Woong Shin, PhD

Research Staff
AI Analytics Methods at Scale (AIMS) Group, Advanced Technology Section (ATS),
National Center for Computational Sciences Division, Oak Ridge National Laboratory
E-Mail: shinw@ornl.gov

Overview

- **Current position:** (October 2020 ~ Current) Research Staff,
  AI Analytics Methods at Scale (AIMS) Group, Advanced Technology Section (ATS),
  National Center for Computational Sciences Division, Oak Ridge National Laboratory
- **Field of Interest**: data driven data center operations and control, large-scale data analytics system, systems for near-real time data analytics, AI and machine learning applications for intelligent facilities, NVRAM based storage system architecture, distributed systems, large-scale databases, operating systems, computer architecture, embedded systems, performance study on systems and architectures, system performance monitoring and tuning, software engineering
- **Experience:** 3 years as a research staff in the US Department of Energy's (DOE) National Laboratory System, working in the supercomputing program. 7 years of experience in the information technology (IT) industry of South Korea as a R&D research engineer and a software engineer working in enterprise information systems and networks

Educational Experience

- **PhD**, Electrical Engineering and Computer Science, Seoul National University, Republic of Korea
  Master's & Ph.D. Integrated Program (September 2010 ~ February 2017)
  Dissertation Title: *"OS I/O Stack Optimizations for Flash Solid-State Drives"*
- **BS c**, Computer Science, Korea University, Republic of Korea, 2003

Work Experience

- **Research Staff,** October 2020 ~ current
  AI Analytics Methods at Scale (AIMS) Group, National Center for Computational Sciences
  Oak Ridge National Laboratory, Oak Ridge, Tennessee, USA
- **Research Staff,** April 2020 ~ September 2020
  Technology Integration Group, National Center for Computational Sciences
  Oak Ridge National Laboratory, Oak Ridge, Tennessee, USA
- **Research Associate**, May 2017 ~ March 2020
  Technology Integration Group, National Center for Computational Sciences
  Oak Ridge National Laboratory, Oak Ridge, Tennessee, USA
- **Postdoctoral Research Associate**, February 2017 ~ April 2017
  Technology Integration Group, National Center for Computational Sciences
  Oak Ridge National Laboratory, Oak Ridge, Tennessee, USA
- **Research Intern**, April 2016 ~ September 2016
  Technology Integration Group, National Center for Computational Sciences
  Oak Ridge National Laboratory, Oak Ridge, Tennessee, USA
- **Research Intern**, September 2012 ~ February 2013
  IBM Austin Research Laboratory, Austin, Texas, USA
- **Research Assistant**, September 2010 ~ February 2017
  Distributed Computing Systems Laboratory, Dept. of Computer Science and Engineering,

Seoul National University, Republic of Korea

- **Research Engineer**, June 2006 ~ February 2010
  System Management Laboratory, Research & Development Center, TmaxSoft, Republic of Korea, (https://www.tmaxsoft.com)
- **Software Developer**, March 2003 ~ June 2006
  Infrastructure Technology Team, Samsung Networks, Republic of Korea
  (Currently merged into Samsung SDS - https://www.samsungsds.com)

Selected Projects

- **Cooling Intelligence for Summit**
  Oak Ridge National Laboratory Leadership Computing Facility (OLCF)
  Technology Integration Group, National Center for Computational Sciences
  Oak Ridge National Laboratory
  Research Associate, May 2018 ~ Current
  - Delivered a near real-time monitoring and data analytics system that aggregates data from the facility and Summit which was required to push maximum energy efficiency. The system provides near real-time visualization of the current state of the Summit system and the cooling plant, adding confidence in adjusting high-impact cooling parameters. OLCF was able to address over-cooling issues imposed by conservative vendor cooling specifications, introducing significant cooling energy savings.
- **Multi-tiered Storage Systems**
  Technology Integration Group, National Center for Computational Sciences
  Oak Ridge National Laboratory
  Research Associate, April 2016 ~ Current
  - Evaluating NVM-e SSDs from multiple vendors (Intel, Samsung & HGST) to understand the impact on our systems when employed as a component (i.e., burst buffer) within an HPC multi-tiered storage environment. Research in designing system software that aims to solve the difficulty of manually managing the namespace, placement and movement of various data across multiple storage tiers such as burst buffers, parallel file systems, object stores and long-term data backup systems.
- **Enhanced OS & SSD Cross Layer I/O Optimization**
  Distributed Computing Laboratory, Seoul National University
  Research Assistant, January 2015 ~ February 2017
  - Research on advanced cross-layer optimization techniques by introducing new SSD functionalities and extending the interface between the OS and the SSD. Employed multiple SSD evaluation platforms for research. Developed SSD FTL code enhancement as well as corresponding block device driver codes for the OpenSSD Cosmos SSD evaluation platform. Also built a DRAM backed FPGA based SSD latency emulation platform with the Xilinx VC709 evaluation board. Developed Linux block device driver code (v3.5.0), FTL enhancements running on SSD microprocessors (ARM Cortex A9, Xilinx Microblaze) as well as the FPGA logic (Xilinx Zynq-7000 & Xilinx Virtex7) for latency emulation, interrupt control and PCI-e communication. OpenSSD platform results presented in USENIX HotStorage'16 [8].
- **Hardware & Software co-optimization**
  for High Performance / High Functionality SSDs in data centers

