

Curriculum Vitae

Jeff LeFevre

November 2017

Computer Science Department
University of California, Santa Cruz
1156 High Street, MS SOE2
Santa Cruz, CA 95064

PHONE: +1 (831) 502-7474
FAX: +1 (831) 459-4829
EMAIL: jlefevre@cs.ucsc.edu
<http://www.cs.ucsc.edu/~jlefevre/>

RESEARCH INTERESTS

I do research and development in the area of cloud-based data management. My academic research concerned database physical design, for which I developed new physical design methods for cloud-based architectures. My work also addressed “multistore” or hybrid architectures composed of multiple unique systems used together for workload co-processing. In industry, I continued this direction toward integrating the HP Vertica database with external systems such as Apache Spark; designing new methods for fast and reliable parallel data transfer. I currently lead the *SyhookDB* elastic database for the cloud project in the Center for Research in Open Source Software at UC Santa Cruz. *SyhookDB* leverages the flexibility and extensibility of object storage, which is ubiquitous in the cloud, toward database physical design to improve processing and enhance scalability as well as support domain-specific processing. In a related project, I do research on domain-specific processing in object storage for high performance computing centers.

EMPLOYMENT HISTORY

2016– Assistant Project Scientist, Computer Science Department, University of California, Santa Cruz
2014–2016 Software Engineer, Advanced R&D Lab, Big Data Platform, Hewlett Packard Enterprise
2014 Associate in Computer Science, Graduate Student Instructor, University of California, Santa Cruz
2011–2013 Summer Internships, Researcher, Data Management Group, NEC Labs America
2007–2010 Summer Internships, Software Engineer, Platforms Group, Google
2005–2006 Summer Internships, Software Engineer, Virtual Storage Architecture Group, Teradata
2004–2005 Systems Analyst, World Wide Protein Data Bank, San Diego Supercomputer Center

EDUCATION

2014 Ph. D., Computer Science, University of California, Santa Cruz (advisor: Neoklis Polyzotis)
Thesis: *Physical Design Tuning Methods for Emerging System Architectures*
2009 M. S., Computer Science & Engineering, University of California, San Diego (advisor: Walt Burkhard)
Thesis: *Improving Disk Array Reliability and Performance*
2004 B. S., Computer Science & Engineering, University of South Florida

HONORS

2013 Student Travel Award, ACM SIGMOD.
2012 Student Travel Award, ACM SIGMOD.
2004 Student Engineer of the Year, IEEE Florida West Coast Section.
2003 Elected to Tau Beta Pi, Florida Gamma Chapter.

PUBLICATIONS

Journals

- J2. Hakan Hacigumus, Jagan Sankaranarayanan, Junichi Tatemura, **Jeff LeFevre**, Neoklis Polyzotis, “Odyssey: a multistore system for evolutionary analytics”, *Proceedings of the VLDB Endowment (VLDB '13)*, **6**(11), August 2013, pages 1180–1181.
- J1. Ivo Jimenez, **Jeff LeFevre**, Neoklis Polyzotis, Huascar Sanchez, Karl Schnaitter, “Benchmarking Online Index-Tuning Algorithms”, *IEEE Data Engineering Bulletin* **34**(4), December 2011, pages 28–35.

