Shuotao Xu

32 Vassar Street 32-G836, Cambridge, MA 02139, USA
shuotao@csail.mit.edu
https://people.csail.mit.edu/shuotao

Research Interests

Hardware Accelerators: Application-specific accelerators using FPGAs
Flash Storage: Novel in-storage system architectures for big-data applications

Education

Massachusetts Institute of Technology Cambridge, MA
Ph.D., Electrical Engineering and Computer Science (GPA: 4.9/5.0) Jun. 2021
- Thesis: *Computing Big-Data Applications Near Flash*
- Advisor: *Professor Arvind*

S.M., Electrical Engineering and Computer Science (GPA: 5.0/5.0) May 2016
- Thesis: *BlueCache: a Scalable Distributed Flash-based Key Value Store*
- Advisor: *Professor Arvind*

University of Illinois at Urbana-Champaign Urbana, IL
B.S., Electrical and Computer Engineering (GPA: 3.94/4.0) May 2012

Research Projects

AQUOMAN, MIT CSAIL 2017-Present
AQUOMAN is an in-storage Analytics QUery Offloading MAchiNe, which provides a more cost-effective high-performance solution to multi-terabyte SQL analytic workloads than an x86-based solution. AQUOMAN "offload"s most SQL operators to SSDs, including multi-way joins. AQUOMAN uses a streaming computation model, which allows AQUOMAN to process queries with a reasonable amount of DRAM for intermediate results. AQUOMAN is a general analytic query processor, which can be integrated in the database management system (DBMS) transparently. A prototype of AQUOMAN using FPGAs has shown that, using TPC-H benchmarks on 1TB data sets, a single instance of 1TB AQUOMAN disk, on average, can free up x86-host's CPU cycles by 70% and its DRAM usage by 60%.

BlueCache, MIT CSAIL 2014-2016
BlueCache offers a 10X to 100X cheaper data-center-scale key-value caching solution based on flash storage and hardware accelerators than a DRAM-based x86 solution. In BlueCache key-value pairs are stored in flash storage and all key-value operations, including the flash controller are directly implemented in hardware. Furthermore, BlueCache includes a fast interconnect between flash controllers to provide a scalable solution. We have built an FPGA prototype and shown that BlueCache can outperform a DRAM-based key-value cache when the latter has more than 7.4% misses.

BlueDBM, MIT CSAIL 2013-2014
BlueDBM is a specialized appliance for big-data analytics, which uses flash for fast and cheap storage, a low-latency high-bandwidth inter-controller network for distributed data access and reconfigurable hardware accelerators to process data within the storage device. We have built a 20-node BlueDBM system with a custom flash board to demonstrate significant speedups in search, graph algorithms and databases.

Journal Publications

- Sang-Woo Jun, Ming Liu, Sungjin Lee, Jamey Hicks, John Ankcorn, Myron King, **Shuotao Xu** and Arvind. "BlueDBM: Distributed Flash Storage for Big Data Analytics", *ACM Transactions on Computer Systems (TOCS'16)*, 2016

Conference Publications

- Xuhao Chen, Tianhao Huang, **Shuotao Xu**, Thomas Bourgeat, Chanwoo Chung and Arvind, "FlexMiner: A Pattern-Aware Accelerator for Graph Pattern Mining", *International Symposium on Computer Architecture (ISCA'21)*, 2021
- **Shuotao Xu**, Thomas Bourgeat, Tianhao Huang, Hojun Kim, Sungjin Lee and Arvind, "AQUOMAN: an Analytic-Query Offloading Machine", *International Symposium on Microarchitecture (MICRO'20)*, 2020
- Sang-Woo Jun, Andy Wright, Sizhuo Zhang, **Shuotao Xu** and Arvind, "GraFBoost: Accelerated Flash Storage for External Graph Analytics", *International Symposium on Computer Architecture (ISCA'18)*, 2018
- **Shuotao Xu**, Sungjin Lee, Sang-Woo Jun, Ming Liu, Jamey Hicks and Arvind, "BlueCache: A Scalable Distributed Flash-based Key-value Store", *International Conference on Very Large Data Bases (VLDB'17)*, 2017
- Sang-Woo Jun, **Shuotao Xu** and Arvind, "Terabyte Sort on FPGA-Accelerated Flash Storage", *International Symposium on Field-Programmable Custom Computing Machines (FCCM'17)*, 2017
- Sungjin Lee, Ming Liu, Sang-Woo Jun, **Shuotao Xu**, Jihong Kim, and Arvind, "Application-Managed Flash", *USENIX Conference on File and Storage Technologies (FAST'16)*, 2016
- Sang-Woo Jun, Ming Liu, **Shuotao Xu** and Arvind. "A Transport-Layer Network for Distributed FPGA Platforms", *International Conference on Field Programmable Logic and Applications (FPL '15)*, 2015
- Sang-Woo Jun, Ming Liu, Sungjin Lee, Jamey Hicks, John Ankcorn, Myron King, **Shuotao Xu** and Arvind. "BlueDBM: an appliance for big data analytics", *International Symposium on Computer Architecture (ISCA'15)*, 2015

Teaching Experience

Teaching Assistant, 6.375 Complex Digital Systems — MIT Spring 2019
Advanced undergraduate class for digital circuit designs — *Class size: 20*
Redesigned the processor lab in Bluespec for a 5-stage pipelined in-order R32I RISC-V Processor. Mentored a final project on a sparse-matrix multiplication accelerator.

Teaching Assistant, 6.004 Computer Structures — MIT Fall 2018
Freshman class for digit designs and computer architecture — *Class size: 200-300*
The leading TA who assisted in re-developing class materials using a new generation of hardware design tools. Redesigned the labs and the final project in Bluespec System Verilog, which leads to a 5-stage pipelined in-order R32I RISC-V processor.

Work Experience

Research Intern, Samsung Semiconductor Inc., San Jose, CA — Summer 2018
Project: Near-data SQL processing — *Mentor: Oscar Pinto*
Explored database acceleration with next-generation in-storage computing devices using FPGAs. Designed and implemented in-storage database accelerators, and integrated them with the MonetDB database software.

Research Intern, Microsoft Research, Redmond, WA — Summer 2017
Project: Distributed Key-value Device — *Mentor: Jae Young Do*
Explored an ARM-based architecture of distributed key-value store. Offloaded compression and decompression to hardware IPs to accelerate key-value operations.

Skill Set

Hardware Development: Verilog, System Verilog, Bluespec System Verilog, VHDL
Software Development: C, C++, Python, MATLAB, Bash, Tcl