Research subcontracted to Samsung Electronics, Seoul National University
Research Assistant, January 2014 ~ December 2014

- Research on cross-layer optimizations for key value storages targeting to reduce latency variations caused by mixed read and write operations on a flash SSD. Developed a multi-SSD storage engine which employs the capability of performing read write separation as well as exploiting garbage collection API calls provided by the SSDs (Customized Samsung SSDs). Developed a custom storage engine integrated with Memcached 1.4.17 to demonstrate the effectiveness of tight cross-layer optimizations. Was able to reduce the 99.9999 percentile tail latency of foreground read operations to be under 1ms. Reported the results in ICPADS'14 [6] & BigComp'15 [7].

- **OS Optimizations for Multiple Flash SSDs**
  Distributed Computing Laboratory, Seoul National University
  Research Assistant, June 2013 ~ June 2015
  - Research on optimizing the Linux 3.2.40 file I/O path to enhance the IOPS scalability when using multiple SSDs. Extended the OS block storage interface, enhanced the interactions between the OS and the device. Developed a hardware abstraction layer that provides a generic interface to multiple types of HBA controllers (AHCI, LSI Megaraid, NVM-e, DRAM SSD) in order to demonstrate the capability of interface extensions involving exposing H/W abstractions such as tags, queues & interrupts. Total 10k LoC Linux Kernel code developed to demonstrate the idea. Achieved 100% IOPS gain (from 300k IOPS to 600k IOPS) with the code. Reported the result in USENIX ATC'14 [5].
- **Optimized Systems for Big Data Analytics: Workload analysis for applying NVRAM components**
  Internship, IBM Austin Research Laboratory, USA
  Research Intern, September 2012 ~ June 2013
  - Research on evaluating the impact of using NVRAM components on systems for big data analytics. Built a research prototype workload characterization module which extracts condensed application I/O models on the fly to solve the difficulties of extracting large trace data from distributed data intensive applications such as Hadoop. Hooked system calls such as open / close, read/write and lseek system calls per I/O thread to extract multiple Markov models from runtime traces with less than 5% performance degradation of the host program. The condensed Markov models were used to reproduce the application behavior. Total 10k LoC C code and 2.1k LoC Python code developed to demonstrate the idea.
- **Development of a File System based on Phase Change Memory to Support more than 200,000 IOPS for Peta-scale Computing**
  Research Grant, Ministry of Knowledge Economy (MKE) of Republic of Korea, Seoul National University
  Research Assistant, May 2011 ~ August 2012
  - Research on optimizing the Linux 2.6.32 file I/O path to better support high performance storage devices, projected to be using new memory technologies. Performed profiling and optimization activities with the Linux 2.6.32 page cache & block layer code using a high-performance PCI-e 2.0 SSD backed with a large DRAM array. Fixed performance degradation when read-ahead was enabled (20%~30%) by patching the read-ahead code in the page cache which deals with multiple sequential streams. Results reported in HotStorage'12 [2] and TOCS [4]. Found unnecessary memory copies with DMA pages residing above the 32bit limit (non DMA32 pages). In response, developed a new type of page cache code which places the page cache under the 32bit limit to eliminate the additional memory copies. Total 3.6k LoC patched for the additional functionality.
- **A Research on a Hierarchical Storage Architecture for Cloud Computing**

Research Grant, National Research Foundation of Korea (NRF), Seoul National University
Research Assistant, August 2010 ~ August 2012

- Research on an RDMA based remote memory disk cache for virtual machines. Developed a Linux kernel (v2.6.32) block device driver (1.4k LoC) which employs kernel level RDMA transactions to exploit remote DRAM exposed by remote memory donors, with the capability of direct DMA memory mapping support for virtualized environments. Employed and enhanced the dm-cache (Linux) code to better use our RDMA block device. Results reported in [1].

