

Aaron J. Elmore

Assistant Professor
Department of Computer Science
1100 E 58th Street
University of Chicago
Chicago, IL 60637
aelmore@cs.uchicago.edu

Research Areas

Database as a Service • Polystore Systems • Self-Managed Elastic Systems

Education

- 2014 PHD in Computer Science, University of California Santa Barbara.
Thesis: Elasticity Primitives for Database as a Service.
Advisors: Divyakant Agrawal • Amr El Abbadi
- 2009 MS in Computer Science, University of Chicago.
- 2002 BS in Electronic Commerce Technologies, DePaul University. Magna Cum Laude.

Grants, Honors, and Awards

- 2015–2016 Intel Science and Technology Center for BigData PI
- 2014 UCSB CS Outstanding Graduate Student
- 2011,2013 SIGMOD Student Travel Grant
- 2013 UCSB Senate Travel Grant
- 2010 Amazon Research Grant, Winter 2011
- 2009 Outstanding Teaching Assistant, UCSB
- 2009–2010 Computer Science Merit Fellowship

Publications

- 2017 S. Huang, L. Xu, Jialin Liu, A.J. Elmore, A. Parameswaran. “OrpheusDB: Bolt-on Versioning for Relational Databases.” *PVLDB* 10(10): 1130-1141, 2017.
- 2017 A. Shanbhag, A. Jindal, S. Madden, J. Quiane, A.J. Elmore. “A Robust Partitioning Scheme for Ad-Hoc Query Workloads”, *SoCC*, 2017.
- 2017 L. Xu, S. Huang, S. Hui, A.J. Elmore, A. Parameswaran. “OrpheusDB: A Lightweight Approach to Relational Dataset Versioning” *SIGMOD*: 1655-1658, 2017.
- 2017 D. Tang, H. Jiang, and A.J. Elmore. “Adaptive Concurrency Control: Despite the Looking Glass, One Concurrency Control Does Not Fit All”, *CIDR*, 2017.
- 2016 M.Serafini, R. Taft, A.J. Elmore, A. Pavlo, A. Abounaga, and M. Stonebraker. “Clay:

- Fine-Grained Adaptive Partitioning for General Database Schemas”, *VLDB*, Vol 10(4): 445-456, 2016.
- 2016 M.Maddox, D. Goehring, A.J. Elmore, S. Madden, A. Parameswaran, and A. Deshpande. “Decibel: The Relational Dataset Branching System”, *VLDB*, Vol 9(9): 624-635, 2016.
- 2016 R. Taft, W. Lang, J. Duggan, A. J. Elmore, M. Stonebraker, and D. DeWitt. “STeP: Scalable Tenant Placement for Managing Database-as-a-Service Deployments”, *SoCC*, 2016.
- 2016 V. Gadepally, J. Duggan, A.J. Elmore, J. Kepner, S. Madden, T. Mattson, and M. Stonebraker. “The BigDAWG Architecture”, *NEDB*, 2016.
- 2015 A.J. Elmore, V. Arora, R. Taft, A. Pavlo, D. Agrawal, and A. El Abbadi . “Squall: Fine-Grained Live Reconfiguration for Partitioned Main Memory Databases”, *SIGMOD*, 2015.
- 2015 R. Taft, E. Mansour, M. Serafini, J. Duggan, A.J. Elmore, A. Abounaga, A. Pavlo, and M. Stonebraker. “E-Store: Fine-Grained Elastic Partitioning for Distributed Transaction Processing Systems”, *VLDB*, Vol. 8, 2015.
- 2015 A. Bhardwaj, S. Bhattacharjee, A. Chavan, A. Deshpande, A.J. Elmore, S. Madden, and A. Parameswaran. “DataHub: Collaborative Data Science & Dataset Version Management at Scale”, *CIDR* 2015.
- 2015 A. Bhardwaj, A. Deshpande, A.J. Elmore, D. Karger, S. Madden, and A. Parameswaran, E. Wu, and R. Zhang. “Collaborative Data Analytics with DataHub”, *VLDB (demo)*, 2015.
- 2015 A.J. Elmore, J. Duggan, M. Stonebraker, M. Balazinska, U. Cetintemel, V. Gadepally, J. Heer, B. Howe, J. Kepner, et al. “A Demonstration of the BigDAWG Polystore System”, *VLDB (demo)*, 2015.
- 2015 J. Duggan, A. J. Elmore, M. Stonebraker, M. Balazinska, B. Howe, J. Kepner, S. Madden, D. Maier, T. Mattson, and S. Zdonik. “The BigDAWG Polystore System”, *SIGMOD Record*, 44(3), 2015.
- 2015 J. Duggan, A.J. Elmore, T. Kraska S. Madden, T. Mattson and M. Stonebraker. “The BigDawg Architecture and Reference Implementation”, *NEDB* 2015.
- 2015 A. Dziedzic, J. Duggan, A.J. Elmore, V. Gadepally, and M. Stonebraker. ”BigDAWG: A Polystore for Diverse Interactive Applications”, *IEEE Viz Data Systems for Interactive Analysis*, 2015
- 2013 Aaron J. Elmore, Carlo Curino, Divyakant Agrawal, Amr El Abbadi. “Towards Database Virtualization for Database as a Service”, *VLDB* 2013 (Tutorial).
- 2013 Aaron J. Elmore, Sudipto Das, Alexander Pucher, Divyakant Agrawal, Amr El Abbadi, Xifeng Yan. “Characterizing Tenant Behavior for Placement and Crisis Mitigation in Multitenant DBMSs”, *ACM International Conference on Management of Data (SIGMOD)* 2013: 517-528.
- 2012 Stacy Patterson, Aaron J. Elmore, Faisal Nawab, Divyakant Agrawal, Amr El Abbadi. “Serializability, not Serial: Concurrency Control and Availability in Multi-Datacenter Data-

