

Curriculum Vitae

Peter Alvaro

September 2018

Computer Science Department
University of California, Santa Cruz
1156 High Street, MS SOE3
Santa Cruz, CA 95064

PHONE: +1 (415) 813 9364
EMAIL: palvaro@ucsc.edu
<http://people.ucsc.edu/~palvaro>

EMPLOYMENT HISTORY

2015– Assistant Professor, Computer Science Department, University of California, Santa Cruz
2018 Consultant, eBay Inc., San Jose, CA
2016 Consultant, Uber Inc., San Francisco, CA
2015 Consultant, Netflix Inc., Los Gatos, CA
2008-2014 Research Assistant, Computer Science Division, University of California at Berkeley
1999-2008 Senior Software Engineer, Ask.com, Oakland, CA

EDUCATION

2015 Ph. D., University of California at Berkeley, Computer Science (advisor: Joseph M. Hellerstein)
Thesis: *Data-centric Programming for Distributed Systems*
2010 M. S., University of California at Berkeley, Computer Science
1997 B. A., Middlebury College, English Literature, *magna cum laude*

HONORS AND AWARDS

2017 **NSF CAREER** Award: “Lineage-driven Fault Injection.”
2017 **Facebook Research Award**
2017 **Keynote Speaker**, O’Reilly Velocity, San Jose, CA, 2017.
2016 **Keynote Speaker**, Reactive Summit, Austin, TX, 2016.
2016 **Keynote Speaker**, Qcon London 2016.
2015 **Keynote Speaker**, Strange Loop, Saint Louis, MO, September 2015.
2014 **Keynote Speaker**, Ricon, Las Vegas, NV, October 2014.
2010 **Graduate Research Fellow**, National Science Foundation
2009 **Finalist**, Qualcomm Innovation Fellowship
1997 **Phi Beta Kappa**, Middlebury College
1997 **Highest Honors** in English Literature, Middlebury College
1997 **Winner**, Reid L. Carr prize, Middlebury College

RESEARCH FUNDING

Grants

2018–2022 PI, *CSR: Medium: Declarative Programmable Storage*, National Science Foundation, \$425,000.
2017–2021 PI, *CAREER: Lineage-driven Fault Injection*, National Science Foundation, \$475,000.

Gifts

2018 eBay, \$125,000.
2017 Facebook, \$50,000.
2017 Huawei, Inc., \$100,000.

SCHOLARLY AND CREATIVE WORK

NOTE: (*) denotes a student co-author, and (‡) denotes a student co-author who was one of my advisees.

Journals

- J6. Haryadi S. Gunawi, Riza O. Suminto, Russell Sears, Casey Gollhofer, Swaminathan Sundararaman, Xing Lin, Tim Emami, Weiguang Sheng, Nematollah Bidokhti, Caitie McCaffrey, Gary Grider, Parks M. Fields, Kevin Harms, Robert B. Ross, Andree Jacobson, Robert Ricci, Kirk Webb, **Peter Alvaro**, Mingzhe Hao, Huaicheng Li, H. Biral Runesha, “Fail-Slow at Scale: Evidence of Hardware Performance Faults in Large Production Systems”, *ACM Transactions on Storage (TOS, to appear)*, November 2018.
- J5. **Peter Alvaro**, Severine Tymon, “Abstracting the Geniuses Away from Failure Testing”, *Communications of the ACM (CACM)*, Volume 61, Issue 1, November 2017.
- J4. **Peter Alvaro**, Neil Conway, Joseph M. Hellerstein, David Maier, “Coordination Analysis and Placement for Distributed Programs”, *Transactions on Database Systems (TODS)*, Volume 42, Issue 4, October 2017 (ACM).
- J3. Tom J. Ameloot, Jan Van den Bussche, William R. Marczak, **Peter Alvaro**, and Joseph M. Hellerstein, “Putting logic-based distributed systems on stable grounds”, *Theory and Practice of Logic Programming*, **16**(2), August 2015.
- J2. Neil Conway, **Peter Alvaro**, Emily Andrews, and Joseph M Hellerstein, “Edelweiss: Automatic Storage Reclamation for Distributed Programming”, *Proceedings of the VLDB Endowment (PVLDB 2014)*, **7**(6), February, 2014.
- J1. **Peter Alvaro**, Tyson Condie, Neil Conway, Joseph M Hellerstein, and Russell Sears, “I do declare: consensus in a logic language”, *ACM SIGOPS Operating Systems Review*, **43**(4), January 2010.

