

Andy Pavlo

Carnegie Mellon University
Department of Computer Science
Gates-Hillman Center 9019
Pittsburgh, PA 15213-3891 USA

Voice: +1-844-88-CMUDB
E-mail: pavlo@cs.cmu.edu
Web: <http://www.cs.cmu.edu/~pavlo/>
Twitter: @andy_pavlo

Education

B.Sc. with Honors, Computer Science Rochester Institute of Technology	Spring 2005 <i>Rochester, NY USA</i>
M.Sc., Computer Science Rochester Institute of Technology	Spring 2006 <i>Rochester, NY USA</i>
M.Sc., Computer Science Brown University	Spring 2009 <i>Providence, RI USA</i>
Ph.D., Computer Science Brown University	Summer 2013 <i>Providence, RI USA</i>

Appointments

Assistant Professor Carnegie Mellon University, Department of Computer Science	September 2013 – present <i>Pittsburgh, PA USA</i>
Courtesy Assistant Professor Carnegie Mellon University, Department of Electrical & Computer Engineering	January 2014 – present <i>Pittsburgh, PA USA</i>

Professional Experience

Expert Witness Much Shelist, P.C.	May 2014 – May 2016 <i>Chicago, IL USA</i>
Research Assistant Brown University, Department of Computer Science	September 2007 – August 2013 <i>Providence, RI USA</i>
Non-Testifying Technical Expert Choate, Hall & Stewart, LLP.	June 2009 – November 2011 <i>Boston, MA USA</i>
Research Intern Vertica, Inc.	June 2008 – August 2008 <i>Billerica, MA USA</i>
Systems Programmer – Condor Project University of Wisconsin–Madison, Department of Computer Sciences	September 2005 – August 2007 <i>Madison, WI USA</i>
Software Developer PAID, Inc.	March 1999 – December 2004 <i>Owings Mills, MD USA</i>

Teaching Experience

- **Advanced Database Systems** (Carnegie Mellon University, 15-721) – Spring 2019
- **Intro to Database Systems** (Carnegie Mellon University, 15-445/645) – Fall 2018
- **Advanced Database Systems** (Carnegie Mellon University, 15-721) – Spring 2018

- **Intro to Database Systems** (Carnegie Mellon University, 15-445/645) – Fall 2017
- **Advanced Database Systems** (Carnegie Mellon University, 15-721) – Spring 2017
- **Database Applications** (Carnegie Mellon University, 15-415/615) – Fall 2016
- **Advanced Database Systems** (Carnegie Mellon University, 15-721) – Spring 2016
- **Database Applications** (Carnegie Mellon University, 15-415/615) – Fall 2015
- **Database Applications** (Carnegie Mellon University, 15-415/615) – Spring 2015
- **Database Applications** (Carnegie Mellon University, 15-415/615) – Spring 2014
- **Advanced Topics in Database Systems** (Carnegie Mellon University, 15-799) – Fall 2013
- **NewSQL – The Course** (Brown University, CSC2270) – Spring 2012
- **Non-standard, Web-scale Databases** (Brown University, CSC2270) – Spring 2011

Mentoring

PhD Students

- Xiangyao Yu (2013 – 2017)
- Joy Arulraj (2013 – 2018)
- Dana Van Aken (2014 – present)
- Prashanth Menon (2015 – present)
- Lin Ma (2015 – present)

Master's Students

- Prashasthi Prabhakar (Carnegie Mellon University, M.Sc. 2017)
- Patrick Huang (Carnegie Mellon University, M.Sc. 2017)
- Matt Perron (Carnegie Mellon University, M.Sc. 2016)
- Jiexi Lin (Carnegie Mellon University, M.Sc. 2016)
- Siddharth Santurkar (Carnegie Mellon University, M.Sc. 2016)
- Haibin Lin (Carnegie Mellon University, M.Sc. 2016)
- Thomas Marshall (Carnegie Mellon University, M.Sc. 2015)
- Atreyee Maiti (Carnegie Mellon University, M.Sc. 2014)
- Xin Jia (Brown University, M.Sc. 2013)
- Yang Zou (Brown University, M.Sc. 2012)
- Yang Lu (Brown University, M.Sc. 2012)
- Visawee Angkanawaraphan (Brown University, M.Sc. 2011)
- Zhe Zhang (Brown University, M.Sc. 2010)
- Ning Shi (Brown University, M.Sc. 2010)

