Jennie Duggan (née Rogers)

MIT CSAIL
32 Vasser Street, Room G904B
Cambridge, MA 02139

Voice: omitted for web
E-mail: jennie@csail.mit.edu
Web: http://people.csail.mit.edu/jennie/

Interests

Large-scale data processing, cloud computing, core database internals, scientific data management, database performance, big data

Education

Brown University December 2012
Ph.D., Computer Science
- Thesis: "Query Performance Prediction for Analytical Workloads"
- Advisor: Uğur Çetintemel

Brown University May 2009
Sci.M., Computer Science
- Thesis: "Towards a Generic Compression Advisor"
- Advisor: Uğur Çetintemel

Rensselaer Polytechnic Institute May 2003
B.S., Computer Science
- Minor: Brain & Brain Behavior

Research Experience

MIT CSAIL January 2013–Present
Postdoctoral Associate

Data placement for elastic array processing (SciDB):
- Researched the efficient reorganization of data for scalable array analytics
- Designed a control loop for incrementally scaling out a scientific database cluster
- Proposed partitioning schemes to maximize the preservation of array space for spatial querying

Data placement for elastic distributed main memory databases (H-Store):
- Investigated the management of execution skew for large-scale transactional workloads
- Created algorithms for determining when and how to reprovision hardware resources for consistent query throughput
- Built a two-level caching scheme that caters to hot spots in a transactional workload

Massively parallelized join optimization (SciDB):
- Proposed a two-phase join for arrays independently optimizing data alignment and cell matching in a shared-nothing architecture
- Improved performance by sending sparse array chunks to their dense counterparts
- Used integer programming to assign join segments to nodes to minimize network bandwidth usage and join execution skew

Brown University 2006–2012
Research Assistant

Performance optimization of static query workloads executing under concurrency:
- Predicted individual query duration when executing as a part of an arbitrary concurrent mix
- Introduced a novel metric for logical I/O which approximates the effects of resource sharing on analytical query execution time

Query performance prediction for ad-hoc analytical workloads:
- Created prediction framework to capture both the performance characteristics of a query under prediction and the shared resources to which it will have access
- Applied machine learning to build individual query models based on their execution plan

Provisioning of cloud databases:
- Researched the matching of Infrastructure-as-a-Service resources to a query workload
- Evaluated the utility of statistics at varying levels of granularity for cloud purchasing decisions

Compression advisor for array databases:
- Designed and implemented a system for selecting Pareto optimal compression configuration for scalable array processing
- Optimized for storage size, encoding, and decoding time to improve total database performance

Rensselaer Polytechnic Institute 2002–2003
Research Assistant

Project Links:
- Assisted in NSF-funded study dedicated to linking the concepts of higher math to real-world applications through interactive web modules
- Developed online education tools

Industrial Experience

Qatar Computing Research Institute October–November 2013
Visiting Scholar
- Implemented data partitioners for the H-Store main memory database
- Worked with full-time scientists on research agenda

NEC Labs of America Summer 2012
Research Intern
- Created models for predicting query workload performance for databases executing on a dynamic hardware platform

Paradigm 4, Inc. Summer 2011
Research Intern
- Designed techniques for automatically selecting compression algorithms in an array database
- Investigated vectorized execution model for array data management to expedite query processing and reduce disk contention

Naval Undersea Warfare Center 2003–2006
Scientist
- Researched path planning algorithms for unmanned surface vessels and evaluated prototypes
- Designed and implemented a distributed networked architecture for battlespace visualization

Naval Research Enterprise Internship Program Summer 2002
Research Intern
- Extended algorithms for autonomous vehicle control using multi-objective optimization

Teaching Experience

Graduate Teaching Assistant Spring 2008, 2010
Brown University CSCI1660–Introduction to Computer Security
- Created homework questions, researched course material, and prepared lecture slides

Guest Lecturer
- "Lock picking for Dummies," Brown University CSCI1660–2007, 2008, 2010
- "Amazon Web Services Tutorial," Brown University CSCI 2950T–2008, 2009, 2011, and Brandeis University COSI 12B–2011

Awards & Honors

- 2009 Google Workshop for Women Engineers Invitee
- 2006–2007 Navy Long Term Training Award
- 2005 Society of Women Engineers Helen Martha Sternberg Award

Academic Service

- Panelist–Brown U. Women in Computer Science (WiCS), "Applying to Graduate School," December 2012
- External Reviewer–Very Large Databases, August 2009
- Student Volunteer Coordinator–SIGMOD, June 2009
- External Reviewer–Extending Database Technology, March 2009
- Mentor–WiCS Mentoring Program, 2006–2007
- Mentor–Women at RPI Mentoring Program, 2001–2003

Publications

Jennie Duggan, Olga Papaemmanouil, Uğur Çetintemel, and Eli Upfal. Contender: A Resource Modeling Approach for Concurrent Query Performance Prediction. To appear in *EDBT* 2014.

Michael Stonebraker, Jennie Duggan, Leilani Battle, and Olga Papaemmanouil. SciDB DMBS Research at M.I.T. In *IEEE Data Eng. Bull.*, 36(4): 21-30 (2013).