stores”, *VLDB* 2012: 1459-1470.

- 2012 Aaron J. Elmore, Sudipto Das, Divyakant Agrawal, Amr El Abbadi. “InfoPuzzle: Exploring Group Decision Making in Mobile Peer-to-Peer Databases”, *VLDB* 2012: 1998-2001.
- 2012 Divyakant Agrawal, Amr El Abbadi, Beng Chin Ooi, Sudipto Das, Aaron J. Elmore. “The evolving landscape of data management in the cloud”, *International Journal of Computational Science and Engineering* 2012.
- 2011 Aaron J. Elmore, Sudipto Das, Divyakant Agrawal, Amr El Abbadi. “Zephyr: Live Migration in Shared Nothing Databases for Elastic Cloud Platforms”, *ACM International Conference on Management of Data (SIGMOD)* 2011: 301-312.
- 2011 Divyakant Agrawal, Amr El Abbadi, Sudipto Das, Aaron J. Elmore. “Database Scalability, Elasticity, and Autonomy in the Cloud”. *16th International Conference on Database Systems for Advanced Applications (DASFAA)* 2011: 2-15.
- 2011 Aaron J. Elmore, Sudipto Das, Divyakant Agrawal, Amr El Abbadi. “Towards an Elastic and Autonomic Multitenant Database”, *6th International Workshop on Networking Meets Databases (NetDB)* 2011.

TALKS

- 2015 “Building an Elastic Main-Memory Database: E-Store” University of Waterloo Database Seminar Series
- 2015 “DataHub: Collaborative Data Science & Dataset Version Management at Scale” GCASR
- 2015 “The BigDawg Architecture and Reference Implementation” NEDB
- 2013 “Self-Managed Elasticity for OLTP and Analytics”, High Performance Transaction Systems (HPTS). Fall 2013.
- 2013 “Towards Database Virtualization for Database as a Service”, *VLDB* 2013.
- 2013 “Characterizing Tenant Behavior for Placement and Crisis Mitigation in Multitenant DBMSs”, *SIGMOD* 2013.
- 2012 “MapReduce: Distributed Programming for Commodity Clusters”, University of Southern California Graduate Datamining Course. Spring 2012.
- 2011 “Towards an Elastic and Autonomic Multitenant Database”, *NetDB* 2011.
- 2010–2012 “Databases in the Cloud”, University of Chicago Graduate Database Course. Fall 2010, Spring 2010, Fall 2011, Fall 2012.

