

Research Interests

My dissertation research lies at the intersection of **Big Data management** and **distributed Cloud Computing systems**. Specifically, I work on Global-Scale Data Management, where data management systems span multiple geographically separated data centers for fault-tolerance and proximity to users. Looking forward, I plan to leverage my background to explore opportunities in building systems for **data science** and infrastructure for **global connectivity** of autonomous and mobile applications, especially in the context of the Internet of Things.

Education

- 2012-present** Ph.D. in Computer Science - The University of California, Santa Barbara
Dissertation topic: Global-Scale Data Management With Strong Transactional Guarantees
Advisors: Divyakant Agrawal and Amr El Abbadi
- 2009-2011** M.S. in Computer Science - University of Science and Technology (KAUST), Saudi Arabia
Thesis title: TMAC: Timestamp-Ordered MAC for CSMA/CA Wireless Mesh Networks
Advisor: Dr. Basem Shihada
- 2005-2009** B.S. in Computer Engineering - King Fahd University (KFUPM), Saudi Arabia

Research Experience

- 2012-present** UC Santa Barbara, Computer Science Department, Santa Barbara, CA
Graduate Student
Advisors: Divyakant Agrawal and Amr El Abbadi
- Global-Scale Data Management**
To ensure the reliability of Cloud and Big Data applications in the face of datacenter-scale failures, and to bring data closer to users to improve performance, data and applications are deployed across multiple datacenters. In this line of work, I study various fundamental aspects of global-scale data management, such as global-scale transactional data processing [P3, P4, P5, P6, P11, P15, P18], data communication [P13], data placement and configuration [P17], and multi-representation and heterogeneous data processing [P19,P20].
- June 2014-Aug 2016** Hewlett-Packard (HP) Labs, Palo Alto, CA
Research Associate, Systems Group
Mentor: Dr. Terence Kelly
- Large-Scale Data Management Over Non-Volatile Memory**
I worked on designing data stores for non-volatile memory architectures. I studied the implications of emerging flush-on-fail CPU technology on the durability cost of transactions [P9, P10, P12]. I also worked on a design of a data store that improves performance by avoiding two main overhead sources: logging and cache-line flushing. This is achieved by embedding logging information as part of the update to the data structure and periodic whole-CPU-cache flushing [X1, P22].
- June 2013-Sep 2013** Microsoft Research (MSR), Redmond, WA
Research Intern, Database Group
Mentor: Dr. David Lomet
- High Performance Temporal Indexing on Modern Hardware**
This project combines the high performance of lock-free indexing with the functionality of multi-versioned B-trees. Specifically, I worked on the Time-Split Bw-tree (TSBw-tree) that integrates the algorithms of the Time-split B-tree within the lock-free implementation of the Bw-tree [P14].

2009-2011 University of Science and Technology (KAUST), Computer Science Department, Saudi Arabia
Graduate Student
Advisor: Dr. Basem Shihada

Timestamp-Ordered MAC for CSMA/CA Wireless Mesh Networks

This project tackles the problem of unfairness in Wireless Mesh Networks, where TCP flows experience different performance characteristics depending on their location in the network. A MAC-layer solution is developed to transparently improve TCP fairness. The proposed MAC layer, called TMAC, uses a timestamp-ordering technique to achieve fairness. A Markov chain model is developed to analytically study the characteristics of TMAC and simulations using ns-3 are conducted to validate the design [P1][P2][P7][T1].

Teaching Experience

Mentoring Students

- *Colin Biafore*, B.S. at UC Santa Barbara (2015-2016)
Project: Graph Summarization for Social Networks Trends Detection
★ The work is presented at SIGMOD 2016 as part of the undergraduate research competition.
★ Colin is the Computer Science Department nominee for UCSB's Undergraduate Research Award.
- *Tanuj Mittal*, M.S. at UC Santa Barbara (2016-2017)
Project: Reads Consistency in Globally Distributed Databases
- *Ravi Kumar*, M.S. at UC Santa Barbara (2016-2017)
Project: Consensus with Globally Distributed Agents
- *Darshan Maiya*, M.S. at UC Santa Barbara (2016-2017)
Project: Elastic Resource Allocation for Machine Learning Tasks
- *Alan Buzdar*, M.S. at UC Santa Barbara (2016)
Project: Processing Geo-Replicated Spatial Data
- *Emily Littleworth*, B.S. at UC Santa Barbara (2015)
Project: Geo-Replication of Graph Databases
- *Charles Weng*, B.S. at UC Santa Barbara (2015)
Project: Geo-Replication of Ordered Logs

Teaching

- Teaching Assistant for a graduate level course (CS271: Advanced Topics in Distributed Systems)
★ Outstanding TA award by UC Santa Barbara's Computer Science department.
- Prepared and presented a 3-hour tutorial in SIGMOD 2016 (to appear in ICDE 2017)
Title: The Challenges of Global-Scale Data Management
- Presented as a guest lecturer for multiple courses at UC Santa Barbara:
CS271 Advanced Topics in Distributed Systems
CS274 Advanced Topics in Databases
CS171 Distributed Systems

Patents

- X1. Faisal Nawab, Joseph Izraelevitz, Terence Kelly, Charles B. Morrey, and Dhruva Chakrabarti
Memory System to Access Uncorrupted Data
Patent Application Filed (March 2016)

Conference and journal publications

Research papers (15): 3 in VLDB, 3 in EDBT, and in SIGMOD, ICDE, CIDR, DISC, ICDCS, and others.

