

RESEARCH INTERESTS

Operating systems and computer architecture, virtualization and high-performance computing.

EDUCATION

Georgia Institute of Technology

PhD in Computer Science

Advisor: Karsten Schwan

GPA 4.0/4.0

Atlanta, GA

Current, since Aug. 2009

Rochester Institute of Technology

BS in Computer Science, minor in German

GPA 3.77/4.0

Rochester, NY

Sept. 2004 - May 2009

RESEARCH EXPERIENCE

Georgia Institute of Technology

Fine-grained Resource Management on Heterogeneous High-performance Clusters

Atlanta, GA

Jan. 2010 - present

GPGPUs have proven to be advantageous for increasing application scalability both in the HPC and enterprise domains. This has resulted in an increase in the array of programming languages and range of physical compute capabilities of current hardware. Yet applications' scalability and portability remain limited with respect to both their degree of customization and the physical limitations of compute nodes to contain any number and composition of accelerators. This research defines the notion of a *GPGPU assembly* for applications on both bare-metal compute nodes as well as within Xen virtual machines in high-performance clusters, presenting to applications a logical configuration of hardware "just right" for their needs—CPUs and a set of local and remote GPGPUs. GPGPUs are provided transparently as locally-available devices, easing programmability and portability. We characterize workloads to best match them with available GPGPUs and employ techniques such as dynamic resource monitoring of CPU, network and memory usage, latency and throughput measurements, as well as instrumentation of binary CUDA kernels to obtain GPGPU-level usage to enable higher-level *mapping* policies. This research provides a dynamic runtime system for evaluating more intelligent management methods of GPGPU clusters, examining global scheduling policies, admission control and dynamic retargeting of execution streams. *Current work.*

PUBLICATIONS/PRESENTATIONS

- "Shadowfax: Dynamically Composed GPGPU Assemblies"; **Alexander Merritt**, Vishakha Gupta, Abhishek Verma, Ada Gavrilovska, Karsten Schwan; *VTDC'11* (paper), San Jose, CA, June 2011.
- "Minimal-overhead Virtualization of a Large Scale Supercomputer"; John Lange, Kevin Pedretti, Peter Dinda, Chang Bee, Patrick Bridges, Philip Soltero, **Alexander Merritt**; *VEE'11* (paper), Newport Beach, CA, March 2011.
- "Techniques for Managing Data Distribution in NUMA Systems"; **Alexander Merritt**, Kevin Pedretti; *SC'10* (poster), New Orleans, LA, November 2010.
- "Techniques for Managing Data Distribution in NUMA Systems"; **Alexander Merritt**, Kevin Pedretti; *Sandia Computer Science Research Institute 2010 Summer Proceedings*, Albuquerque, NM, August 2010.

PROFESSIONAL EXPERIENCE

Intel Corporation

Summer Intern

- Researching and developing operating system abstractions to support client platform composition.

Hillsboro, OR

May 2012 - Aug. 2012

Sandia National Laboratories, CSRI

Summer Intern

- Ported Xen & Linux IOMMU code to the Kitten light-weight kernel to enable access to GPGPUs from within Palacios-hosted VMs.

Albuquerque, NM

May 2011 - Aug. 2011

Sandia National Laboratories, CSRI

Summer Intern

- Evaluated current state of memory distribution techniques in the Linux operating system using

Albuquerque, NM

May 2010 - Aug. 2010

- high-performance benchmarks on NUMA hardware, and Pin to visualize memory access phases.
- Added support for large pages in the Palacios virtual machine monitor.

Goodrich Corporation, ISR Systems

Chelmsford, MA

Summer Intern

June 2008 - Nov. 2008

- Implemented a C++ simulator to unit-test an image-processing single-board computer running VxWorks.
- Designed a pixel data generator within a simulator to cope with memory constraints by the hardware.

Rochester Software Associates, WebCRD

Rochester, NY

Winter Intern

Nov. 2007 - Feb. 2008

- Implemented a customer-requested feature in an online printing management system using Java.

Harris Corporation, RF Comm. Division

Rochester, NY

Summer Intern

Nov. 2006 - Aug. 2007

- Executed and updated manual/automated tests on military radios.
- Investigated toolchain environment for enabling the preservation of radio memory state, post deployment. This would provide access to a stack trace for debugging system failures.

ACADEMIC RECOGNITION

- President's Fellowship, Georgia Institute of Technology. 2009-present
- Dean's List, Golisano College of Computing and IS, Rochester Institute of Technology. 2004-2009

TEACHING EXPERIENCE

Georgia Institute of Technology

Atlanta, GA

Teaching Assistant

2009 - 2010

- High-Performance Computing Architecture. Wrote a simulator in C for a dynamic CPU cache and provided guidance for students implementing the project.
- Introduction to Database Systems. Held office hours, provided project feedback, grading.

Rochester Institute of Technology

Rochester, NY

Private Tutor - German Language

Dec. 2006 - Feb. 2007

Assisted one student with the reading, writing and spoken word of the German language.

SKILLS

Software Platforms	GNU/Linux, Xen
Hardware Platforms	x86/x86_64, NVIDIA GPGPU, Multi-core & NUMA
Languages	C, Shell scripting, Python, Java, C++
Libraries	POSIX IPC, NVIDIA CUDA, BSD Sockets
Tools	git, Pin, R (plotting)

COURSE PROJECTS

- Created a C API to export dynamic CUDA PTX instrumentation within a cluster monitoring tool.
- Modified a custom user-level thread scheduler to mimic the Xen credit scheduler.
- Implemented a simulator in Pin for a pipelined, out-of-order execution SMT CPU.
- Developed a parallel ray-tracer in CUDA with another student.
- Parallelized a Sobel image edge-detection algorithm using OpenMP and DevIL.
- Implemented and integrated a simplified DMA IDE disk driver into a bootstrap OS framework.

COURSES

Distributed Computing, High-Performance Computer Architectures, Advanced Operating Systems, Computer Networking, Real-Time and Embedded Systems; Systems Programming, Computer Graphics.

REFERENCES

- **Karsten Schwan**, schwan@cc.gatech.edu. Regent's Professor and CERCS Director
College of Computing at the Georgia Institute of Technology, Atlanta, GA.
- **Kevin Pedretti**, ktpedre@sandia.gov. Senior Member of Technical Staff, Scalable System Software
Sandia National Laboratories, P.O. Box 5800, Albuquerque, NM 87185.
- **Hyesoon Kim**, hyesoon@cc.gatech.edu. Assistant Professor
College of Computing at the Georgia Institute of Technology, Atlanta, GA.