Awards & Honors

- VLDB Distinguished Associate Editor, Summer 2018
- ACM SIGMOD Best Paper Award, Summer 2018
- ACM SIGMOD Distinguished Program Committee Member, Summer 2018
- CMU Joel & Ruth Spira Teaching Award, Spring 2018
- Sloan Research Fellowship, Winter 2018
- Google Faculty Research Award, Winter 2018
- ACM SIGMOD Distinguished Program Committee Member, Summer 2017
- Facebook Faculty Research Award, Summer 2015
- Google Faculty Research Award, Summer 2015
- ACM SIGMOD Jim Gray Dissertation Award, Summer 2014
- Amazon Research Grant, Winter 2012
- ACM SIGMOD Student Travel Grant, Spring 2010
- Class Clown, Mt. Hebron High School, Spring 1999

Service

To the University

- CSD Faculty Search Committee – 2018
- CSD MS Admissions Committee – 2017
- CSD PhD Admissions Committee – 2016
- CSD PhD Admissions Committee – 2015
- CSD PhD Admissions Committee – 2014

To the Profession

- Program Committee – CIDR 2019
- Program Committee – SIGMOD 2019
- Program Committee – VLDB 2019
- Associate Editor – VLDB 2018
- Program Committee – SIGMOD 2018
- Co-Chair – FADS@VLDB 2017
- Program Committee / Demo Co-Chair – SIGMOD 2017
- Program Committee – CIDR 2017
- Program Committee Area Chair – ICDE 2016
- Program Committee / PhD Workshop – VLDB 2016
- Co-Chair – IMDM 2016
- Program Committee – SIGMOD 2016
- Steering Committee / Program Committee – HPTS 2015
- Co-Chair – IMDM 2015
- Program Committee / Demo Committee / Dissertation Award Committee – SIGMOD 2015
- Program Committee / Demo Committee – VLDB 2015
- Co-Chair – IMDM 2014
- Program Committee – SIGMOD 2014
- Program Committee – SoCC 2014
- Program Committee – CloudDP 2014
- Program Committee – VLDB 2014
- External Reviewer – LNCS Transactions on Large-Scale Data- and Knowledge-Centered Systems, March 2013
- External Reviewer – Springer Journal of Distributed and Parallel Databases, January 2013
- External Reviewer – IBM Journal of Research and Development, July 2012
- Demo Judge – SIGMOD 2010
- Demo Judge – SIGMOD 2009

Publications

Refereed

- [1] Timo Kersten, Viktor Leis, Alfons Kemper, Thomas Neumann, Andrew Pavlo, and Peter Boncz. “Everything You Always Wanted to Know About Compiled and Vectorized Queries But Were Afraid to Ask”. In: *Proc. VLDB Endow.* 11 (13 Sept. 2018), pp. 2209–2222.
- [2] Xiangyao Yu, Yu Xia, Andrew Pavlo, Daniel Sanchez, Larry Rudolph, and Srinivas Devadas. “Harmonizing Concurrency Control and Caching in a Distributed OLTP Database Management System”. In: *Proc. VLDB Endow.* 11.10 (June 2018), pp. 1289–1302.
- [3] Ziqi Wang, Andrew Pavlo, Hyeontaek Lim, Viktor Leis, Huanchen Zhang, Michael Kaminsky, and David G. Andersen. “Building a Bw-Tree Takes More Than Just Buzz Words”. In: *Proceedings of the 2018 ACM International Conference on Management of Data*. SIGMOD ’18. 2018, pp. 473–488.
- [4] Lin Ma, Dana Van Aken, Ahmed Hefny, Gustavo Mezerhane, Andrew Pavlo, and Geoffrey J Gordon. “Query-based Workload Forecasting for Self-Driving Database Management Systems”. In: *Proceedings of the 2018 ACM International Conference on Management of Data*. SIGMOD ’18. 2018, pp. 631–645.

