

# Sang-Woo Jun

---

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

<b>EDUCATION</b>	<b>Massachusetts Institute of Technology</b> Ph.D. in Electrical Engineering & Computer Science <ul style="list-style-type: none"><li>• Expected to finish Summer, 2018</li><li>• Advisor: <b>Arvind</b></li></ul>	2013–Present
	<b>Massachusetts Institute of Technology</b> M.S. in Electrical Engineering & Computer Science <ul style="list-style-type: none"><li>• Thesis: Scalable Multi-Access Flash Store for Big Data Analytics</li><li>• Advisor: <b>Arvind</b></li></ul>	2011–2013
	<b>Seoul National University</b> B.S. in Computer Science and Engineering	2004–2011
<b>RESEARCH INTERESTS</b>	<b>Hardware Acceleration:</b> <ul style="list-style-type: none"><li>• Application-specific hardware acceleration using FPGA</li><li>• Programming system, debugging, and libraries</li><li>• Hardware acceleration in the cloud</li></ul> <b>Flash Storage:</b> <ul style="list-style-type: none"><li>• Out-of-core analytic systems to reduce cost while maintaining performance</li><li>• Distributed flash storage with in-storage accelerators</li><li>• Operating system support for fast storage</li></ul>	
<b>PROFESSIONSL EXPERIENCE</b>	<b>Oracle Inc</b> , Big Data Discovery, Cambridge, MA Summer Intern	2016
	<b>Nexon Inc</b> , Maple Story development team, Seoul, Korea Game Server/Client Software Engineer, Client Security Engineer	2005–2007
<b>PUBLICATIONS</b>	“Using accelerated flash storage for external graph analytics,” <b>Sang-Woo Jun</b> , Andy Wright, Sizhuo Zhang, Shuotao Xu, Arvind (Under Review) International Symposium on Computer Architecture ( <b>ISCA</b> )	2018
	“BlueCache: A Scalable Distributed Flash-based Key-value Store,” <sup>1</sup> Shuotao Xu, Sungjin Lee, <b>Sang-Woo Jun</b> , Ming Liu, Jamey Hicks, Arvind, International Conference on Very Large Data Bases ( <b>VLDB</b> )	2017
	“Terabyte Sort on FPGA-Accelerated Flash Storage,” <b>Sang-Woo Jun</b> , Shuotao Xu, Arvind, Field-Programmable Custom Computing Machines ( <b>FCCM</b> )	2017

---

<sup>1</sup>**Media coverage:** MIT News, Inside Science, etc.

- “BlueDBM: Distributed Flash Storage for Big Data Analytics,” 2016  
**Sang-Woo Jun**, Ming Liu, Sungjin Lee, Jamey Hicks, John Ankcorn, Myron King, Shuotao Xu, Arvind,  
ACM Transactions on Computer Systems (**TOCS**)
- “Application-Managed Flash” 2016  
Sungjin Lee, Ming Liu, **Sangwoo Jun**, Shuotao Xu, Jihong Kim,  
File and Storage Technologies (**FAST**)
- “minFlash: A Minimalistic Clustered Flash Array,” 2016  
Ming Liu, **Sang-Woo Jun**, Sungjin Lee, Jamey Hicks, Arvind,  
Design, Automation and Test in Europe (**DATE**)
- “Large-scale high-dimensional nearest neighbor search using Flash memory with in-store processing,” 2015  
**Sang-Woo Jun**, Chanwoo Chung, Arvind,  
International Conference on Reconfigurable Computing and FPGAs (**ReConFig**)
- “A Transport-Layer Network for Distributed FPGA Platforms,” 2015  
**Sang-Woo Jun**, Ming Liu, Shuotao Xu, Arvind,  
International Conference on Field Programmable Logic and Applications (**FPL**)
- “BlueDBM: An Appliance for Big Data Analytics,”<sup>2</sup> 2015  
**Sang-Woo Jun**, Ming Liu, Sungjin Lee, Jamey Hicks, John Ankcorn, Myron King, Shuotao Xu, Arvind,  
International Symposium on Computer Architecture (**ISCA**)
- “Scalable Multi-Access Flash Store for Big Data Analytics,”, 2014  
**Sang-Woo Jun**, Ming Liu, Kermin Fleming, Arvind,  
International Symposium on Field-Programmable Gate Arrays (**FPGA**)
- “ZIP-IO: Architecture for Application-Specific Compression of Big Data,” 2012  
**Sang-Woo Jun**, Kermin Fleming, Michael Adler, Joel Emer,  
International Conference on Field Programmable Technology (**FPT**)

**TEACHING  
EXPERIENCE**

- Teaching Assistant** Fall 2012  
6.s195 Computer Architecture Laboratory  
Department of Electrical Engineering and Computer Science, MIT, Cambridge, USA
- One of two TAs developing in-class activities and programming lab assignments.
  - Weekly recitations, special lectures and mentored 18 students for final projects.
  - Graded student programming assignments.

**BOOK  
CHAPTER**

- Affordable mongo graph analytics using hardware accelerated flash storage,** 2018  
in *Managing Heterogeneous Data and Polystore Databases*,  
edited by Vijay Gadepally, Sam Madden and and Michael Stonebraker  
(In progress)

---

<sup>2</sup>**Media coverage:** MIT News, Engadget, Enterprise Tech, The Next Platform, etc.

**SCHOLARSHIPS AND GRANTS** MIT-Spain La Caxia Foundation 2017  
Seed fund for collaboration with Barcelona Supercomputing Center

**Kwanjeong Educational Foundation Scholarship** 2011–2016  
Tuition, Kwanjeong Educational Foundation, Korea

**TOEFL Scholarship** 2010  
ETS (Educational Testing Service), USA

**National Scholarship for Science and Engineering** 2004–2010  
Full Tuition, Aid Foundation, Korean Government

**SKILLSET**

**Programming Skills**

- FPGA accelerator development: Bluespec, Verilog
- System software development: Device driver, File system
- Scalable analytics: Hadoop, Spark, GraphLab, Accumulo

**Languages:** Fluent in Korean and English