

Matei Zaharia

<http://cs.stanford.edu/~matei>

matei@cs.stanford.edu

EDUCATION

University of California at Berkeley 2007–2013
PhD, Computer Science
Topic: [An Architecture for Fast and General Data Processing on Large Clusters](#)
Advisors: Scott Shenker and Ion Stoica

University of Waterloo 2003–2007
Bachelor of Mathematics, Honors Computer Science
Minor in Combinatorics and Optimization

AWARDS

- NSF CAREER Award, 2017
- VMware Systems Research Award, 2016
- ACM Doctoral Dissertation Award, 2014
- U. Waterloo Faculty of Mathematics Young Alumni Achievement Medal, 2014
- Daytona GraySort World Record, 2014
- David J. Sakrison Prize for research, UC Berkeley, 2013
- Best Paper Award, SIGCOMM 2012
- Best Paper Award, NSDI 2012
- Honorable Mention for Community Award, NSDI 2012
- Best Demo Award, SIGMOD 2012
- Google PhD Fellowship, 2011–2012
- Tong Leon Lim Pre-Doctoral Prize for highest distinction in computer science preliminary examinations, 2009
- National Sciences and Engineering Research Council of Canada (NSERC) Post-graduate Scholarship, 2009–2011
- Julie Payette NSERC Research Scholarship, 2007–2009
- Governor General's Academic Silver Medal for highest academic standing upon graduation from the University of Waterloo

RESEARCH POSITIONS

Assistant Professor, Stanford University 2016–present
Projects in computer systems and big data:

- *Weld*, a runtime for data-parallel computations that can optimize across widely used libraries like NumPy and Pandas and leverage heterogeneous hardware
- *Splinter*, a system for making private queries against public data (*e.g.*, online maps) with sub-second response times for realistic applications
- *NoScope*, a system for 100–1000× faster neural network based queries over video
- *Usable machine learning* systems that automatically perform model selection, tuning and monitoring for specific problems such as recommendation engines

Assistant Professor, Massachusetts Institute of Technology 2015–2016
Projects in computer systems and big data:

- *Vuvuzela and Stadium*, private messaging systems resistant to traffic analysis

- *Spark SQL and DataFrames*, a relational API for the Spark engine allowing rich optimization of user code underneath a familiar interface
- *High-performance analytics* projects including GraphFrames (relational API for pattern matching) and Yggdrasil (distributed decision tree learning)

PhD Student, University of California at Berkeley 2007–2013

Cluster computing:

- *Spark*, a system that provides fault-tolerant distributed memory abstractions to run iterative algorithms and interactive queries on large clusters
- *Spark Streaming*, a highly scalable stream processing engine
- *Shark*, a fast and fault-tolerant SQL engine built on Spark
- *Mesos*, a system that enables resource sharing across diverse applications by letting them control their own scheduling
- *Dominant resource fairness (DRF)*, a generalization of weighted fair sharing for multiple resources that preserves most of its attractive properties
- *Delay scheduling*, an algorithm for achieving data locality in parallel applications that is now included in Hadoop MapReduce
- *Longest approximate time to end (LATE)*, an algorithm for mitigating slow nodes in MapReduce that is also now included in Hadoop

Bioinformatics:

- *SNAP*, a sequence alignment algorithm that is 10–100× faster and simultaneously more accurate than state-of-the-art tools
- *SURPI*, a low-latency pathogen identification pipeline based on SNAP that can be used to test samples against large databases of known organisms

Networking:

- *Dominant resource fair queueing (DRFQ)*, a generalization of fair queueing for systems that consume multiple resources per packet (*e.g.*, middleboxes)
- *Orchestra*, a new architecture for optimizing communication in datacenters that considers the semantics of parallel transfer operations (sets of flows)
- *Cloud Terminal*, a minimal software platform that provides secure access to sensitive applications (*e.g.*, banks) even if a user’s OS is compromised

**INDUSTRY
EXPERIENCE**

Chief Technologist, Databricks 2013–present

- Databricks is a startup company offering a cloud-hosted big data processing service based on Apache Spark to over 400 customers