- [5] Huanchen Zhang, David G. Andersen, Michael Kaminsky, Andrew Pavlo, Hyeontaek Lim, Viktor Leis, and Kimberly Keeton. “SuRF: Practical Range Query Filtering with Fast Succinct Tries”. In: *Proceedings of the 2018 ACM International Conference on Management of Data*. SIGMOD ’18. 2018, pp. 323–336.
- [6] Prashanth Menon, Todd C. Mowry, and Andrew Pavlo. “Relaxed Operator Fusion for In-Memory Databases: Making Compilation, Vectorization, and Prefetching Work Together At Last”. In: *Proc. VLDB Endow.* 11.1 (Sept. 2017), pp. 1–13.
- [7] Yingjun Wu, Joy Arulraj, Jiexi Lin, Ran Xian, and Andrew Pavlo. “An Empirical Evaluation of In-Memory Multi-Version Concurrency Control”. In: *Proc. VLDB Endow.* 10.7 (Mar. 2017), pp. 781–792.
- [8] Rachael Harding, Dana Van Aken, Andrew Pavlo, and Michael Stonebraker. “An Evaluation of Distributed Concurrency Control”. In: *Proc. VLDB Endow.* 10.5 (Jan. 2017), pp. 553–564.
- [9] Dana Van Aken, Andrew Pavlo, Geoffrey J. Gordon, and Bohan Zhang. “Automatic Database Management System Tuning Through Large-scale Machine Learning”. In: *Proceedings of the 2017 ACM International Conference on Management of Data*. SIGMOD ’17. 2017, pp. 1009–1024.
- [10] Lianghong Xu, Andrew Pavlo, Sudipta Sengupta, and Gregory R. Ganger. “Online Deduplication for Databases”. In: *Proceedings of the 2017 ACM International Conference on Management of Data*. SIGMOD ’17. 2017, pp. 1355–1368.
- [11] Andrew Pavlo, Gustavo Angulo, Joy Arulraj, Haibin Lin, Jiexi Lin, Lin Ma, Prashanth Menon, Todd Mowry, Matthew Perron, Ian Quah, Siddharth Santurkar, Anthony Tomasic, Skye Toor, Dana Van Aken, Ziqi Wang, Yingjun Wu, Ran Xian, and Tiejing Zhang. “Self-Driving Database Management Systems”. In: *CIDR 2017, Conference on Innovative Data Systems Research*. 2017.
- [12] Joy Arulraj, Matthew Perron, and Andrew Pavlo. “Write-Behind Logging”. In: *Proc. VLDB Endow.* 10.4 (Dec. 2016), pp. 337–348.
- [13] Marco Serafini, Rebecca Taft, Aaron J Elmore, Andrew Pavlo, Ashraf Aboulmaga, and Michael Stonebraker. “Clay: Fine-Grained Adaptive Partitioning for General Database Schemas”. In: *Proc. VLDB Endow.* 10.4 (Dec. 2016), pp. 445–456.
- [14] Lin Ma, Joy Arulraj, Sam Zhao, Andrew Pavlo, Subramanya R. Dulloor, Michael J. Giardino, Jeff Parkhurst, Jason L. Gardner, Kshitij Doshi, and Stanley Zdonik. “Larger-than-memory Data Management on Modern Storage Hardware for In-memory OLTP Database Systems”. In: *Proceedings of the 12th International Workshop on Data Management on New Hardware*. DaMoN ’16. 2016, 9:1–9:7.
- [15] Joy Arulraj, Andrew Pavlo, and Prashanth Menon. “Bridging the Archipelago Between Row-Stores and Column-Stores for Hybrid Workloads”. In: *Proceedings of the 2016 International Conference on Management of Data*. SIGMOD ’16. 2016, pp. 583–598.
- [16] Huanchen Zhang, David G. Andersen, Andrew Pavlo, Michael Kaminsky, Lin Ma, and Rui Shen. “Reducing the Storage Overhead of Main-Memory OLTP Databases with Hybrid Indexes”. In: *Proceedings of the 2016 International Conference on Management of Data*. SIGMOD ’16. 2016, pp. 1567–1581.
- [17] Xiangyao Yu, Andrew Pavlo, Daniel Sanchez, and Srinivas Devadas. “TicToc: Time Traveling Optimistic Concurrency Control”. In: *Proceedings of the 2016 International Conference on Management of Data*. SIGMOD ’16. 2016, pp. 1629–1642.
- [18] John Meehan, Nesime Tatbul, Stan Zdonik, Cansu Aslantas, Ugur Çetintemel, Jiang Du, Tim Kraska, Samuel Madden, David Maier, Andrew Pavlo, Michael Stonebraker, Kristin Tufte, and Hao Wang. “S-Store: Streaming Meets Transaction Processing”. In: *PVLDB* 8.13 (2015), pp. 2134–2145.
- [19] Lianghong Xu, Andrew Pavlo, Sudipta Sengupta, Jin Li, and Gregory R. Ganger. “Reducing Replication Bandwidth for Distributed Document Databases”. In: *Proceedings of the Sixth ACM Symposium on Cloud Computing*. SoCC ’15. 2015, pp. 222–235.
- [20] Aaron J. Elmore, Vaibhav Arora, Rebecca Taft, Andrew Pavlo, Divyakant Agrawal, and Amr El Abbadi. “Squall: Fine-Grained Live Reconfiguration for Partitioned Main Memory Databases”. In: *Proceedings of SIGMOD*. SIGMOD ’15. 2015, pp. 299–313.
- [21] Joy Arulraj, Andrew Pavlo, and Subramanya Dulloor. “Let’s Talk About Storage & Recovery Methods for Non-Volatile Memory Database Systems”. In: *Proceedings of the 2015 International Conference on Management of Data*. SIGMOD ’15. 2015, pp. 707–722.