Community Services

- **2020 Energy Efficient High-Performance Computing Working Group Workshop**
  Program Committee / Session Lead & Moderator - ARM Operational Experiences
- **2020 Energy Efficient High-Performance Computing Working Group - State of Practice Workshop**
  Program Committee
- **2020 Cluster Computing**
  External reviewer
- **Energy Efficient High-Performance Computing Working Group**
  Operational data analytics team (https://eehpcwg.llnl.gov)
  November 2019 ~ Current
- **2019 International Conference for High Performance Computing, Networking, Storage and Analysis (SC19)**
  External reviewer in behalf of Dr. Sudharshan Vazhkudai
-

Publications

- **[12]** Justin Thaler, ***Woong Shin***, Steven Roberts, Jim Rogers, Todd Rosedahl, "Hybrid Approach to HPC Cluster Telemetry and Hardware Log Analytics", in The 2020 IEEE High Performance Extreme Computing Virtual Conference, September 2020
- **[11]** Michael Ott, ***Woong Shin***, Norman Bourassa, Torsten Wilde, Stefan Ceballos, Melissa Romanus, Natalie Bates, "Global Experiences with HPC Operational Data Measurement, Collection and Analysis", in The 2020 IEEE International Conference on Cluster Computing (CLUSTER) workshops, September 2020
- **[10]** ***Woong Shin***, Christopher D Brumgard, Bing Xie, Sudharshan S Vazhkudai, Devarshi Ghoshal, Sarp Oral, Lavanya Ramakrishnan, "Data Jockey: Automatic data management for HPC multi-tiered storage systems", in The 33rd IEEE International Parallel and Distributed Processing Symposium (IPDPS'19), Rio de Janeiro, Brazil, May 2019
- **[9]** ***Woong Shin***, "OS I/O Path Optimizations for Flash Solid-State Drives", Ph.D. Thesis, Seoul National University, Feb 2017
- **[8]** ***Woong Shin***, Jaehyun Park, Heon, Y. Yeom, "Unblinding the OS to Optimize User-Perceived Flash SSD latency", in The 8th USENIX Workshop on Hot Topics in Storage and File Systems (USENIX HotStorage'16), Denver, CO, Jun 2016
- **[7]** ***Woong Shin***, Myeongcheol Kim, Kyudong. Kim, and Heon. Y. Yeom, "Providing QoS through Host Controlled Flash SSD Garbage Collection and Multiple SSDs,", in The Second International Conference on Big Data and Smart Computing (BigComp'15), Jeju Island, South Korea, 9-12 Feb, 2015,
  (Best paper award)

- **[6]** ***Woong Shin***, Myeongcheol Kim, Jinyoung. Choi, Hyeonsang. Eom, and Heon. Y. Yeom, “HIOPS-KV : Exploiting Multiple Flash Solid-state Drives for Key Value Stores,” in The fourth International Workshop on Extreme Scale Computing Application Enablement - Modeling and Tools (ESCAPE’14) in conjunction with ICPADS'14, Hsinchu, Taiwan, 16-19 Dec, 2014.
- **[5]** ***Woong Shin***, Quichen Chen, Myoung Won Oh, Hyeonsang Eom and Heon Y. Yeom, “OS I/O Path Optimizations for Flash Solid-state Drives”, USENIX ATC’14, Jun 2014
- **[4]** Young Jin Yu, Dong In Shin, ***Woong Shin***, Nae Young Song, Jae Woo Choi, Hyeong Seog Kim, Hyeonsang Eom and Heon Y. Yeom, “Optimizing Block I/O Subsystem for Fast Storage Devices”, ACM Transactions on Computer Systems (TOCS), Jun 2014. (SCI)
- **[3]** NaeYoung Song, Young Jin Yu, ***Woong Shin***, Hyeonsang Eom, Heon Y. Yeom, “Low-latency Memory-Mapped I/O for Data-Intensive Applications on Fast Storage Devices”, High Performance Computing, Networking, Storage and Analysis (SCC), 2012
- **[2]** Young Jin Yu, Dong In Shin, ***Woong Shin***, Nae Young Song, Hyeonsang Eom, and Heon Y. Yeom, "Exploiting Peak Device Throughput from Random Access Workload", USENIX HotStorage'12, Jun 2012
- **[1]** Hyuck Han, Young Choon Lee, ***Woong Shin***, Hyungsoo Jung, Heon Y. Yeom, Albert Y. Zomaya, "Cashing in on the Cache in the Cloud", IEEE Transactions on Parallel and Distributed Systems (TPDS), Aug 2012. (SCI)