Sang-Woo Jun

32-G836, Stata Center, MIT
Cambridge, MA 02139
wjun@csail.mit.edu
http://people.csail.mit.edu/wjun

EDUCATION

Ph.D., Electrical Engineering & Computer Science (2013–Present)
Massachusetts Institute of Technology, Cambridge, MA, USA
Computation Structures Group,
Computer Science and Artificial Intelligence Laboratory *(Advisor: Arvind)*
Total GPA: 4.6/5 (108 Units)

Master of Science, Electrical Engineering & Computer Science (2011–2013)
Massachusetts Institute of Technology, Cambridge, MA, USA

Bachelor of Science, Computer Science and Engineering (2004–2009)
Seoul National University, Seoul, Korea
Total GPA: 3.53/4.3 (154 Units)
Major GPA: 3.67/4.3 (103 Units)
Major GPA (Upper division): 3.84/4.3 (53 Units)

SKILLSET

Technical Expertise:

- Low-level hardware design on FPGAs
- System software development
- Large scale systems including Hadoop, Accumulo and SQL
- Various languages including C, C++, Bluespec, Verilog, Java, PHP, Python, Javascript, x86 assembly, and OCaml.

Personal Projects:

- Operating System development for x86 (bootloader, kernel, VM, FS)
- Linux device driver development (PCIe/Network/Storage)
- distributed file sharing tools
- reverse-engineering software binaries
- Mobile games

Linguistic Abilities: Fluent in English and Korean.
G-ETAT Korean-English Translator's License

Translation of Professional Books:

- Greg Hoglund and Gary McGraw, Exploiting Online Games, Addison-Wesley (2007), 2008,
- Steve Suehring, JavaScript Step-byStep, O'Reilly (2008), 2008
- Gregor Hohpe and Bobby Woolf, Enterprise Integration Patterns, Addison-Wesley (2003)

RESEARCH PROJECTS

High-Performance Analytics using Accumulo with Smart Flash
2015–Present
We aim to accelerate large scale data analytics using the Accumulo key-value store, by using fast distributed flash storage augmented with FPGA-based in-store processors.
Jointly funded by Samsung and Lincoln Laboratory

BlueDBM - Distributed Flash Store for Big Data Analytics
2011–Present
BlueDBM is a complete HW/SW stack of a distributed flash storage system with FPGA-based reonfigurable controllers and a hardware-accelerated storage network, aiming to aid in high-throughput analytics and rapid development of hardware accelerators.
BlueDBM has been used to accelerate applications such as text search, graph analytics, image search and distributed web cache.
Jointly funded by Quanta, Samsung, Intel and Lincoln Laboratory.
(`http://people.csail.mit.edu/wjun/ssd.htm`)

Personalized Power/Performance Optimization for Android Smartphones
2010–2011
National Research Laboratory project funded by the Ministry of Education, Science and Technology, Korean Government.

Visualizing and Detecting Concurrency Bugs in Multicore Software
2009–2009
Part of the National Information Technology Research Center project (on Development of Embedded Software for Multicore-based Avionics), funded by the Ministry of Knowledge and Economy, Korean Government.

Foretelling Possibilities of Heart Diseases by Data Mining Electro Cardiograms,
2009 Capstone design class, Leader of Student Team, Joint Project with Samsung Advanced Institute of Technology (SAIT), The collaborated research staffs of SAIT were impressed with my teams results and offered internship positions of my whole team to continue this research.

WORK EXPERIENCE

Software Engineer, Security Engineer 2005–2007
Nexon Inc, Seoul, Korea

- Client and server developer for MapleStory, Nexon's biggest MMORPG. Experience with high load client-server architectures.
- Client security engineer for MaplerStory, preventing and detecting reverse engineering of client software.

TEACHING EXPERIENCE

Teaching Assistant Fall 2012
6.s195 Computer Architecture Laboratory
Department of Electrical Engineering and Computer Science, MIT, Cambridge, USA

- Worked with another TA to develop in-class activities and programming lab assignments.
- Gave TA lectures and mentored students for their final projects.
- Graded student programming assignments.

SELECTED PUBLICATIONS

Notable publications emphasized red

Sang-Woo Jun, Ming Liu, Sungjin Lee, Jamey Hicks, John Ankcorn, Myron King, Shuotao Xu, Arvind. 2016
"BlueDBM: Distributed Flash Storage for Big Data Analytics," to appear in ACM Transactions on Computer Systems (TOCS)

Sungjin Lee, Ming Liu, **Sangwoo Jun**, Shuotao Xu, Jihong Kim, 2016
"Application-Managed Flash" File and Storage Technologies (FAST)

Ming Liu, **Sang-Woo Jun**, Sungjin Lee, Jamey Hicks, Arvind, 2016
"minFlash: A Minimalistic Clustered Flash Array," Design, Automation and Test in Europe (DATE)

Sang-Woo Jun, Chanwoo Chung, Arvind, 2015
"Large-scale high-dimensional nearest neighbor search using Flash memory with in-store processing," International Conference on Reconfigurable Computing and FPGAs (ReConFig)

Sang-Woo Jun, Ming Liu, Shuotao Xu, Arvind. 2015
"A Transport-Layer Network for Distributed FPGA Platforms," International Conference on Field Programmable Logic and Applications (FPL)

Sang-Woo Jun, Ming Liu, Sungjin Lee, Jamey Hicks, John Ankcorn, Myron King, Shuotao Xu, Arvind. 2015
"BlueDBM: An Appliance for Big Data Analytics," International Symposium on Computer Architecture (ISCA)

Sang-Woo Jun, Ming Liu, Kermin Fleming, Arvind. 2014
"Scalable Multi-Access Flash Store for Big Data Analytics," International Symposium on Field-Programmable Gate Arrays (FPGA)

Sang-Woo Jun, Kermin Fleming, Michael Adler, Joel Emer. 2012
"ZIP-IO: Architecture for Application-Specific Compression of Big Data," International Conference on Field Programmable Technology (FPT)

Mun-Hye Jang, Ok-Kyoon Ha, **Sang-Woo Jun**, Yong-Kee Jun. 2009
"A Tool for Detecting First Races in OpenMP Programs," International Conference on Parallel Computing Technologies

SELECTED MEDIA COVERAGE

Ditching RAM may lead to low-cost supercomputers Jul. 2015
URL: `http://www.engadget.com/2015/07/12/mit-flash-only-supercomputers/`
Jon Fingas, Engadget

Cutting cost and power consumption for big data Jul. 2015
URL: `http://news.mit.edu/2015/cutting-cost-power-big-data-0710`
Larry Hardesty, MIT News

MIT Prototypes Zippy Flash-FPGA Hybrid Storage Array Feb. 2014
URL: http://www.enterprisetech.com/2014/02/03
/mit-prototypes-zippy-flash-fpga-hybrid-storage-array/
Timothy Prickett Morgan, Enterprise Tech

Storage system for 'big data' dramatically speeds access to information Jan. 2014
URL: http://news.mit.edu/2014/
storage-system-for-big-data-dramatically-speeds-access-to-information-0131
Helen Knight, MIT News

SCHOLARSHIPS

- **National Scholarship for Science and Engineering** (Full tuition, 2004-2010), Korean Student Aid Foundation, Korean Government.
- **TOEFL Scholarship** (Full tuition, 2010), ETS (Educational Testing Service), USA.
- **Kwanjeong Educational Foundation Scholarship** (Tuition, 2011-Present), Kwanjeong Educational Foundation, Korea