- [22] Xiangyao Yu, George Bezerra, Andrew Pavlo, Srinivas Devadas, and Michael Stonebraker. “Staring into the Abyss: An Evaluation of Concurrency Control with One Thousand Cores”. In: *PVLDB* 8.3 (2014), pp. 209–220.
- [23] Rebecca Taft, Essam Mansour, Marco Serafini, Jennie Duggan, Aaron J. Elmore, Ashraf Aboulnaga, Andrew Pavlo, and Michael Stonebraker. “E-Store: Fine-Grained Elastic Partitioning for Distributed Transaction Processing”. In: *PVLDB* 8.3 (2014), pp. 245–256.
- [24] Justin DeBrabant, Joy Arulraj, Andrew Pavlo, Michael Stonebraker, Stanley B. Zdonik, and Subramanya Duloor. “A Prolegomenon on OLTP Database Systems for Non-Volatile Memory”. In: *ADMS @ VLDB*. 2014, pp. 57–63.
- [25] Djellel Eddine Difallah, Andrew Pavlo, Carlo Curino, and Philippe Cudré-Mauroux. “OLTP-Bench: An Extensible Testbed for Benchmarking Relational Databases”. In: *PVLDB* 7.4 (2013), pp. 277–288.
- [26] Justin DeBrabant, Andrew Pavlo, Stephen Tu, Michael Stonebraker, and Stanley B. Zdonik. “Anti-Caching: A New Approach to Database Management System Architecture”. In: *PVLDB* 6.14 (2013), pp. 1942–1953.
- [27] Carlo A. Curino, Djellel E. Difallah, Andrew Pavlo, and Philippe Cudre-Mauroux. “Benchmarking OLTP/Web Databases in the Cloud: The OLTP-bench Framework”. In: *CloudDB*. 2012, pp. 17–20.
- [28] Andrew Pavlo, Carlo Curino, and Stanley Zdonik. “Skew-Aware Automatic Database Partitioning in Shared-Nothing, Parallel OLTP Systems”. In: *SIGMOD*. 2012, pp. 61–72.
- [29] Andrew Pavlo, Evan P.C. Jones, and Stan Zdonik. “On Predictive Modeling for Optimizing Transaction Execution in Parallel OLTP Systems”. In: *Proc. VLDB Endow.* 5 (2 Oct. 2011), pp. 85–96.
- [30] Andrew Pavlo, Erik Paulson, Alexander Rasin, Daniel J. Abadi, David J. DeWitt, Samuel Madden, and Michael Stonebraker. “A Comparison of Approaches to Large-Scale Data Analysis”. In: *Proceedings of SIGMOD*. 2009, pp. 165–178.
- [31] Ewa Deelman, Miron Livny, Gaurang Mehta, Andy Pavlo, Gurmeet Singh, Mei-Hui Su, Karan Vahi, and R. Kent Wenger. “Pegasus and DAGMan From Concept to Execution: Mapping Scientific Workflows onto Today’s Cyberinfrastructure”. In: *HPC and Grids in Action*. Vol. 16. Advances in Parallel Computing. 2008, pp. 56–74.
- [32] Andrew Pavlo, Peter Couvares, Rebekah Gietzel, Anatoly Karp, Ian D. Alderman, Miron Livny, and Charles Bacon. “The NMI Build & Test Laboratory: Continuous Integration Framework for Distributed Computing Software”. In: *LISA*. 2006, pp. 263–273.