Consultant, Quantifind Fall 2012

- Helped optimize Spark applications for predictive analytics from social data

Contractor, Cloudera 2008–2011

- Developed new features for the Hadoop Fair Scheduler to support customer workloads and contributed these changes into the open source Hadoop project

Software Engineer Intern, Facebook Summer 2008

- Developed a fair job scheduler for the Hadoop MapReduce platform; system is in production use and is part of open-source Hadoop distribution
- Built internal Hadoop monitoring and accounting dashboard

Software Engineer Intern, Google Summer 2007

- Developed new features for a code searching system, GTags, used internally by Google engineers and available externally as open source

OPEN SOURCE

Vice President, Apache Spark Project 2013–present

- I started Spark as a research project at UC Berkeley and serve as its Vice President at the Apache Software Foundation

Committer, Apache Mesos Project 2011–present

- I helped to start Mesos as an academic project at Berkeley

Committer, Apache Hadoop Project 2009–present

- Committer status is given to individuals that make significant contributions
- My work included straggler mitigation logic and the Hadoop Fair Scheduler

PUBLICATIONS

Conference Papers

- N. Tyagi, Y. Gilad, D. Leung, M. Zaharia and N. Zeldovich. Stadium: A Distributed Metadata-Private Messaging System, *to appear at SOSP 2017*
- D. Kang, J. Emmons, F. Abuzaid, P. Bailis and M. Zaharia. [NoScope: Optimizing Neural Network Queries over Video at Scale](#), *VLDB 2017*
- F. Wang, C. Yun, S. Goldwasser, V. Vaikuntanathan and M. Zaharia. [Splinter: Practical Private Queries on Public Data](#), *NSDI 2017*
- S. Palkar, J. Thomas, A. Shanbhag, M. Schwarzkopf, S. Amarasinghe and M. Zaharia. [Weld: A Common Runtime for High-Performance Data Analysis](#), *CIDR 2017*
- F. Abuzaid, J. Bradley, F. Liang, A. Feng, L. Yang, M. Zaharia and A. Talwalkar. Yggdrasil: An Optimized System for Training Deep Decision Trees at Scale, *to appear at NIPS 2016*
- R.B. Zadeh, X. Meng, A. Staple, B. Yavuz, L. Pu, S. Venkataraman, E. Sparks, A. Ulanov and M. Zaharia. [Matrix Computations and Optimizations in Apache Spark](#), *KDD 2016 (best paper runner-up)*.
- S. Venkataraman, Z. Yang, D. Liu, E. Liang, X. Meng, R. Xin, A. Ghodsi, M. Franklin, I. Stoica and M. Zaharia. [SparkR: Scaling R Programs with Spark](#), *SIGMOD 2016 Industry Track*.
- Q. Pu, H. Li, M. Zaharia, A. Ghodsi, and I. Stoica. [FairRide: Near-Optimal, Fair Cache Sharing](#), *NSDI 2016*
- J. van den Hooff, D. Lazar, M. Zaharia and N. Zeldovich. [Vuvuzela: Scalable Private Messaging Resistant to Traffic Analysis](#), *SOSP 2015*
- M. Armbrust, T. Das, A. Davidson, A. Ghodsi, A. Or, J. Rosen, I. Stoica, P. Wendell, R. Xin, and M. Zaharia. [Scaling Spark in the Real World: Performance and Usability](#), *VLDB 2015 Industry Track*
- M. Armbrust, R. Xin, C. Lian, Y. Huai, D. Liu, J. Bradley, X. Meng, T. Kaftan, M. Franklin, A. Ghodsi and M. Zaharia. [Spark SQL: Relational Data Processing in Spark](#), *SIGMOD 2015 Industry Track*
- H. Li, A. Ghodsi, M. Zaharia, S. Shenker and I. Stoica. [Tachyon: Reliable, Memory Speed Storage for Cluster Computing Frameworks](#), *SOCC 2014*
- M. Zaharia, T. Das, H. Li, T. Hunter, S. Shenker, and I. Stoica. [Discretized Streams: Fault-Tolerant Streaming Computation at Scale](#), *SOSP 2013*
- K. Ousterhout, P. Wendell, M. Zaharia and I. Stoica. [Sparrow: Distributed, Low-Latency Scheduling](#), *SOSP 2013*

