

Sang-Woo Jun

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

EDUCATION	Massachusetts Institute of Technology Ph.D. in Electrical Engineering & Computer Science <ul style="list-style-type: none">• Expected to finish Summer, 2018• Advisor: Arvind	2013–Present
	Massachusetts Institute of Technology M.S. in Electrical Engineering & Computer Science <ul style="list-style-type: none">• Thesis: Scalable Multi-Access Flash Store for Big Data Analytics• Advisor: Arvind	2011–2013
	Seoul National University B.S. in Computer Science and Engineering	2004–2011
RESEARCH INTERESTS	Hardware Acceleration: <ul style="list-style-type: none">• Application-specific hardware acceleration using FPGA• Programming system, debugging, and libraries• Hardware acceleration in the cloud Flash Storage: <ul style="list-style-type: none">• Out-of-core analytic systems to reduce cost while maintaining performance• Distributed flash storage with in-storage accelerators• Operating system support for fast storage	
PROFESSIONSL EXPERIENCE	Oracle Inc , Big Data Discovery, Cambridge, MA Summer Intern	2016
	Nexon Inc , Maple Story development team, Seoul, Korea Game Server/Client Software Engineer, Client Security Engineer	2005–2007
PUBLICATIONS	[JuWr18] “Using accelerated flash storage for external graph analytics,” Sang-Woo Jun , Andy Wright, Sizhuo Zhang, Shuotao Xu, Arvind (Under Review) International Symposium on Computer Architecture (ISCA)	2018
	[XuLe17] “BlueCache: A Scalable Distributed Flash-based Key-value Store,” ¹ Shuotao Xu, Sungjin Lee, Sang-Woo Jun , Ming Liu, Jamey Hicks, Arvind, International Conference on Very Large Data Bases (VLDB)	2017
	[JuXu17] “Terabyte Sort on FPGA-Accelerated Flash Storage,” Sang-Woo Jun , Shuotao Xu, Arvind, Field-Programmable Custom Computing Machines (FCCM)	2017

¹**Media coverage:** MIT News, Inside Science, etc.

- [JuLi16] “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**)
- [LeLi16] “Application-Managed Flash,” 2016
 Sungjin Lee, Ming Liu, **Sangwoo Jun**, Shuotao Xu, Jihong Kim,
 File and Storage Technologies (**FAST**)
- [LiJu16] “minFlash: A Minimalistic Clustered Flash Array,” 2016
 Ming Liu, **Sang-Woo Jun**, Sungjin Lee, Jamey Hicks, Arvind,
 Design, Automation and Test in Europe (**DATE**)
- [JuCh15] “Large-scale High-Dimensional Nearest Neighbor Search using 2015
 Flash Memory with In-Store Processing,”
Sang-Woo Jun, Chanwoo Chung, Arvind,
 International Conference on Reconfigurable Computing and FPGAs (**ReConFig**)
- [JuLX15] “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**)
- [JuLi15] “BlueDBM: An Appliance for Big Data Analytics,” ² 2015
Sang-Woo Jun, Ming Liu, Sungjin Lee, Jamey Hicks, John Ankcorn,
 Myron King, Shuotao Xu, Arvind,
 International Symposium on Computer Architecture (**ISCA**)
- [JuLi14] “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**)
- [JuFl12] “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**)

TALKS

- “Terabyte Sort on FPGA-Accelerated Flash Storage” 2017
 at Field-Programmable Custom Computing Machines (**FCCM**)
- “Large-scale High-Dimensional Nearest Neighbor Search using Flash Memory 2015
 with In-Store Processing,”
 at International Conference on Reconfigurable Computing and FPGAs (**ReConFig**)
- “BlueDBM: A Multi-access, Distributed Flash Store for Big Data Analytics” 2015
 at CSAIL Affiliates Program, MIT
- “BlueDBM Hardware Details” 2015
 at Samsung Inc., Seoul, Korea
- “BlueDBM: An Appliance for Big Data Analytics” 2015
 at International Symposium on Computer Architecture (**ISCA**)

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

	“Triple Store on BlueDBM” for Lincoln Laboratory at MIT	2014
	“BlueDBM: A Multi-Access, Distributed Flash Store for Big Data Analytics” at EMC Flash Forum, Hopkinton, MA	2014
	“Scalable Multi-Access Flash Store for Big Data Analytics” at International Symposium on Field-Programmable Gate Arrays (FPGA)	2014
	“BlueDBM: Distributed Flash Store for Big Data Analytics” at Xilinx Inc., San Jose, CA	2013
	“ZIP-IO: Architecture for Application-Specific Compression of Big Data” at International Conference on Field Programmable Technology (FPT)	2012
TEACHING EXPERIENCE	Teaching Assistant 6.s195 Computer Architecture Laboratory Department of Electrical Engineering and Computer Science, MIT, Cambridge, USA <ul style="list-style-type: none"> • 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. 	Fall 2012
BOOK CHAPTER	Affordable mongo graph analytics using hardware accelerated flash storage, in <i>Managing Heterogeneous Data and Polystore Databases</i> , edited by Vijay Gadepally, Sam Madden and Michael Stonebraker (In progress)	2018
SCHOLARSHIPS AND GRANTS	MIT-Spain La Caixa Foundation Seed fund for collaboration with Barcelona Supercomputing Center	2017
	Kwanjeong Educational Foundation Scholarship Tuition, Kwanjeong Educational Foundation, Korea	2011–2016
	TOEFL Scholarship ETS (Educational Testing Service), USA	2010
	National Scholarship for Science and Engineering Full Tuition, Aid Foundation, Korean Government	2004–2010
SKILLSET	Programming Skills <ul style="list-style-type: none"> • FPGA accelerator development: Bluespec, Verilog • System software development: Device driver, File system • Scalable analytics: Hadoop, Spark, GraphLab, Accumulo Languages: Fluent in Korean and English	