Book Chapters

- [1] Franz Faerber, Alfons Kemper, Per-Åke Larson, Justin J. Levandoski, Thomas Neumann, and Andrew Pavlo. “Main Memory Database Systems”. In: *Foundations and Trends in Databases* 8.1-2 (2017), pp. 1–130.

Demonstrations / Tutorials

- [1] Bohan Zhang, Dana Van Aken, Justin Wang, Tao Dai, Shuli Jiang, Jacky Lao, Siyuan Sheng, Andrew Pavlo, and Geoffrey J. Gordon. “A Demonstration of the OtterTune Automatic Database Management System Tuning Service”. In: *PVLDB* 11.12 (2018), pp. 1910–1913.
- [2] Joy Arulraj and Andrew Pavlo. “How to Build a Non-Volatile Memory Database Management System”. In: *Proceedings of the 2017 ACM International Conference on Management of Data*. SIGMOD ’17. 2017, pp. 1753–1758.
- [3] Dana Van Aken, Djellel Eddine Difallah, Andrew Pavlo, Carlo Curino, and Philippe Cudre-Mauroux. “BenchPress: Dynamic Workload Control in the OLTP-Bench Testbed”. In: *Proceedings of SIGMOD*. 2015, pp. 1069–1073.
- [4] Ugur Çetintemel, Jiang Du, Tim Kraska, Samuel Madden, David Maier, John Meehan, Andrew Pavlo, Michael Stonebraker, Erik Sutherland, Nesime Tatbul, Kristin Tufte, Hao Wang, and Stanley B. Zdonik. “S-Store: A Streaming NewSQL System for Big Velocity Applications”. In: *PVLDB* 7.13 (2014), pp. 1633–1636.

- [5] Robert Kallman, Hideaki Kimura, Jonathan Natkins, Andrew Pavlo, Alexander Rasin, Stanley Zdonik, Evan P. C. Jones, Samuel Madden, Michael Stonebraker, Yang Zhang, John Hugg, and Daniel J. Abadi. “H-Store: A High-Performance, Distributed Main Memory Transaction Processing System”. In: *Proc. VLDB Endow.* 1.2 (2008), pp. 1496–1499.
- [6] Christopher Homan, Andrew Pavlo, and Jonathan Schull. “Smoother Transitions Between Breadth-First-Spanning-Tree-Based Drawings”. In: *14th International Symposium on Graph Drawing*. Vol. 4372. Lecture Notes in Computer Science. 2006, pp. 442–445.