Professional Experience

June 2015– *Assistant Professor*

UNIVERSITY OF CHICAGO

Chicago, IL

Assistant Professor of Computer Science and the College of the University of Chicago.

- 2014–2015 *Postdoctoral Associate*
MIT Cambridge, MA
Working with Sam Madden and Mike Stonebraker on various research projects related to elastic databases and user-driven data analytics.
- 2013 *Software Engineering Intern*
TRIFACTA San Francisco, CA
Worked with Joe Hellerstein, Adam Silberstein, and Sean Kandel on tools for interactive data analytics, which resulted in two patents applications.
- 2012 *Research Intern*
MICROSOFT RESEARCH, EXTREME COMPUTING GROUP Redmond, WA
Worked with Phil Bernstein on scale-out performance for a cloud computing programming model.
- 2010 *Software Engineering Intern*
AMAZON WEB SERVICES Seattle, WA
Developed an internal, distributed, automated, deployment for monitoring service configurations.
- 2010–2013 *Research Assistant*
NATIONAL CENTER FOR ECOLOGICAL ANALYSIS AND SYNTHESIS (NCEAS) Santa Barbara, CA
Helped design and develop JEDI, a web enabled spatial information system to encourage collaboration on jellyfish bloom populations between ecologists.
- 2010–2014 *Research Assistant*
DISTRIBUTED SYSTEMS LAB, UCSB Santa Barbara, CA
Researched large scale, distributed and multitenant database systems.
- 2008–2009 *Research Assistant*
COMPUTATION INSTITUTE / UNIVERSITY OF CHICAGO Chicago, IL
Developed GWT web based module for Generalized Labels Over Scientific data Studies (GLOSS), a collaborative data tagging system. Researched bootstrapping based information extraction algorithm for identifying taggable elements based on training data.
- 2007–2012 *Chief Architect*
CUSTOMORE Chicago, IL
Designed and implemented a pay-per-click (PPC) management and tracking platform.
- 2006–2008 *Software Engineer*
1SYNC / GS1US Chicago, IL
Helped lead design and implementation of re-platforming the Universal Product Code (UPC/barcode) registry application. Developed UI modules for extensible data objects for Item Management, a GDSN data pool synchronization platform— with 2.7 million items synchronized between retailers and suppliers, including many Fortune 50 companies.
- 2005–2006 *Software Engineer*
JC WHITNEY Chicago, IL
Developed J2EE applications for internal systems and customer facing website, which was in the top 100 E-commerce sites based on revenue. One of three developers responsible for 24/7 support. Led search engine optimization efforts which increased search traffic by 33%.
- 2002–2005 *Software Engineer*
THE INCREMENTUM GROUP, LLC Chicago, IL
Lead development for customer service, reporting and billing applications.

2000–2002 *Computer Science Tutor*
DEPAUL UNIVERSITY
Tutored undergraduate and graduate students on various CS topics.

Chicago, IL

Professional Activities

2017 Student Research Competition Co-Chair SIGMOD
2016–2017 Proceedings Co-Chair VLDB
2016 Co-chair Demonstration Track SIGMOD
2016 VLDB, CIDR, and SIGMOD PC Member
2014 CIKM and ICDE PC Member
2014 ACM SIGMOD Demo PC Member
2012–2013 UCSB Computer Science Faculty Recruitment Graduate Representative
2010–2011 UCSB Computer Science Graduate Representative President
2010–2011 UCSB Computer Science Graduate Admissions Committee Representative
2011–2013 External reviewer for ODBASE 2011, VLDB 2012, COMAD 2012-2013, Middleware 2012
2011–2012 Reviewer for Transactions on Computers, Transactions on Storage
2011 Helped organize the NSF Workshop “Science of Cloud” held in March 2011

TEACHING

2015–2016 Instructor for CMCS 23500/33550 “Introduction to Database Systems” and CAPP 30235
“Databases for Public Policy”
2014 Co-lecturer of MIT 6.830 “Database Systems”
2013 Co-instructor of UCSB CS595 seminar “Recent Trends in Computing Research”
2009–2012 UCSB Teaching Assistant for CS20 “Programming Methods”, CS174a “Database Systems”,
CS271 “Adv. Topics in Distributed Systems”, CS274 “Adv. Topics in Databases”