Refereed Conference & Workshop Papers

- C22. Michael Whittaker*, Cristina Teodoropol*, **Peter Alvaro**, Joseph M. Hellerstein, “Debugging Distributed Systems with Why-Across-Time Provenance”, *Proceedings of the Ninth ACM Symposium on Cloud Computing (SoCC 2018)*, Carlsbad, CA (ACM), 2018.
- C21. Daniel Bittman*, Matthew Gray*, Justin Raizes*, Sinjoni Mukhopadhyay*, Matt Bryson*, **Peter Alvaro**, Darrell Long and Ethan Miller, “Designing Data Structures to Minimize Bit Flips on NVM”, *Proceedings of the 7th IEEE Non-Volatile Memory Systems and Applications Symposium (NVMSA 2018)*, Hakodate, Japan (IEEE), 2018.
- C20. Michael A. Sevilla*, Reza Nasirigerdeh*, Carlos Maltzahn, Jeff LeFevre, Noah Watkins*, **Peter Alvaro**, Margaret Lawson, Jay Lofstead, Jim Pivarskia, “Tintenfisch: File System Namespace Schemas and Generators”, *Proceedings of the 10th USENIX Workshop on Hot Topics in Storage and File Systems (HotStorage 2018)*, Boston, MA (USENIX), 2018.
- C19. Michael Sevilla*, Carlos Maltzahn, **Peter Alvaro**, Reza Nasirigerdeh*, Bradley Settlemyer, Danny Perez, David Rich and Galen Shipman, “Programmable Caches with a Data Management Language & Policy Engine”, *Proceedings of the 18th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGrid 2018)*, Washington, DC (IEEE/ACM), 2018.
- C18. Michael Sevilla*, Ivo Jimenez*, Noah Watkins*, Jeff Lefevre, Shel Finkelstein, **Peter Alvaro**, Patrick Donnelly and Carlos Maltzahn, “Cudele: An API and Framework for Programmable Consistency and Durability in a Global Namespace”, *Proceedings of the 32nd IEEE International Parallel and Distributed Processing Symposium (IPDPS 2018)*, Vancouver, BC, Canada (IEEE), May 2018.
- C17. Haryadi S. Gunawi, Riza O. Suminto, Russell Sears, Casey Gollhofer, Swaminathan Sundararaman, Xing Lin, Tim Emami, Weiguang Sheng, Nematollah Bidokhti, Caitie McCaffrey, Gary Grider, Parks M. Fields, Kevin Harms, Robert B. Ross, Andree Jacobson, Robert Ricci, Kirk Webb, **Peter Alvaro**, Mingzhe Hao, Huaicheng Li, H. Biral Runesha, “Fail-Slow at Scale: Evidence of Hardware Performance Faults in Large Production Systems”, *Proceedings of the 16th USENIX Conference on File and Storage Technologies (FAST 2018)*, Oakland, CA (USENIX), February 2018.
- C16. Kamala Ramasubramanian‡, Kathryn Dahlgren‡, Asha Karim‡, Sanjana Maiya*, Sarah Borland*, **Peter Alvaro**, “Growing a Protocol”, *Proceedings of the 9th USENIX Workshop on Hot Topics in Cloud Computing (HotCloud 2017)*, Santa Clara, CA (USENIX), July 2017.