Miscellaneous Writings

- [1] Dana Van Aken, Andrew Pavlo, and Geoffrey J. Gordon. *Amazon AWS AI Blog – Tuning Your DBMS Automatically with Machine Learning*. <https://aws.amazon.com/blogs/ai/tuning-your-dbms-automatically-with-machine-learning/>. June 2017.
- [2] Andrew Pavlo. “What Are We Doing With Our Lives?: Nobody Cares About Our Concurrency Control Research”. In: *Proceedings of the 2017 ACM International Conference on Management of Data*. SIGMOD ’17. 2017, p. 3.
- [3] Peter Bailis, Camille Fournier, Joy Arulraj, and Andrew Pavlo. “Research for Practice: Distributed Consensus and Implications of NVM on Database Management Systems”. In: *Commun. ACM* 59.11 (Oct. 2016), pp. 52–55. URL: <https://db.cs.cmu.edu/papers/2016/cacm-nvm2016.pdf>.
- [4] Andrew Pavlo and Matthew Aslett. “What’s Really New with NewSQL?” In: *SIGMOD Rec.* 45.2 (June 2016), pp. 45–55. URL: <https://db.cs.cmu.edu/papers/2016/pavlo-newsql-sigmodrec2016.pdf>.
- [5] Andrew Pavlo. “Emerging Hardware Trends in Large-Scale Transaction Processing”. In: *IEEE Internet Computing* 19.3 (May 2015), pp. 68–71. URL: http://www.cs.cmu.edu/~pavlo/static/papers/IC_20150501_May_2015.pdf.
- [6] Michael Stonebraker, Andrew Pavlo, Rebecca Taft, and Michael L. Brodie. “Enterprise Database Applications and the Cloud: A Difficult Road Ahead”. In: *2014 IEEE International Conference on Cloud Engineering*. 2014, pp. 1–6.
- [7] Andrew Pavlo. “Not Your Traditional Data Management - HPTS Conference Report”. In: *USENIX ;login:* (Feb. 2012), pp. 76–78. URL: <https://www.usenix.org/publications/login/february-2012/14th-international-workshop-high-performance%20-transaction-systems>.
- [8] Xin Jia, Andrew Pavlo, and Stanley B. Zdonik. “Tastes Great, Less Filling: Low-Impact OLAP MapReduce Queries on High-Performance OLTP Systems”. In: *TinyToCS* 1 (2012). URL: <http://tinytocs.org/vol1/papers/tinytocs-v1-jia.pdf>.
- [9] Andrew Pavlo and Ning Shi. “Graffiti Networks: A Subversive, Internet-Scale File Sharing Model”. In: *CoRR* abs/1101.0350 (2011). URL: <http://arxiv.org/abs/1101.0350>.
- [10] Michael Stonebraker, Daniel Abadi, David J. DeWitt, Sam Madden, Erik Paulson, Andrew Pavlo, and Alexander Rasin. “MapReduce and Parallel DBMSs: Friends or Foes?” In: *Communications of the ACM* 53.1 (2010), pp. 64–71.
- [11] Andrew Pavlo, Christopher Homan, and Jonathan Schull. “A parent-centered radial layout algorithm for interactive graph visualization and animation”. In: *CoRR* abs/cs/0606007 (2006). URL: <http://arxiv.org/abs/cs/0606007>.

Synergistic Activities

Open-Source Software:

- Peloton Database Management System (<https://pelotondb.io>)
- OtterTune (<https://ottertune.cs.cmu.edu>)
- Carnegie Mellon Database Application Catalog (<http://cmdbac.cs.cmu.edu>)
- H-Store Database Management System (<http://hstore.cs.brown.edu>)

- OLTP-Bench Database Benchmark Suite (<https://github.com/oltpbenchmark>)

Database Seminar Series:

- Fall 2018: Hardware-Accelerated Database Lectures (<https://db.cs.cmu.edu/seminar2018>)
- Fall 2017: Time Series Database Lectures (<https://db.cs.cmu.edu/seminar2017>)
- Fall 2015: Databaseology Lectures (<https://db.cs.cmu.edu/seminar2015>)
- Fall 2014: Seven Databases in Seven Weeks (<https://db.cs.cmu.edu/seminar2014>)

Internship Programs:

- Summer 2019 (<https://db.cs.cmu.edu/internship2019>)
- Summer 2018 (<https://db.cs.cmu.edu/internship2018>)
- Summer 2017 (<https://db.cs.cmu.edu/internship2017>)
- Summer 2017 Pennsylvania Governor's School for the Sciences (Pre-College)