- R. Xin, J. Rosen, M. Zaharia, M. Franklin, S. Shenker, and I. Stoica. [Shark: SQL and Rich Analytics at Scale](#), *SIGMOD 2013*
- A. Ghodsi, M. Zaharia, S. Shenker and I. Stoica. [Choosy: Max-Min Fair Sharing for Datacenter Jobs with Constraints](#), *EuroSys 2013*
- A. Ghodsi, V. Sekar, M. Zaharia and I. Stoica. [Multi-Resource Fair Queueing for Packet Processing](#), *SIGCOMM 2012 (best paper award)*
- L. Martignoni, P. Poosankam, M. Zaharia, J. Han, S. McCamant, D. Song, V. Paxson, A. Perrig, S. Shenker, I. Stoica. [Cloud Terminal: Secure Access to Sensitive Applications from Untrusted Systems](#), *USENIX ATC 2012*
- M. Zaharia, M. Chowdhury, T. Das, A. Dave, J. Ma, M. McCauley, M.J. Franklin, S. Shenker, I. Stoica. [Resilient Distributed Datasets: A Fault-Tolerant Abstraction for In-Memory Cluster Computing](#), *NSDI 2012 (best paper award)*
- A.N. Rafferty, M. Zaharia and T.L. Griffiths. [Optimally Designing Games for Cognitive Science Research](#), *Annual Conf. of the Cognitive Science Society, 2012*
- T. Hunter, T. Moldovan, M. Zaharia, S. Merzgui, J. Ma, M.J. Franklin, P. Abbeel, and A.M. Bayen. [Scaling the Mobile Millennium System in the Cloud](#), *ACM Symposium on Cloud Computing (SoCC) 2011*
- M. Chowdhury, M. Zaharia, J. Ma, M.I. Jordan and I. Stoica. [Managing Data Transfers in Computer Clusters with Orchestra](#), *SIGCOMM 2011*
- A. Ghodsi, M. Zaharia, B. Hindman, A. Konwinski, S. Shenker and I. Stoica. [Dominant Resource Fairness: Fair Allocation of Multiple Resources Types](#), *NSDI 2011*
- B. Hindman, A. Konwinski, M. Zaharia, A. Ghodsi, A.D. Joseph, R. Katz, S. Shenker and I. Stoica. [Mesos: A Platform for Fine-Grained Resource Sharing in the Data Center](#), *NSDI 2011*
- M. Zaharia, D. Borthakur, J. Sen Sarma, K. Elmeleegy, S. Shenker and I. Stoica. [Delay Scheduling: A Simple Technique for Achieving Locality and Fairness in Cluster Scheduling](#), *EuroSys 2010*
- R. Luk, M. Zaharia, M. Ho, B. Levine and P. Aoki, [ICTD for Healthcare in Ghana: Two Parallel Case Studies](#), *ICTD 2009*
- M. Zaharia, A. Konwinski, A.D. Joseph, R. Katz and I. Stoica, [Improving MapReduce Performance in Heterogeneous Environments](#), *OSDI 2008*
- S. Guo, M.H. Falaki, E.A. Oliver, S. Ur Rahman, A. Seth, M. Zaharia, U. Ismail, and S. Keshav, [Design and Implementation of the KioskNet System](#), *ICTD 2007*
- A. Seth, D. Kroeker, M. Zaharia, S. Guo and S. Keshav, [Low-cost Communication for Rural Internet Kiosks Using Mechanical Backhaul](#), *MOBICOM 2006*