- C15. Noah Watkins*, Michael A. Sevilla*, Ivo Jimenez*, Kathryn Dahlgren[‡], **Peter Alvaro**, Shel Finkelstein, Carlos Maltzahn, “DeclStore: Layering is for the Faint of Heart”, *Proceedings of the 9th USENIX Workshop on Hot Topics in Storage and File Systems (HotStorage 2017)*, Santa Clara, CA (USENIX), July 2017.
- C14. Michael A. Sevilla*, Noah Watkins*, Ivo Jimenez*, **Peter Alvaro**, Shel Finkelstein, Jeff LeFevre, Carlos Maltzahn, “Malacology: A Programmable Storage System”, *Proceedings of the 12th European conference on Computer systems (Eurosys 2017)*, Belgrade, Serbia, April, 2017.
- C13. **Peter Alvaro**, Kolton Andrus, Chris Sanden, Casey Rosenthal, Ali Basiri, Lorin Hochstein, “Automating Failure Testing Research at Internet Scale”, *Proceedings of the 7th Annual Symposium on Cloud Computing (SoCC 2016)*, Santa Clara, CA: ACM, October, 2016.
- C12. **Peter Alvaro**, Joshua Rosen, and Joseph M Hellerstein, “Lineage-driven fault injection”, *Proceedings of the 2015 ACM SIGMOD International Conference on Management of Data (SIGMOD 2015)*, Melbourne, Victoria, Australia: ACM, May 2015.
- C11. **Peter Alvaro**, Neil Conway, Joseph M. Hellerstein, and David Maier, “Blazes: Coordination analysis for distributed programs”, *Proceedings of the 30th International Conference on Data Engineering (ICDE 2014)*, Chicago, IL: IEEE, April 2014.
- C10. **Peter Alvaro**, Peter Bailis, Neil Conway, and Joseph M Hellerstein, “Consistency without borders”, *Proceedings of the 4th Annual Symposium on Cloud Computing (SoCC 2013)*, Santa Clara, CA: ACM, October, 2013.
- C9. Neil Conway, William R Marczak, **Peter Alvaro**, Joseph M Hellerstein, and David Maier, “Logic and lattices for distributed programming”, *Proceedings of the 3rd Annual Symposium on Cloud Computing (SoCC 2012)*, San Jose, CA: ACM, October, 2012.
- C8. **Peter Alvaro**, Tom J. Ameloot, Joseph M. Hellerstein, William R Marczak, and Jan Van den Bussche, “A declarative semantics for Dedalus”, *Proceedings of the Second international conference on Datalog in Academia and Industry (Datalog2.0 2012)*, Vienna, Austria, September 2012.
- C7. William R Marczak, **Peter Alvaro**, Neil Conway, Joseph M Hellerstein, and David Maier, “Confluence analysis for distributed programs: A model-theoretic approach”, *Proceedings of the Second international conference on Datalog in Academia and Industry (Datalog2.0 2012)*, Vienna, Austria, September 2012.
- C6. **Peter Alvaro**, Andrew Hutchinson, Neil Conway, William R Marczak, and Joseph M Hellerstein, “BloomUnit: Declarative testing for distributed programs”, *Proceedings of the Fifth International Workshop on Testing Database Systems (DBTest 2012)*, Scottsdale, AZ: ACM, May 2012.
- C5. **Peter Alvaro**, Neil Conway, Joe Hellerstein, and William R Marczak, “Consistency Analysis in Bloom: a CALM and Collected Approach”, *Proceedings of the 5th Biennial Conference on Innovative Data Systems Research (CIDR 2011)*, Asilomar, CA, January, 2011.
- C4. Haryadi S Gunawi, Thanh Do, Pallavi Joshi, **Peter Alvaro**, Joseph M Hellerstein, Andrea C Arpaci-Dusseau, Remzi H Arpaci-Dusseau, Koushik Sen, and Dhruva Borthakur, “FATE and DESTINI: a framework for cloud recovery testing”, *Proceedings of the 8th USENIX Symposium on Networked Systems Design and Implementation (NSDI 2011)*, Boston, MA: USENIX, April 2011.
- C3. **Peter Alvaro**, Tyson Condie, Neil Conway, Khaled Elmeleegy, Joseph M Hellerstein, and Russell Sears, “Boom analytics: exploring data-centric, declarative programming for the cloud”, *Proceedings of the 5th European conference on Computer systems (Eurosys 2010)*,
- C2. Tyson Condie, Neil Conway, **Peter Alvaro**, Joseph M Hellerstein, Khaled Elmeleegy, and Russell Sears, “MapReduce online”, *Proceedings of the 7th USENIX Symposium on Networked Systems Design and Implementation (NSDI 2010)*, San Jose, CA: USENIX, April 2010.
- C1. **Peter Alvaro**, William R Marczak, Neil Conway, Joseph M Hellerstein, David Maier, and Russell Sears, “Dedalus: Datalog in time and space”, *Proceedings of the First international Conference on Datalog Reloaded (Datalog 2010)*, Oxford, England, September 2010.

