

HAOYU ZHANG

Princeton University
Department of Computer Science
35 Olden St, Princeton, NJ 08540

haoyuz@cs.princeton.edu
+1 (609) 258-2398
<http://www.haoyuzhang.org>

ACADEMICS

- 09/2013 – *Present* Ph.D. student in Computer Science, **Princeton University**
M.A. in Computer Science, awarded in Summer 2015
Advisor: Michael J. Freedman GPA: 3.95 / 4.00
- 09/2009 – 06/2013 B.Sc. in Computer Science, **Peking University**, Beijing, China
GPA (overall): 3.77 / 4.00 GPA (major): 3.85 / 4.00 Rank: 3 / 150
B.A. in Economics (double-major), **Peking University**
- 09/2012 – 02/2013 Visiting Student in Faculty of Electrical Engineering
Technion—Israel Institute of Technology, Haifa, Israel

RESEARCH INTERESTS

My research interests lie in scalable, reliable and highly efficient distributed systems. My recent research focuses on cluster resource scheduling for big-data analytics and machine learning systems, strongly consistent and ultrafast key-value service using programmable data-plane networks (P4), and controller fault tolerance in software-defined networking (SDN).

EXPERIENCE

- 06/2017 – 08/2017 **Ph.D. Software Engineer Intern** Facebook, Menlo Park
Mentors: Brian Cho and Ergin Seyfe
- Designed and implemented an optimized shuffle service for Spark big-data analytics jobs on resource-disaggregated clusters at Facebook scale.
 - The system led to 10x reduction in IO requests and 40% better resource efficiency in production.
- 06/2015 – 08/2015 **Research Intern** Microsoft Research, Redmond
Mentors: Ganesh Ananthanarayanan and Peter Bodik
- Designed and implemented VideoStorm, a video analytics system that processes thousands of vision queries on live video streams over large clusters, with input from operational traffic cameras.
 - VideoStorm is currently deployed and running in Bellevue Traffic Department, and soon in more cities.
- 09/2014 – 05/2015 **Assistant Instructor, Introduction to Programming Systems** Princeton University
- Precept instructor and grader. Course topics span C, Unix and assembly programming.

PROFESSIONAL SKILLS

- Programming** C/C++, Java, Python, C#, Scala, Assembly, Javascript, SQL, Bash scripting
- Unix Tools** \LaTeX , Makefile, version-control tools, plotting tools, profiling tools, Vim, etc.
- Mathematics** Mathematical analysis, discrete mathematics, probability and statistics, linear algebra
- Systems Field** Operating systems, networking, database, and computer architecture
- Other Fields** Machine learning, compilation, software engineering, and digital image processing

PUBLICATIONS

- **Haoyu Zhang**, Brian Cho, Ergin Seyfe, Avery Ching, Michael J. Freedman. Riffle: Optimized Shuffle Service for Large-Scale Data Analytics. In *European Conference on Computer Systems (EuroSys '18)*. Porto, Portugal, April 2018
- Xin Jin, Xiaozhou Li, **Haoyu Zhang**, Nate Foster, Jeongkeun Lee, Robert Soulé, Changhoon Kim, Ion Stoica. NetChain: Scale-Free Sub-RTT Coordination. In *Proc. 15th USENIX Symposium on Networked Systems Design and Implementation (NSDI '18)*. Renton, WA, USA, April 2018
- **Haoyu Zhang**, Logan Stafman, Andrew Or, Michael J. Freedman. SLAQ: Quality-Driven Scheduling for Distributed Machine Learning. In *Proc. ACM Symposium on Cloud Computing (SoCC '17)*. Santa Clara, CA, USA, September 2017
- Xin Jin, Xiaozhou Li, **Haoyu Zhang**, Robert Soulé, Nate Foster, Jeongkeun Lee, Changhoon Kim, Ion Stoica. NetCache: Balancing Key-Value Stores with Fast In-Network Caching. In *Proc. 26th ACM Symposium on Operating Systems Principles (SOSP '17)*. Shanghai, China, October 2017
- **Haoyu Zhang**, Ganesh Ananthanarayanan, Peter Bodik, Matthai Philipose, Paramvir Bahl, Michael J. Freedman. Live Video Analytics at Scale with Approximation and Delay-Tolerance. In *Proc. 14th USENIX Symposium on Networked Systems Design and Implementation (NSDI '17)*. Boston, MA, USA, March 2017
- Naga Katta, **Haoyu Zhang**, Michael J. Freedman, Jennifer Rexford. Ravana: Controller Fault-Tolerance in Software-Defined Networking. In *Proc. ACM Symposium on SDN Research (SOSR '15)*. Santa Clara, CA, USA, June 2015
- **Haoyu Zhang**. Automating the Implementation for Web-based Distributed Systems from State Machine Modeling Language. *Senior Thesis*, Peking University. May 2013

Workshop Papers and Posters

- **Haoyu Zhang**, Logan Stafman, Andrew Or, Michael J. Freedman. SLAQ: Quality-Driven Scheduling for Distributed Machine Learning. In *1st SysML Conference*. Stanford, CA, USA, February 2018
- Xin Jin, Xiaozhou Li, **Haoyu Zhang**, Robert Soulé, Nate Foster, Jeongkeun Lee, Changhoon Kim, Ion Stoica. NetCache: Balancing Key-Value Stores with Fast In-Network Caching. In *SOSP Posters*. Shanghai, China, October 2017
- **Haoyu Zhang**, Logan Stafman, Andrew Or, Michael J. Freedman. SLAQ: Quality-Driven Scheduling for Distributed Machine Learning. In *SoCC Posters*. Santa Clara, CA, USA, September 2017
- Xin Jin, Xiaozhou Li, **Haoyu Zhang**, Robert Soulé, Nate Foster, Jeongkeun Lee, Changhoon Kim, Ion Stoica. NetCache: Fast In-Network Caching for Key-Value Stores. In *P4 Workshop*. Stanford, CA, USA, May 2017

PROFESSIONAL SERVICES

2018	Reviewer: IEEE JSAC
2017	Reviewer: ACM TACO, NIPS MLSys Workshop; PC Member: MobiSys PhD Forum

HONORS

09/2017	Best Paper Award, ACM Symposium on Cloud Computing
09/2013	Princeton University Graduate Fellowship
10/2012	8508 Alumni Scholarship Award
10/2011	Tian Chuang Scholarship Award
10/2010	Outstanding Student Award at Peking University