.edu/hpc-bootcamp> 2016, 2017

Secondary Session Leader for the SC BoF on Best Practices in Mentoring Undergraduate Research in Supercomputing 2016

SELECTED PUBLICATIONS & TALKS T. Cogumbreiro, R. Surendran, F. Martins, V. Sarkar, V. Vasconcelos, M. Grossman. "Deadlock Avoidance in Parallel Programs with Future Tasks". OOPSLA 2017. October 2017.

V. Sarkar, M. Grossman, Z. Budimlic, S. Imam. "Preparing an Online Java Parallel Computing Course". 7th NSF/TCPP Workshop on Parallel and Distributed Computing Education (EduPar-17). May 2017.

M. Grossman, H. Pritchard, V. Sarkar, Z. Budimlic, N. Graham. "Performance of a Task-Parallel PGAS Programming Model using OpenSHMEM and UCX". OpenFabrics Alliance Workshop 2017. March 2017.

M. Grossman, V. Kumar, N. Vrvilo, Z. Budimlic, V. Sarkar. "A Pluggable Framework for Composable HPC Scheduling Libraries". The Seventh International Workshop on Accelerators and Hybrid Exascale Systems (AsHES), co-located with IPDPS. May 2017.

M. Grossman, H. Pritchard, Z. Budimlic, V. Sarkar. "Graph500 on OpenSHMEM: Using A Practical Survey of Past Work to Motivate Novel Algorithmic Developments". Los Alamos National Laboratory Technical Report. December 2016.

M. Grossman, M. Aziz, H. Chi, A. Tibrewal, S. Imam, V. Sarkar. "Pedagogy and Tools for Teaching Parallel Computing at the Sophomore Undergraduate Level". Journal of Parallel and Distributed Computing Special Issue on Parallel, Distributed, and High Performance Computing Education. 2016.

Y. Peng, M. Grossman, V. Sarkar. "Static Cost Estimation for Data Layout Selection on GPUs." 7th International Workshop on Performance Modeling, Benchmarking, and Simulation of High Performance Computer Systems. November 2016.

M. Grossman, V. Kumar, Z. Budimlic, V. Sarkar. "Integrating Asynchronous Task Parallelism with OpenSHMEM." OpenSHMEM Workshop. August 2016.

M. Grossman, J. Shirako, V. Sarkar. "OpenMP as a High-Level Specification Language for Parallelism." International Workshop on OpenMP (IWOMP). October 2016.

M. Grossman, M. Araya-Polo, F. Frank, F.O. Alpak, V. Sarkar. "Analysis of Sparse Matrix-Vector Multiply for Large Sparse Linear Systems." European Conference on the Mathematics of Oil Recovery. August 2016.

M. Grossman, V. Sarkar. "SWAT: A Programmable, In-Memory, Distributed, High-Performance Computing Platform." International ACM Symposium on High-Performance Parallel and Distributed Computing (HPDC). May 2016.

M. Grossman, M. Araya-Polo, F. Frank, C. Thiel, V. Sarkar. "A Survey of Sparse Matrix-Vector Multiply Performance on Large Matrices". Rice Oil & Gas High Performance Computing Workshop. March 2016.

- M. Grossman, V. Sarkar. “Efficient Checkpointing of Multi-Threaded Applications as a Tool for Debugging, Performance Tuning, and Resiliency.” IEEE International Parallel and Distributed Processing Symposium (IPDPS). May 2016.
- M. Grossman, M. Araya-Polo. “Efficient Static and Dynamic Memory Management Techniques for Multi-GPU Systems.” Workshop on Runtime Systems for Extreme Scale Programming Models and Architectures (RESPA). November 2015.
- M. Grossman, M. Araya-Polo. “Distributed, Heterogeneous Scheduling Techniques Motivated by Production Geophysical Applications.” Workshop on Many-Task Computing on Clouds, Grids, and Supercomputers (MTAGS). November 2015.
- M. Aziz, H. Chi, Anant Tibrewal, M. Grossman, V. Sarkar. “Auto-Grading for Parallel Programs.” Workshop on Education for High-Performance Computing (EduHPC). November 2015.
- M. Grossman, S. Imam, V. Sarkar. “HJ-OpenCL: Reducing the Gap Between the JVM and Accelerators.” 2015 International Conference on Principles and Practices of Programming on the Java Platform (PPPJ). September 2015.
- M. Grossman, M. Breternitz, V. Sarkar. “HadoopCL2: Motivating the Design of a Distributed, Heterogeneous Programming System With Machine-Learning Applications.” IEEE Transactions on Parallel and Distributed Systems.
- M. Grossman. “Efficient, Automatic Application Checkpointing as a Powerful Tool for CUDA Development.” NVIDIA GPU Technology Conference (GTC) 2015.
- X. Li, M. Grossman, D. Kaeli. “Mahout on Heterogeneous Clusters Using HadoopCL.” 2nd Workshop on Parallel Programming for Analytics Applications. February 2015.
- M. Grossman, G. Gonzalez, M. Araya-Polo. “Efficient Memory and Bandwidth Management for Industrial Strength Kirchhoff Migration.” NVIDIA GPU Technology Conference (GTC) 2014.
- M. Grossman, M. Breternitz, V. Sarkar. “Characterizing APU Performance in HadoopCL on Heterogeneous Distributed Platforms.” AMD Developer Summit 2013.
- A. Hayash, M. Grossman, J. Zhao, J. Shirako, V. Sarkar. “Speculative Execution of Parallel Programs with Precise Exception Semantics on GPUs.” 26th International Workshop on Languages and Compilers for Parallel Computing (LCPC). September 2013.
- A. Hayashi, M. Grossman, J. Zhao, J. Shirako, V. Sarkar. “Accelerating Java Programs with OpenCL Generation.” 2013 International Conference on Principles and Practices of Programming on the Java Platform. September 2013.
- M. Grossman, M. Breternitz, V. Sarkar. “HadoopCL: MapReduce on Distributed Heterogeneous Platforms Through Seamless Integration of Hadoop and OpenCL.” 2013 International Workshop on High Performance Data Intensive Computing (HPDIC). May 2013.
- M. Grossman, M. Araya-Polo, G. Gonzalez. “Hybrid Implementation of 3D Kirchoff Migration.” NVIDIA GPU Technology Conference (GTC) 2013.

S. Chatterjee, Z. Budimli, V. Cav, M. Chabbi, M. Grossman, S. Tarlar, Y. Yan, V. Sarkar. “Integrating Asynchronous Task Parallelism with MPI.” 27th IEEE International Parallel & Distributed Processing Symposium (IPDPS).

M. Grossman, V. Sarkar, M. Breternitz. “Accelerating Medical Imaging Applications Using OpenCL and APUs.” AMD Fusion Developer Summit. June 2012.

M. Grossman, M. Araya-Polo. “GPU-based Travel Time Computation.” SIAM Conference on Parallel Processing for Scientific Computing. February 2012.

S. Chatterjee, M. Grossman, A. Sbirlea, V. Sarkar. “Dynamic Task Parallelism with a GPU Work-Stealing Runtime System.” 24th International Workshop on Languages and Compilers for Parallel Computing. September 2011. **Awarded Best Student Paper**

M. Grossman, A. Sbirlea, Z. Budimlic, V. Sarkar “CnC-CUDA: Declarative Programming for GPUs.” 23rd International Workshop on Languages and Compilers for Parallel Computing. October 2010.

Y. Yan, M. Grossman, V. Sarkar “JCUDA: A Programmer Friendly Interface for Accelerating Java Programs with CUDA.” 2009 Euro-Par Conference Series. August 2009.