Talks

- “Make Your Database Dream of Electric Sheep: Designing for Autonomous Operation,” *Amazon AWS / Percona Live / Microsoft Research Faculty Summit*, March 2018.
- “What Non-Volatile Memory Means for the Future of Database Systems,” *Two Sigma / MongoDB*, August 2017.
- “What Are We Doing With Our Lives?: Nobody Cares About Our Concurrency Control Research,” *SIGMOD Keynote*, May 2017.
- “How to Build a Non-Volatile Memory Database Management System,” *SIGMOD Tutorial*, May 2017.
- “Self-Driving Database Management System,” *CIDR / Google / DataBricks*, January 2017.
- “Peloton: The Self-Driving Database Management System,” *DataEngConf (NYC)*, November 2016.
- “Peloton: The Self-Driving Database Management System,” *PDL Visit Day*, May 2016.
- “What Non-Volatile Memory Means for the Future of Database Systems,” *In-Memory Computing Summit*, May 2016.
- “Andy Pavlo’s Mostly Professional and Collected Presentation about Getting Down & Dirty in OLTP Databases with Intel’s NVM SDV,” *ISTC for Big Data Retreat*, September 2015.
- “Reducing Replication Bandwidth for Distributed Document Databases,” *MongoDB*, July 2015.
- “Moving the Abyss: Database Management on Future 1000-core Processors,” *NSF Workshop on Exploiting Parallelism and Scalability*, June 2015.
- “I Don’t Want to be the Mitt Romney of Databases,” *University of Michigan / Two Sigma / IBM Research / University of Waterloo*, April–November 2015.
- “Database Systems Do Not Scale to 1000 CPU Cores and Other Tales of the Macabre,” *CMU Qatar / Qatar Computing Research Institute / New England Database Summit*, December 2014.
- “Let’s Talk About Storage & Recovery Methods for Non-Volatile Memory OLTP Database Systems,” *ISTC for Big Data Retreat*, August 2014.
- “H-Store Retrospective,” *SIGMOD Jim Gray Doctoral Dissertation Award*, June 2014.
- “Cache Rules Everything Around Me,” *Univ. of California, Santa Cruz / Univ. of California, Berkeley / Stanford University / University of Pittsburgh*, January 2014
- “OLTP on the NVM SDV: YMMV,” *ISTC for Big Data Retreat*, December 2013.
- “What’s Really New With NewSQL,” *Northwest Database Society / Google Kirkland*, October 2013.
- “Random Prognostications: An Evening with Andy Pavlo,” *HPTS 2013*, September 2013.
- “On Predictive Modeling for Distributed Databases,” *VLDB*, August 2012.
- “Automatic Database Partitioning in Parallel OLTP Systems,” *SIGMOD*, May 2012.

- “Making Fast Databases Faster,” *Columbia University / Yale University*, April 2012.
- “Methods for the Optimization of Parallel OLTP Systems,” *VoltDB, Inc.*, April 2012.
- “Making Fast Databases Faster,” *New England Database Summit 2012*, February 2012.
- “Magical Parallel OLTP Databases,” *Massachusetts Institute of Technology*, November 2011.
- “Life After the Stonebraker Stores,” *HPTS*, October 2011.
- “MapReduce and Parallel DBMSs: Together At Last,” *Brown IPP Symposium on Cloud Computing*, May 2010.
- “MapReduce and Parallel DBMSs: Together At Last,” *New England Database Summit 2010*, January 2010.
- “H-Store: A Specialized Architecture for High-throughput OLTP Applications,” *International Workshop on High Performance Transaction Systems*, October 2009.
- “Graffiti Networks: A Subversive, Internet-Scale File Sharing Model,” *DC401 - Rhode Island Defcon Group*, October 2009.
- “MapReduce and Parallel DBMSs: A Comparison of Approaches to Large-Scale Data Analysis,” *University of Maryland College Park*, September 2009.
- “Challenges in Dynamic Deployment of Condor Across Distributed Environments,” *University of Wisconsin–Madison*, May 2007.
- “NMI Build & Test Laboratory: Continuous Integration Framework for Distributed Computing Software,” *USENIX LISA*, December 2006.

Panels

- “Junior Faculty Panel,” *CMU/MIT Rising Stars Workshop*, November 2016.
- “Computing Futures,” *CMU CSD50 Celebration*, October 2015.
- “Can Big Data Platforms EVER Deliver Real-time Analytics?,” *BIRTE @ VLDB*, August 2015.
- “Mobility and Big Data: Adding an M to the Vs, or not?,” *IEEE International Conference on Mobile Data Management*, June 2015.