Ellis H. Wilson III, Ph.D.

Contact Information

Information Sciences and Technology Building 351
Department of Computer Science and Engineering
The Pennsylvania State University
State College, PA, 16801 USA

Voice: (484) 947-7419
E-mail: ellis@cse.psu.edu
WWW: www.ellisv3.com

Research Interests

Distributed file systems and storage, non-volatile memories and their architectures, multilevel and distributed caching

Education

The Pennsylvania State University, State College, Pennsylvania USA

Ph.D., Computer Science and Engineering, August 2014

- Dissertation: A Protean Attack on the Compute-Storage Gap in High-Performance Computing
- Adviser: Mahmut T. Kandemir

La Salle University, Philadelphia, Pennsylvania USA

B.S., Computer Science, May 2009

Honors and Awards

Supercomputing: Best Paper Finalist & Best Student Paper Finalist, 2013

Supercomputing: 2nd Place, ACM Student Research Competition, 2012

The Pennsylvania State University: College of Engineering Fellowship, 2009

La Salle University: graduated Magna Cum Laude, 2009

NCAA Division I Cross Country: Atlantic 10 All-Conference, 2007
NCAA Division I Cross Country: Atlantic 10 Academic All-Conference, 2007

Boy Scouts: Eagle Scout, 2005

Academic Experience

The Pennsylvania State University, State College, Pennsylvania USA

Graduate Student **August 2009 - May 2014**
Ph.D. research and coursework.

Research Assistant **January 2010 - May 2014**
Collaborative research towards improving performance of traditional and non-volatile storage, particularly in distributed and high performance environments.

Teaching Assistant **August 2009 - December 2009**
Assisted in a computer science undergraduate course on spreadsheets and databases. Responsible for assisting students in and out of class on coursework, creating and grading homework assignments, and proctoring exams.

La Salle University, Philadelphia, Pennsylvania USA

Undergraduate Student **August 2005 - May 2009**
Undergraduate coursework.

Special Topics Research **January 2008 - May 2009**
Developed linux-based live operating system for automated creation of small- to medium-sized clusters from workstation computers. Distribution targetted at researchers without access to clusters but with access to computer labs, which are closed at night and can be repurposed for research.

Publications

Ellis Wilson, Myoungsoo Jung, Mahmut Kandemir. ZombieNAND: Resurrecting Dead NAND Flash for Improved SSD Longevity. IEEE International Symposium on Modeling Analysis and Simulation of Computer and Telecommunication Systems (MASCOTS '14).

Ellis Wilson, Garth Gibson, Mahmut Kandemir. Will They Blend?: Exploring Big Data Computation atop Traditional HPC NAS Storage. The International Conference on Distributed Computing

Systems (ICDCS '14).

Myoungsoo Jung, Ellis H. Wilson III, Wonil Choi, John Shalf, Hasan Metin Aktulga, Chao Yang, Erik Saule, Umit V. Catalyurek, Mahmut Kandemir. Exploring the Future of Out-Of-Core Computing with Compute-Local Non-Volatile Memory. Supercomputing 2013 (SC '13).

Ellis H. Wilson III. Performing Cloud Computation on a Parallel File System. ACM Student Research Competition at Supercomputing 2012 (ACM SRC '12).

Myoungsoo Jung, Ellis H. Wilson III, Mahmut Kandemir. Physically Addressed Queueing (PAQ): Improving Parallelism in Solid State Disks. The 39th International Symposium on Computer Architecture (ISCA '12).

Myoungsoo Jung, Ellis Wilson, David Donofrio, John Shalf, Mahmut Kandemir. NANDFlashSim: Intrinsic Latency Variation Aware NAND Flash Memory System Modeling and Simulation at Microarchitecture Level. The 28th IEEE Conference on Massive Data Storage (MSST '12).

Professional Experience

Panasas, Inc., Pittsburgh, Pennsylvania USA

Software Engineer **June, 2014 - Present**
Filling a full-time software engineering position at Panasas starting June 16th, 2014, assisting in largely C-based development of their high-performance distributed file storage appliance spinning disk and solid-state products.

Summer Researcher **May 2013 - August 2013**
Co-designed, developed, and debugged a file system tool written part in C and part in Perl, intended for analysis of the benefits of various deduplication and compression algorithms on any POSIX file system. Developed a novel deduplication algorithm. Continued collaboration with Panasas employees and Parallel Data Lab members at Carnegie Mellon University.

Summer Researcher **May 2012 - August 2012**
Continued exploration of Hadoop atop a network file system using the BigData benchmark TeraSort on a medium-size cluster. Developed new Hadoop file system, RAINFS, for improved performance in such environments. Continued collaboration with Panasas employees and Parallel Data Lab members at Carnegie Mellon University.

Summer Researcher **May 2011 - August 2011**
Was formally trained to administer Panasas ActiveScale storage and did background research on Hadoop MapReduce and Hadoop Distributed File System. Used such knowledge to explore usage of Hadoop atop a network file system as a supplement or to completely replace traditionally used local hard disk drives. Collaborated with both employees at Panasas and Parallel Data Lab members at Carnegie Mellon University.

Argonne National Laboratory, Chicago, Illinois USA

Summer Researcher **May 2010 - August 2010**
Worked on several research projects, including development of a cluster simulator for high fidelity replay of MPI traces through simulated PVFS client and server models. Also assisted in implementation of a specialized data structure for the next major release of the PVFS file system, which would improve performance of numerous, small metadata I/O requests to the storage subsystem.

Temple University, Philadelphia, Pennsylvania USA

Summer Research Assistant **May 2008 - August 2008 and May 2009 - August 2009**
Assisted faculty and graduate students at Temple's Center for Advanced Photonics Research in large data collection and processing. Also developed a specialized genetic algorithm and support vector machine for their closed-loop feedback ultrafast laser guidance systems. Managed a small linux cluster.

Computer Skills

- Languages: C, C++, Java, Perl, PHP, Octave (Matlab), MPI, several shell scripting languages, several web markup languages
- Applications: OmNet++, LaTeX, LaTeXBeamer, common office suite software, GIMP, Xen
- Operating Systems: Linux/GNU, occasional use of BSD and Windows.