Technical Reports

- C1. Noah Watkins, Michael Sevilla, Ivo Jimenez, Neha Ojha, **Peter Alvaro**, Carlos Maltzahn, “Bradoss: Declarative, Programmable Object Storage”, Technical Report UCSC-SOE-16-12, December 2016.

- 2016 **Keynote:** “Orchestrated Chaos”, Reactive Summit 2016.
- 2016 **Keynote:** “Monkeys in Lab Coats: Applying Failure Testing Research @Netflix”, QCon London 2016.
- 2016 “Cause, that’s why”, Papers We Love, San Francisco, CA, December 2016.
- 2016 “Applying Database Research in the Wild”, Stanford University DAWN Seminar, Stanford, CA, November 2016.
- 2016 “Applying Database Research in the Wild”, UC Berkeley Database Seminar, Berkeley, CA, November 2016.
- 2016 “Applying Failure Testing Research @Netflix”, QCon San Francisco, San Francisco, CA, November 2016.
- 2016 “Orchestrated Chaos”, Chaos Community Day, Amazon, Seattle, WA, August 2016.
- 2016 “Lineage-driven Fault Injection”, VMWare Research Group Chat, VMWare, Palo Alto, CA, August 2016.
- 2016 “Databases are dead: long live databases!”, Norcal DB Day, Google, Inc. Mountain View, CA May, 2016.
- 2015 **Keynote:** “I see what you mean”, Strange Loop 2015, Saint Louis, MO, September, 2015.
- “Lineage-driven fault injection”, ACM SIGMOD 2015, Melbourne, Victoria, Australia, June, 2015.
- “Lineage-driven fault injection”, Pivotal, Palo Alto, CA, November, 2015.
- “Lineage-driven fault injection”, Arista Networks, San Francisco, CA, November, 2015.
- 2014 **Keynote:** “Outwards from the middle of the maze”, RICON 2014, Las Vegas, NV.
- “Blazes: coordination analysis for distributed programs”, ICDE 2014, Chicago, IL, April, 2014.
- “Ineluctable modality of the distributed”, Papers We Love SF, San Francisco, CA, August, 2014.
- 2013 “Consistency without borders”, ACM SoCC 2013, Santa Clara, CA, October, 2013.
- “Bloom and CALM: Programming the Cloud”, Norcal DB Day 2013, Stanford, CA, April, 2013.
- “Disorderly Distributed Programming with Dedalus and Bloom”, UCLA Programming Language Seminar, March, 2013.
- Panelist, *Think Distributed* Panel, Ricon 2012, San Francisco, CA, December, 2013.
- 2012 “BloomUnit: Declarative testing for distributed programs”, ACM DBTest 2012, Scottsdale, AZ, May 2012.
- “Bloom: disorderly programming for a distributed world”, lang.next, Microsoft, Inc., Redmond, WA, April, 2012.
- 2011 “Bloom: disorderly programming. CALM analysis”, Ask.com, Oakland, CA, March, 2011.
- “Bloom: CALMly building skyscrapers on quicksand”, Twitter, Inc., San Francisco, CA, January 2011.
- 2010 “Show and Tell: Building a consistent, replicated shopping cart in bloom”, TDS Research Group, MIT, Cambridge, MA, November 2010.
- “I Do Declare: Consensus in a Logic Language”, OSQ Retreat, Santa Cruz, CA, May 2010.
- 2009 “I Do Declare”, NetDB 2009, Big Sky, MT, October 2009.
- “BOOM: Data-Centric Programming For The Data Center”, Infolunch, Stanford University, Stanford, CA, April 2009.