Journal Articles

- M. Zaharia, R. Xin, P. Wendell, T. Das, M. Armbrust, A. Dave, X. Meng, J. Rosen, S. Venkataraman, M. Franklin, A. Ghodsi, J. Gonzalez, S. Shenker, I. Stoica, [Apache Spark: A Unified Engine for Big Data Processing](#), *Communications of the ACM*, 59(11): 56–65, 2016.
- X. Meng, J. Bradley, B. Yuvaz, E. Sparks, S. Venkataraman, D. Liu, J. Freeman, D. Tsai, M. Amde, S. Owen, D. Xin, R. Xin, M. Franklin, R. Zadeh, M. Zaharia, and A. Talwalkar. [MLlib: Machine Learning in Apache Spark](#), *JMLR*, 17(34): 1–7, 2016
- S.N. Naccache, S. Federman, N. Veeeraraghavan, M. Zaharia, D. Lee, E. Samayoa, J. Bouquet, A.L. Greninger, K. Luk, B. Enge, D.A. Wadford, S.L. Messenger, G.L. Genrich, K. Pellegrino, G. Grard, E. Leroy, B.S. Schneider, J.N. Fair, M.A.

Martinez, P. Isa, J.A. Crump, J.L. DeRisi, T. Sittler, J. Hackett Jr., S. Miller and C.Y. Chiu, [A Cloud-Compatible Bioinformatics Pipeline for Ultrarapid Pathogen Identification from Next-Generation Sequencing of Clinical Samples](#), *Genome Research*, 24(7): 1180–1192, 2014

- A.N. Rafferty, M. Zaharia, and T.L. Griffiths. [Optimally designing games for behavioural research](#), *Proc. of the Royal Society Series A*, 470, 2014
- S. Guo, M. Derakhshani, M.H. Falaki, U. Ismail, R. Luk, E.A. Oliver, S. Ur Rahman, A. Seth, M. Zaharia, and S. Keshav, [Design and Implementation of the KioskNet System](#), *Computer Networks*, 55(1): 264–281, 2011
- M. Armbrust, A. Fox, R. Griffith, A.D. Joseph, R.H. Katz, A. Konwinski, G. Lee, D.A. Patterson, A. Rabkin, I. Stoica and M. Zaharia, [Above the Clouds: A View of Cloud Computing](#), *Communications of the ACM*, 53(4): 50–58, 2010.
- M. Zaharia and S. Keshav, [Gossip-based Search Selection in Hybrid Peer-to-Peer Networks](#), *J. Concurrency and Computation: Practice and Experience*, 20(2): 139–153, 2007

Workshop Papers

- C. Coleman, D. Narayanan, D. Kang, T. Zhao, J. Zhang, L. Nardi, P. Bailis, K. Olukotun, C. ReIA and M. Zaharia. [DAWNBench: An End-to-End Deep Learning Benchmark and Competition](#), *to appear at AIsys 2017*.
- A. Dave, A. Jindal, L.E. Li, R. Xin, J. Gonzalez and M. Zaharia. [GraphFrames: An Integrated API for Mixing Graph and Relational Queries](#), *GRADES 2016*.
- M. Vartak, H. Subramanyam, W.E. Lee, S. Viswanathan, S. Husnoo, S. Madden and M. Zaharia. [ModelDB: A System for Machine Learning Model Management](#), *HILDA 2016*.
- H. Li, A. Ghodsi, M. Zaharia, E. Baldeschwieler, S. Shenker, I. Stoica. [Tachyon: Memory Throughput I/O for Cluster Computing Frameworks](#), *LADIS 2013*
- M. Zaharia, T. Das, H. Li, S. Shenker and I. Stoica. [Discretized Streams: An Efficient and Fault-Tolerant Model for Stream Processing on Large Clusters](#), *USENIX HotCloud 2012*
- M. Zaharia, B. Hindman, A. Konwinski, A. Ghodsi, A.D. Joseph, R. Katz, S. Shenker and I. Stoica. [The Datacenter Needs an Operating System](#), *USENIX HotCloud 2011*
- M. Zaharia, M. Chowdhury, M.J. Franklin, S. Shenker and I. Stoica, [Spark: Cluster Computing with Working Sets](#), *USENIX HotCloud 2010*
- B. Hindman, A. Konwinski, M. Zaharia and I. Stoica, [A Common Substrate for Cluster Computing](#), *USENIX HotCloud 2009*
- M. Zaharia, A. Chandel, S. Saroiu and S. Keshav, [Finding Content in File-Sharing Networks When You Can't Even Spell](#), *IPTPS 2007*
- M. Zaharia and S. Keshav, [Gossip-Based Search Selection in Hybrid Peer-to-Peer Networks](#), *IPTPS 2006*

