

EMPLOYMENT

- Illinois Institute of Technology* 2016–Present Assistant Professor, ILLINOIS INSTITUTE OF TECHNOLOGY Department of Computer Science
Chicago, IL
Assistant Professor in the Department of Computer Science.
- Northwestern University* 2010–2016 Ph.D. Student, NORTHWESTERN UNIVERSITY
Department of Electrical Engineering and Computer Science
Evanston, IL
Conducted research in unconventional and experimental computer systems, with an emphasis on operating systems and high-performance computing.
- VMWare, Inc.* Summer 2013 Research Intern, VMWARE, INC.
Proactive Distributed Resource Management Team
Palo Alto, CA
Investigated the ability to leverage application communication patterns in parallel codes to implement proactive resiliency in a virtualized environment (particularly for VMWare vSphere).
Reference: Rean GRIFFITH rean@caa.columbia.edu
- Fujitsu Ltd.* Aug-Sep 2012 Technical Computing Intern, FUJITSU LTD.
Technical Computing Solutions Unit
Chiba, Japan
Tested, packaged, and installed the Fujitsu cross-compiler toolkit for the PRIMEHPC FX10 Supercomputer on access nodes. Developed test-suite of hybrid parallel applications (MPI/OpenMP/FFTW) aimed at customers developing cross-compiled programs for the PRIMEHPC FX10.
Reference: Shinya FUKUMOTO fukumoto.shinya@jp.fujitsu.com
- Sandia National Laboratories* Summer 2012 Graduate Technical Research Intern, SANDIA LABS
Scalable Systems Software Unit
Albuquerque, NM
Ported the Palacios Virtual Machine Monitor to the Cray XK6. Developed a novel, RDMA-based high-performance networking component within the Palacios VMM to mitigate network virtualization overhead in HPC applications.
Reference: Kevin PEDRETTI ktpedre@sandia.gov
- USAA* Summer 2008 Database Support Intern, USAA
Database Group
San Antonio, TX
Helped consolidate data from several independent legacy data warehouses (Oracle, DB2, SQL Server, MySQL). Helped consolidate data from several independent legacy data warehouses (Oracle, DB2, SQL Server, MySQL). Migrated data reporting mechanisms from hard-coded scripts and Excel sheets to a fully customizable reporting system using Microsoft products, allowing management to quickly and efficiently drill down on important data-usage statistics. Reference: Ranjith RAGHUNATH ranjith_nath@hotmail.com

PUBLICATIONS

Refereed Conference Papers

- ICAC 2017 **K.C. Hale**, C. Hetland, and P. Dinda. Multiverse: Easy Conversion of Runtime Systems into OS Kernels via Automatic Hybridization. *Proceedings of the 14th International Conference on Autonomic Computing*, July, 2017.
- HPDC 2016 **K.C. Hale**, C. Hetland, and P. Dinda. Automatic Hybridization of Runtime Systems. *Proceedings of the 25th International ACM Symposium on High-performance Parallel and Distributed Computing*, June, 2016.
- VEE 2016 **K.C. Hale** and P. Dinda. Enabling Hybrid Parallel Runtimes Through Kernel and Virtualization Support. *Proceedings of the 12th ACM SIGPLAN/SIGOPS International Conference on Virtual Execution Environments*, April, 2016.
- HPDC 2015 **K.C. Hale** and P. Dinda. A Case for Transforming Parallel Runtimes into Operating System Kernels. *Proceedings of the 24th International ACM Symposium on High-performance Parallel and Distributed Computing*, June, 2015.
- ROSS 2014 M. Swiech, **K.C. Hale**, and P. Dinda. VMM Emulation of Intel Hardware Transactional Memory. *Proceedings of the 4th International Workshop on Runtime and Operating Systems for Supercomputers*, June, 2014.
- HPDC 2014 L. Xia, **K.C. Hale**, and P. Dinda. ConCORD: Easily Exploiting Memory Content Redundancy Through the Content-aware Service Command. *Proceedings of the 23rd ACM Symposium on High-performance Parallel and Distributed Computing*, June, 2014.
- ICAC 2014 **K.C. Hale** and P. Dinda. Guarded Modules: Adaptively Extending the VMM's Privilege Into the Guest. *Proceedings of the 11th International Conference on Autonomic Computing*, June, 2014.
- ICAC 2012 **K.C. Hale**, L. Xia, and P. Dinda. Shifting GEARS to Enable Guest-context Virtual Services. *Proceedings of the 9th International Conference on Autonomic Computing*, September, 2012.
- NoCArc 2009 **K.C. Hale**, B. Grot, and S. Keckler. Segment Gating for Static Energy Reduction in Networks-On-Chip. *Proceedings of the 2nd International Workshop on Network-on-Chip Architectures*, December, 2009.