UNIVERSITY SERVICE

Departmental and School Service

- 2017–18 Chair, Computer Science Dept. Faculty Search Subcommittee (Distributed Systems)
- 2017–18 Computer Science Dept. Faculty Search Committee.
- 2017 Curriculum development: *Computer Systems Fundamentals*
- 2016–17 Computer Science Dept. Faculty Search Committee.
- 2015–16 Computer Science Dept. Graduate Admission Committee.

MENTORING AND STUDENT ADVISING**Doctoral Students**

Dates	Relationship	Degree Year	Name and Activities
2015 Fall - Present	Primary Supervisor		Kamala Ramasubramanian
2016 Fall - Present	Primary Supervisor		Tuan Tran
2016 Fall - Present	Primary Supervisor		Kathryn Dahlgren
2018 Fall	Other Advisor	2018	Noah Watkins Member: Dissertation Committee
2018 Spring	Other Advisor	2018	Michael Sevilla Member: Dissertation Committee
2016 Fall	Other Advisor	2016	Richard Halpert Member: Dissertation Committee
2016 Fall	Other Advisor	2016	Kun Qian Member: Dissertation Committee
2016 Spring	Other Advisor	2016	Dustin Rhodes Member: Qualifying Examination Committee
2016 Spring	Other Advisor	2016	Christine Strong Member: Dissertation Committee

Masters Students

Dates	Relationship	Degree Year	Name and Activities
2018 Spring	Other Advisor		Austen Barker Member: Masters Project Committee
2018 Spring	Other Advisor		James Byron Member: Masters Thesis Committee
2017 Fall - 2018 Spring	Primary Supervisor	2018	Lennart Oldenburg
2018 Spring	Other Advisor		Devashish Purendare Member: Masters Project Committee
2018 Spring	Other Advisor		Kenneth Chang Member: Masters Project Committee
2018 Spring	Other Advisor		Oceane Bell Member: Masters Project Committee
2018 Winter	Other Advisor	2018	Umang Sardesai Chair: Masters Project Committee
2016 Fall - 2017 Fall	Primary Supervisor	2017	Ashutosh Raina
2017 Fall	Other Advisor		Zheyuan Chen Member: Masters Project Committee
2017 Fall	Other Advisor	2017	Haiyu Yang Member: Masters Project Committee
2017 Fall	Other Advisor	2017	Bettie Jea Member: Masters Project Committee
2017 Fall	Other Advisor	2018	Pinglei Guo Chair: Masters Project Committee
2017 Fall	Other Advisor	2017	Nikhil Kini Member: Masters Project Committee
2017 Winter	Other Advisor	2016	Neha Ojha Member: Masters Project Committee
2017 Winter	Other Advisor	2016	Greeshma Swaminathan Member: Masters Project Committee
2016 Fall	Masters Project Supervisor	2016	Sanjana Maiya Chair: Masters Project Committee
2017 Winter	Other Advisor	2017	Trivikram Bollempalli Member: Masters Project Committee
2016 Fall	Other Advisor	2017	Abishek Grover Member: Masters Thesis Committee
2016 Spring	Other Advisor		Kathryn Dahlgren Member: Masters Thesis Committee

Undergraduate Students

Dates	Relationship	Degree Year	Name and Activities
2017 Winter-2017 Fall	Primary Supervisor	2017	Asha Karim

COURSES TAUGHT

Courses in bold italics are those for which I either developed or significantly revised the curriculum.

Undergraduate

Winter 2018 CMPS 128: Distributed Systems
Fall 2017 CMPS 128: Distributed Systems
Winter 2017 CMPS 128: Distributed Systems
Spring 2016 ***CMPS 128: Distributed Systems***
Spring 2013 CS 194-017: Programming the Cloud (UC Berkeley)
Spring 2011 ***CS 194-017: Programming the Cloud*** (UC Berkeley)

Graduate

Spring 2018 CMPS 232: Graduate Distributed Systems
Fall 2016 ***CMPS 232: Graduate Distributed Systems***
Winter 2015 ***CMPS 290S: Advanced Topics in Computer Systems: Distributed Storage Systems and Programming Models***