Invited Papers

- M. Zaharia, M. Chowdhury, T. Das, A. Dave, J. Ma, M. McCauley, M.J. Franklin, S. Shenker, I. Stoica. [Fast and Interactive Analytics over Hadoop Data with Spark](#), *USENIX ;login.*, August 2012.
- B. Hindman, A. Konwinski, M. Zaharia, A. Ghodsi, A.D. Joseph, R. Katz, S. Shenker and I. Stoica. [Mesos: Flexible Resource Sharing for the Cloud](#), *USENIX ;login.*, August 2011.

- S. Guo, M.H. Falaki, E.A. Oliver, S. Ur Rahman, A. Seth, M. Zaharia, and S. Keshav, [Very Low-Cost Internet Access Using KioskNet](#), *ACM Computer Communication Review*, October 2007

Demonstrations

- C. Engle, A. Lupher, R. Xin, M. Zaharia, M. Franklin, S. Shenker, I. Stoica. [Shark: Fast Data Analysis Using Coarse-grained Distributed Memory](#), *SIGMOD 2012* (best demo award)

Technical Reports

- Y. Zhang, V. Kiriansky, C. Mendis, M. Zaharia and S. Amarasinghe. [Optimizing Cache Performance for Graph Analytics](#), *CoRR abs/1608.01362v2*, August 2016
- M. Zaharia, T. Das, H. Li, T. Hunter S. Shenker, and I. Stoica. [Discretized Streams: A Fault-Tolerant Model for Scalable Stream Processing](#), *UC Berkeley Technical Report UCB/EECS-2012-259*, December 2012
- R. Xin, J. Rosen, M. Zaharia, M.J. Franklin, S. Shenker, and I. Stoica. [Shark: SQL and Rich Analytics at Scale](#), *UC Berkeley Technical Report UCB/EECS-2012-213*, November 2012
- M. Zaharia, S. Katti, C. Grier, V. Paxson, S. Shenker, I. Stoica, and D. Song. [Hypervisors as a Foothold for Personal Computer Security: An Agenda for the Research Community](#), *UC Berkeley Technical Report UCB/EECS-2012-12*, January 2012
- M. Zaharia, W.J. Bolosky, K. Curtis, A. Fox, D. Patterson, S. Shenker, I. Stoica, R.M. Karp, and T. Sittler, [Faster and More Accurate Sequence Alignment with SNAP](#), *CoRR abs/1111.5572v1*, November 2011
- G. Ananthanarayanan, K. Heimerl, M. Zaharia, M. Demmer, T. Koponen, A. Tavakoli, S. Shenker and I. Stoica, [A New Communication API](#), *UC Berkeley Technical Report UCB/EECS-2009-84*, May 2009
- M. Zaharia, D. Borthakur, J. Sen Sarma, K. Elmeleegy, S. Shenker and I. Stoica, [Job Scheduling for Multi-User MapReduce Clusters](#), *UC Berkeley Technical Report UCB/EECS-2009-55*, April 2009
- M. Armbrust, A. Fox, R. Griffith, A.D. Joseph, R.H. Katz, A. Konwinski, G. Lee, D.A. Patterson, A. Rabkin, I. Stoica and M. Zaharia, [Above the Clouds: A Berkeley View of Cloud Computing](#), *UC Berkeley Technical Report UCB/EECS-2009-28*, February 2009
- M. Zaharia and S. Keshav, [Fast and Optimal Scheduling Over Multiple Network Interfaces](#), *University of Waterloo Technical Report CS-2007-36*, October 2007
- M. Zaharia and S. Keshav, [Adaptive Peer-to-Peer Search](#), *University of Waterloo Technical Report 2004-55*, November 2004