Jennie Duggan, Yun Chi, Hakan Hacigümüs, Shenghuo Zhu, and Uğur Çetintemel. Packing Light: Portable Workload Performance Prediction for the Cloud. In *ICDE Workshops*, pages 258-265, 2012.

Jennie Duggan, Uğur Çetintemel, Olga Papaemmanouil, and Eli Upfal. Performance Prediction for Concurrent Database Workloads. In *SIGMOD*, pages 337-348, 2011.

Jennie Rogers, Roman Simakov, Emad Soroush, Pavel Velikhov, Magdalena Balazinska, David DeWitt, Bobbi Heath, David Maier, Samuel Madden, Jignesh Patel, Michael Stonebraker, Stanley Zdonik, Artyom Smirnov, Konstantin Knizhnik, and Paul Brown. Overview of SciDB: Large Scale Array Storage, Processing, and Analysis. In *SIGMOD*, pages 963-968, 2010.

Jennie Rogers, Olga Papaemmanouil, and Uğur Çetintemel. A Generic Auto-Provisioning Framework for Cloud Databases. In *ICDE Workshops*, pages 63-68, 2010.

Philippe Cudre-Mauroux, Hideaki Kimura, Kian-Tat Lim, Jennie Rogers, Roman Simakov, Emad Soroush, Pavel Velikhov, Daniel Wang, Magdalena Balazinska, Jacek Becla, David DeWitt, Bobbi Heath, David Maier, Samuel Madden, Jignesh Patel, Michael Stonebraker, and Stanley Zdonik. A Demonstration of SciDB: A Science-Oriented DBMS. In *Proc. VLDB Endow.*, 2(2), pages 1534-1537, 2009.

Yanif Ahmad, Olga Papaemmanouil, Uğur Çetintemel, and Jennie Rogers. Simultaneous Equation Systems for Query Processing on Continuous-Time Data Streams. In *ICDE*, pages 666-675, 2008.

Michael Stonebraker, Chuck Bear, Uğur Çetintemel, Mitch Cherniack, Tingjian Ge, Nabil Hachem, Stavros Harizopoulos, John Lifter, Jennie Rogers, and Stanley Zdonik. One Size Fits All? Part 2: Benchmarking Studies. In *CIDR*, pages 173-184, 2007.

Under Preparation / Submission

Jennie Duggan, Olga Papaemmanouil, Leilani Battle, and Michael Stonebraker. Skew-Aware Join Execution for Array Databases.

Jennie Duggan and Michael Stonebraker. Incremental Elasticity for Array Databases.

Jennie Duggan. The Case for Personal Data-Driven Decision Making.

Marco Serafini, Essam Mansour, Jennie Duggan, Rebecca Taft, Aaron Elmore, Ashraf Aboulnaga, Andrew Pavlo, Michael Stonebraker. Incremental Data Placement for Main Memory Database Systems.

Miscellaneous Writings

Jennie Duggan, contributor to "Chapter 2: Physical Security," in Michael Goodrich and Roberto Tamassia. Introduction to Computer Security. Addison Wesley, October 2010.

Jennie Duggan. "Tunnels, Bunkers and Nukes: My Underground Vacation." *Conduit!* 20(1), Brown Computer Science.

The Database Group. "Major Database Conference Comes to Providence." *Conduit!* 18(2), Brown Computer Science.

Invited Talks

- "Incremental Data Placement for Array Databases," *Qatar Computing Research Institute*, October 2013.
- "Incremental Elasticity for Scientific Databases," *University of Massachusetts Lowell*, August 2013.
- "Performance Prediction for Concurrent Database Workloads," SIGMOD, June 2011.
- "A Generic Auto-Provisioning Framework for Cloud Databases," Self-Managing Database Systems, January 2010.

Departmental Service

Faculty-Graduate Liaison, Brown CS 2010–2011

- Addressed departmental matters between grad students and faculty, including allocation of funds, career resources, and curriculum issues
- Aided students navigating the graduate program

Faculty Search Czar, Brown CS 2008, 2011

- Organized and ran student meetings with applicants
- Reviewed and ranked candidates and presented findings to the department chair and committee

Graduate Student Admissions Czar, Brown CS 2012

- Read, rated, and commented on student applications
- Contributed to a short list of candidates

Michael Stonebraker
Adjunct Professor of Computer Science
Massachusetts Institute of Technology
32 Vassar St.
Cambridge, MA 02139 USA
Email: stonebraker@csail.mit.edu

Samuel Madden
Professor of Computer Science
Massachusetts Institute of Technology
32 Vassar St.
Cambridge, MA 02139 USA
Email: madden@csail.mit.edu

Olga Papaemmanouil
Assistant Professor of Computer Science
Brandeis University
415 South St,
Waltham, MA, 02454
Email: olga@cs.brandeis.edu

Uğur Çetintemel
Professor of Computer Science
Brown University
Box 1910, 115 Waterman St.
Providence, RI 02912 USA
Email: ugur@cs.brown.edu

Eli Upfal
Professor of Computer Science
Brown University
Box 1910, 115 Waterman St.
Providence, RI 02912 USA
Email: eli@cs.brown.edu