Refereed Conference & Workshop Papers

- C9. Michael Sevilla, Noah Watkins, Ivo Jimenez, Peter Alvaro, Shel Finkelstein, **Jeff LeFevre**, Carlos Maltzahn, “Malacology: A Programmable Storage System”, *Proceedings of the Twelfth European Conference on Computer Systems (EuroSys '17)*, April 2017.
- C8. **Jeff LeFevre**, Rui Liu, Cornelio Inigo, Malu Castellanos, Lupita Paz, Edward Ma, Meichun Hsu, “Building the Enterprise Fabric for Big Data with Vertica and Spark”, *Proceedings of the 2016 ACM SIGMOD International Conference on Management of Data (SIGMOD '16)*, June 2016.
- C7. Shreya Prasad, Arash Fard, Vishrut Gupta, Jorge Martinez, **Jeff LeFevre**, Vincent Xu, Meichun Hsu, Indrajit Roy, “Large-scale Predictive Analytics in Vertica: Fast Data Transfer, Distributed Model Creation, and In-database Prediction”, *Proceedings of the 2015 ACM SIGMOD International Conference on Management of Data (SIGMOD '15)*, June 2015.
- C6. **Jeff LeFevre**, Jagan Sankaranarayanan, Hakan Hacigumus, Junichi Tatemura, Neoklis Polyzotis, Michael J. Carey, “MISO: Souping Up Big Data Query Processing with a Multistore System”, *Proceedings of the 2014 ACM SIGMOD International Conference on Management of Data (SIGMOD '14)*, June 2014.
- C5. **Jeff LeFevre**, Jagan Sankaranarayanan, Hakan Hacigumus, Junichi Tatemura, Neoklis Polyzotis, Michael J. Carey, “Opportunistic Physical Design for Big Data Analytics”, *Proceedings of the 2014 ACM SIGMOD International Conference on Management of Data (SIGMOD '14)*, June 2014.
- C4. **Jeff LeFevre**, Jagan Sankaranarayanan, Hakan Hacigumus, Junichi Tatemura, Neoklis Polyzotis, “Towards a Workload for Evolutionary Analytics”, *Proceedings of the Second Workshop on Data Analytics in the Cloud (ACM SIGMOD)*, June 2013.
- C3. Mariano P. Consens, Kleoni Ioannidou, **Jeff LeFevre**, Neoklis Polyzotis, “Divergent Physical Design Tuning for Replicated Databases”, *Proceedings of the 2012 ACM SIGMOD International Conference on Management of Data (SIGMOD '12)*, June 2012.
- C2. Joe B. Buck, Noah Watkins, **Jeff LeFevre**, Kleoni Ioannidou, Carlos Maltzahn, Neoklis Polyzotis, Scott Brandt, “SciHadoop: Array-based Query Processing in Hadoop”, *Proceedings of 2011 International Conference for High Performance Computing, Networking, Storage and Analysis (SC '11)*, November 2011.
- C1. David Kephart, **Jeff LeFevre**, “CodeGen: The Generation and Testing of DNA Code Words”, *Proceedings of the 2004 IEEE Congress on Evolutionary Computation*, June 2004.

PROFESSIONAL ACTIVITIES

Conference Organization

- 2018 **Program Committee:** 19th IEEE International Conference on Mobile Data Management (MDM 2018).
- 2016 **Program Committee:** 32nd International Conference on Data Engineering (ICDE 2016).

Reviewer of Technical Papers and Proposals

- 2017 *International Journal of Distributed and Parallel Databases (DAPD).*
- 2017 *IEEE Transactions on Knowledge and Data Engineering (TKDE).*
- 2016 *IEEE Transactions on Knowledge and Data Engineering (TKDE).*
- 2015 *IEEE Transactions on Knowledge and Data Engineering (TKDE).*
- 2015 *Proceedings of the VLDB Endowment.*

Membership in Professional Associations

ACM SIGMOD.
Tau Beta Pi.

Invited Talks

- 2017 “Leveraging object storage toward database elasticity for the cloud”, Workshop on Data Organisation, Management, and Access; The Flatiron Institute, New York, NY, November 2017.
- 2017 “SkyhookDB: leveraging object storage toward database elasticity for the cloud”, PostgresOpen 2017 Lightning Talk; San Francisco, CA, September 2017.
- 2016 “Building the Enterprise Fabric for Big Data with Vertica and Apache Spark”, UC Merced EECS Seminar; Merced, CA, September 2016.

COMMUNITY SERVICE

2015–2017 Volunteer, Santa Cruz Homeless Garden Project

COURSES TAUGHT**Undergraduate**

Winter 2014 CMPS 181, Database Systems II

OTHER TEACHING**Undergraduate**

Winter 2016 Guest Lecturer, CMPS 181, Database Systems II

Fall 2008 Teaching Assistant, CMPS 111, Operating Systems

Graduate

Spring 2017 Guest Lecturer (4 sessions), CMPS 229, Storage Systems