

Interests

Distributed and networked systems

Education

- 2015 Ph.D. in Computer Science
University of Illinois at Urbana-Champaign
- 2009 B.Tech.(Honours) in Computer Science and Engineering
Indian Institute of Technology at Kharagpur

Employment

- 2019 – present Staff Software Engineer
Google
- 2015 – 2019 Staff Researcher
Samsung Research America, Mountain View
- 2009 – 2015 Research Assistant, Dept. of Computer Science
University of Illinois at Urbana-Champaign
- Summer 2014 Intern, Microsoft CISL, Silicon Valley
- Summer 2013 Intern, Microsoft Research, Silicon Valley
- Fall 2012 Teaching Assistant, Fundamental Algorithms (CS 473), UIUC
- Spring 2012 Teaching Assistant, CS Undergrad Research Lab, UIUC
- Summer 2011 Intern, Microsoft Research, Cambridge
- Summer 2008 Intern, Microsoft Research, Redmond
- 2007 – 2008 Software Engineer (part-time), Minekey Inc.

Selected Research Projects

- **Scaling analytics via approximation** *Samsung Research, 2015 – 2019*
Worked on using data approximation to reduce query latency and scale ingest performance when supporting analytics and learning on high-velocity, high-volume data. I built SummaryStore [SOSP'17, open-source on github], a high-performance approximate single-node data store capable of supporting low-latency analytics on more than a PB of stored data with high accuracy. I'm currently working on using data approximation to better scale data replication and availability in distributed data stores.
- **Geo-distributed analytics** *UIUC & Microsoft, 2013 – 2015*
I built systems to support analytics on massive volumes of data partitioned across several data centers around the world, incorporating a range of query plan and data transfer optimizations to reduce bandwidth cost. The systems I built [NSDI'15, CIDR'15] operate as drop-in multi data center replacements for Apache Hive and Oozie, well-known single data center analytics frameworks.
- **Reducing latency via redundancy** *UIUC, 2011 – 2013*
Investigated request replication as a general technique for reducing latency. I evaluated the tradeoff between added overhead and reduced latency [HotNets'12, CoNEXT'13] both empirically, through large-scale experiments in several systems including key-value stores (Couchbase, memcached), DNS resolvers, the TCP handshake, and multipath networks; and analytically, in a queueing-theoretic model.
- **Mitigating denial-of-service attacks via congestion pricing** *UIUC, 2010*
Evaluated congestion pricing as a way to mitigate DoS attacks, via a game-theoretic analysis backed by large-scale simulations.

Publications

- Nitin Agrawal, Ashish Vulimiri. “Building highly-available geo-distributed data stores for continuous learning”, *Systems for ML Workshop, NIPS 2018*
- Nitin Agrawal, Ashish Vulimiri. “Low-latency analytics on colossal data streams with SummaryStore”, *SOSP 2017*
- Nitin Agrawal, Ashish Vulimiri. “Learning with less: Can approximate storage systems save learning from drowning in data?”, *AI Systems Workshop, SOSP 2017*
- Ashish Vulimiri, Carlo Curino, Brighten Godfrey, Thomas Jungblut, Jitu Padhye, George Varghese. “Global analytics in the face of bandwidth and regulatory constraints”, *NSDI 2015*
- Ashish Vulimiri, Carlo Curino, Brighten Godfrey, Thomas Jungblut, Konstantinos Karanasos, Jitu Padhye, George Varghese. “WANalytics: Geo-distributed analytics for a data-intensive world”, demo at *SIGMOD 2015*
- Ashish Vulimiri, Carlo Curino, Brighten Godfrey, Konstantinos Karanasos, George Varghese. “WANalytics: Analytics for a geo-distributed data-intensive world”, *CIDR 2015*
- Ashish Vulimiri, Brighten Godfrey, Radhika Mittal, Justine Sherry, Sylvia Ratnasamy, Scott Shenker. “Low latency via redundancy”, *CoNEXT 2013*
- Ashish Vulimiri, Oliver Michel, Brighten Godfrey, Scott Shenker. “More is less: Reducing latency via redundancy”, *HotNets 2012*
- Ashish Vulimiri, Gul A. Agha, Brighten Godfrey, Karthik Lakshminarayanan. “How well can congestion pricing neutralize denial-of-service attacks?”, *SIGMETRICS 2012*
- Anjan Sarkar, Ashish Vulimiri, Suman Paul, Mohammed J. Iqbal, Avishek Banerjee, Rahul Chatterjee, Shibendu S. Ray. “Unsupervised and supervised classification of hyperspectral image data using projection pursuit and Markov random field segmentation”, *International Journal of Remote Sensing*, Vol. 33 Issue 18, 2012
- Ashish Vulimiri, Arobinda Gupta, Prमित Roy, Skanda N. Muthaiah, Arzad A. Kherani. “Application of secondary information for misbehavior detection in VANETs”, *IFIP Networking 2010*
- Anjan Sarkar, Ashish Vulimiri, Shantanu Bose, Suman Paul, Shibendu S. Ray. “Unsupervised hyperspectral image analysis with projection pursuit and MRF segmentation approach”, in *2008 International Conference on Artificial Intelligence and Pattern Recognition (AIPR-08)*

Professional Activities

- Member of the program committee for USENIX Conf. on File and Storage Technologies (FAST) 2018
- Reviewer for: Trans. on Computer Systems, Trans. on Parallel and Distributed Systems, Trans. on Big Data, Trans. on Cloud Computing, Trans. on Networking, Trans. on Communications, IEEE Network, ACM Computer Communications Review