Non-overlapping Technical Reports

- March 2016 **K.C. Hale** and P. Dinda. Pushing Software Events to the Hardware Limit. Technical Report NU-EECS-16-02, Department of Electrical Engineering and Computer Science, Northwestern University, March, 2016.
- April 2014 **K.C. Hale** and P. Dinda. Details of the Case for Transforming Parallel Runtimes into Operating System Kernels. Technical Report NU-EECS-15-01, Department of Electrical Engineering and Computer Science, Northwestern University, April, 2014.
- November 2011 J. Lange, P. Dinda, **K.C. Hale**, and L. Xia. An Introduction to the Palacios Virtual Machine Monitor—Version 1.3. Technical Report NU-EECS11-10, Department of Electrical Engineering and Computer Science, Northwestern University, November, 2011.

Miscellaneous Posters and Talks

- GCASR 2017 P. Nookala, I. Raicu, P. Dinda, and **K.C. Hale**. Performance Analysis of Queue-based Data Structures. Poster at *the 6th Annual Greater Chicago Area Systems Research Workshop*, April, 2017.
- ROSS 2016 **K.C. Hale** and P. Dinda. Accelerating Asynchronous Events for Hybrid Parallel Runtimes. Invited talk at *the 6th International Workshop on Runtime and Operating Systems for Supercomputers*, June, 2016.
- GCASR 2016 **K.C. Hale** and P. Dinda. Multiverse: Automatic Hybridization of Parallel Runtime Systems. At *the 5th Annual Greater Chicago Area Systems Research Workshop*, April, 2016.
- HPDC 2015 **K.C. Hale** and P. Dinda. A Case for Transforming Parallel Runtimes into Operating System Kernels. At *the 23rd ACM Symposium on High-performance Parallel and Distributed Computing*, June, 2015.
- GCASR 2015 G. Tziantzioulis, **K.C. Hale**, B. Pashaj, N. Hardavellas, and P. Dinda. SeaFire: Specialized Computing on Dark Silicon with Heterogeneous Hardware Multi-Pipelining. At *the 4th Annual Greater Chicago Area Systems Research Workshop*, April, 2015.
- GCASR 2015 **K.C. Hale** and P. Dinda. A Case for Transforming Parallel Runtimes into Operating System Kernels. At *the 4th Annual Greater Chicago Area Systems Research Workshop*, April, 2015.
- GCASR 2014 **K.C. Hale** and P. Dinda. Guarded Modules: Adaptively Extending the VMM's Privilege Into the Guest. At *the 3rd Annual Greater Chicago Area Systems Research Workshop*, May, 2014.
- February 2013 **K.C. Hale**. Dynamic Linking Considered Harmful. Talk given at the NU Computer Systems Reading Group, February, 2013.
- ICAC 2012 **K.C. Hale**, L. Xia, and P. Dinda. Shifting GEARS to Enable Guest-context Virtual Services. At *the 9th ACM International Conference on Autonomic Computing*, September, 2012.

SOFTWARE

Nautilus Aerokernel

Nautilus is an extremely lightweight kernel layer designed as a *privileged* library operating system (called an Aerokernel) that demonstrates the Hybrid Runtime model, wherein the combined parallel runtime and Aerokernel are transformed into a specialized OS kernel. It is a many-core capable OS layer that runs on commodity x64 hardware and the Intel Xeon Phi. I am the primary developer of Nautilus. Descriptions of related systems appear below.

<http://nautilus.halek.co>

Estimated Lines of Code: 35,000

- **Nemo**: Nemo is an event system in Nautilus for Hybrid Runtimes. Nemo accelerates asynchronous software event delivery by several orders of magnitude by leveraging hardware features typically only available to the OS (but not the runtime) in kernel-mode.

Estimated Lines of Code: 200

- **Multiverse Runtime:** Multiverse is a runtime system aimed at alleviating the effort required to build and port Hybrid Runtimes. It allows users to explore the benefits of HRTs by automatically running their legacy Linux programs in the Nautilus Aerokernel without any porting effort.

Estimated Lines of Code: 4,000

Philix

Philix is a tool that I designed for booting 3rd party OS kernels on the Intel Xeon Phi. It allows kernel developers to rapidly prototype new kernel mechanisms on the Phi without implementing Intel's SCIF protocol.

<http://philix.halek.co>

Estimated Lines of Code: 4,000

Palacios VMM

Palacios is an open-source, embeddable Virtual Machine Monitor actively developed by researchers at several institutions. I have built several systems within the context of Palacios and have made many contributions to the codebase, a subset of which are described below. The code for all of these systems can be found in the development branch of the Palacios repository at the website listed below.

<http://v3vee.org/palacios>

- **GEARS:** Guest Examination and Revision Services (GEARS) is a set of tools that allows developers to create *guests-context virtual services*, VMM-based services that extend *into* the guest.

Estimated Lines of Code: 2,500

- **Guarded Modules:** Guarded Modules extend the concept of guest-context virtual services by granting them *privileged* access to hardware and VMM state. Guarded Modules protect this privilege from the rest of the guest by maintaining a software border with compile-time and run-time techniques.

Estimated Lines of Code: 1,000

- **Virtualized DVFS:** This system allows fine-grained control of the Dynamic Voltage and Frequency Scaling (DVFS) hardware during VM exits, leveraging inferred information about guests to make informed power management decisions.

Estimated Lines of Code: 500

- **Virtual HPET:** This is a virtual implementation of the High-Precision Event Timer, a fine-grained platform timer present on most contemporary high-performance hardware. I added support for the HPET to allow us to run experimental systems like OSv on Palacios.

Estimated Lines of Code: 1,000

- **QEMU backend:** QEMU provides a rich diversity of virtual devices. This is one contributor to the simplicity of, e.g. the KVM codebase, as it leverages these device implementations. I wanted to similarly be able to leverage these devices for Palacios. This system implements that functionality with a software bridge between Palacios and QEMU.

Estimated Lines of Code: 2,000

- **VMM-emulated RTM:** This was the first VMM-emulated implementation of the Restricted Transactional Memory (RTM) component of the Intel Transactional Synchronization Extensions (TSX). Its performance is roughly 60x relative to Intel's emulator.

Estimated Lines of Code: 1,300

- **Palacios on the Cray XK6:** I ported the Palacios VMM to run on the Cray XK6 series of supercomputer nodes. This comprised several bug fixes and enhancements to the Palacios codebase.
- **Other contributions:** I have also participated in regular development and maintenance of the Palacios codebase. This includes bug fixes, enhancements to the extension architecture, guest configuration and loading, software interrupt and system call interception, and others.

Estimated Lines of Code: 12,000

SETI Lab

For our Introduction to Computer Systems Course (EECS 213) at Northwestern, we wanted a new lab to give students an earlier, practical introduction to parallel programming. To accomplish this, I designed and implemented SETI Lab, which is a lab that draws inspiration from SETI@Home. Students compete to parallelize signal analysis code and find alien signals in synthetic radio telescope data.

Estimated Lines of Code: 5,500

AWARDS AND HONORS

2015 · Best Short Presentation Award · *A Case for Transforming Parallel Runtimes into Operating System Kernels* · HPDC 2015

2005-2010 · Member of Turing Scholars Honors Computer Science Degree Program

2008-2010 · Member of Ronald E. McNair Post-Baccalaureate Achievement Program

2010-2011 · Murphy Graduate Fellowship Recipient

TEACHING

TA for Introduction to Databases (NU EECS 339)

TA for Introduction to Computer Systems (NU EECS 213), 2 quarters

Designed a new parallel computing lab called SETI Lab for the NU Introduction to Computer Systems (EECS 213) course. Students are tasked with parallelizing signal analysis in the search for synthetic “alien” signals.

Co-advised masters student Shiva Rao

Topic: Feasibility of Making DVFS Decisions in the VMM
Now Senior Software Engineer at Altera

Co-advised masters student Madhav Suresh

Topic: Parallel language synchronization techniques; Deterministic and stochastic barrier synchronization

Guided and assisted undergraduate students in independent study projects:

Conor Hetland & Jonathan Ford

Topic: Prototype port of the Nautilus AeroKernel to the Intel Xeon Phi

Akhil Guliani, Billy Gross, and Panitan Wongse-ammatt

Topic: Device file virtualization in the Palacios VMM

SERVICE AND ACTIVITIES

Leader and organizer of the Computer Systems Reading Group at Northwestern University

Member of Northwestern team for NCCDC security regional competition (placed 2nd)

Member of Northwestern team for ICTF worldwide security competition (placed 8th out of 72)

2013-Present · ACM Student Member (SIGOPS, SIGHPC)

ISCA 2009 · Student Volunteer

SC 2011 · Student Volunteer

DOE Exascale Research Conference 2012 · Student Scribe

SC 2012 · Student Volunteer

EXTERNAL REVIEWING

DATE 2012

ISPASS 2012, 2017

HPDC 2012, 2013, 2014, 2015

SC 2012, 2016

ICAC 2013

Parallel Computing (2014)

ICDCS 2015

OOPSLA 2016

TPDS (2016)

HPCA 2017

CGO 2017

TECHNICAL PROGRAM COMMITTEE MEMBERSHIPS

VHPC 2015, 2016, 2017

FiCloud 2016

CloudCom 2016, 2017

ICS 2017 (External Review Committee)

Local Chair, ICS 2017

OTHER INFORMATION

Languages

ENGLISH · Native

JAPANESE · Advanced (conversationally fluent, reading and writing)

SPANISH · Basic (simple words and phrases only)

REFERENCES

Available upon request.