Journal papers (2): Elsevier's Ad Hoc Networks, and Springer's Databases in Networked Information Systems.

Others (5): 1 Tutorial in SIGMOD and ICDE, 1 Demo in SIGMOD, and 3 extended abstracts.

- P22. Faisal Nawab, Joseph Izraelevitz, Terence Kelly, Charles B. Morrey, Dhruva Chakrabarti, Michael L. Scott
Dali: A Periodically Persistent Hash Map
The International Symposium on DIStributed Computing (**DISC**), 2017.
- P21. Vaibhav Arora, Faisal Nawab, Divyakant Agrawal, Amr El Abbadi
Typhon: Consistency Semantics for Multi-Representation Data Processing
IEEE International Conference on Cloud Computing (**CLOUD**), 2017.
- P20. Vaibhav Arora, Faisal Nawab, Divyakant Agrawal, Amr El Abbadi
Multi-Representation Based Data Processing Architecture for IoT Applications
IEEE International Conference on Distributed Computing (**ICDCS**), 2017.
- P19. Faisal Nawab, Divyakant Agrawal, Amr El Abbadi, Sanjay Chawla
COP: Planning Conflicts for Faster Parallel Transactional Machine Learning
International Conference on Extending Database Technology (**EDBT**), 2017.
- P18. Faisal Nawab, Divyakant Agrawal, Amr El Abbadi
The Challenges of Global-scale Data Management
ACM International Conference on Management of Data (**SIGMOD**), pages 2223–2227, 2016.
IEEE International Conference on Data Engineering (**ICDE**), 2017 (**Tutorial**)
- P17. Victor Zakhary, Faisal Nawab, Divyakant Agrawal, Amr El Abbadi
DB-Risk: The Game of Global Database Placement
ACM International Conference on Management of Data (**SIGMOD**), pages 2185–2188, 2016.
(**Demo paper**)
- P16. Colin Biafore, Faisal Nawab
Graph Summarization for Geo-correlated Trends Detection in Social Networks
ACM International Conference on Management of Data (**SIGMOD**), pages 2247–2248, 2016.
(**Undergraduate Research Competition**)
- P15. Faisal Nawab, Vaibhav Arora, Divyakant Agrawal, Amr El Abbadi
Minimizing Commit Latency of Transactions in Geo-Replicated Data Stores
ACM International Conference on Management of Data (**SIGMOD**), pages 1279–1294, 2015.
- P14. David Lomet, Faisal Nawab
High Performance Temporal Indexing on Modern Hardware
IEEE International Conference on Data Engineering (**ICDE**), pages 1203–1214, 2015.
★ Selected as one of the best papers in ICDE.
- P13. Faisal Nawab, Vaibhav Arora, Divyakant Agrawal, Amr El Abbadi
Chariots: A Scalable Shared Log for Data Management in Multi-datacenter Cloud Environments
International Conference on Extending Database Technology (**EDBT**), pages 13–24, 2015.
- P12. Faisal Nawab, Dhruva Chakrabarti, Terence Kelly, Brad Morrey
Procrastination Beats Prevention: Timely Sufficient Persistence for Efficient Crash Resilience
International Conference on Extending Database Technology (**EDBT**), pages 689–694, 2015.
- P11. Divy Agrawal, Amr El Abbadi, Vaibhav Arora, Ceren Budak, Theodore Georgiou, Hatem Mahmoud, Faisal Nawab, Cetin Shahin, Shiyuan Wang
Mind Your Ps and Vs: A Perspective on the Challenges of Big Data Management and Privacy Concerns
IEEE International Conference on Big Data and Smart Computing (**BigComp**), pages 1–6, 2015.
- P10. Faisal Nawab, Dhruva Chakrabarti, Terence Kelly, Brad Morrey
Zero-overhead NVM Crash Resilience
Non-Volatile Memories Workshop (**NVMW**), 2015 (**Extended abstract**)
- P9. Faisal Nawab, Dhruva Chakrabarti, Terence Kelly, Brad Morrey
Zero-overhead NVM Crash Resilience
USENIX Conference on File and Storage Technologies (**FAST**), 2015 (**Extended abstract**)
- P8. Hatem Mahmoud, Vaibhav Arora, Faisal Nawab, Divyakant Agrawal, Amr El Abbadi
MaaT: Effective and Scalable Coordination of Distributed Transactions in the Cloud
International Conference on Very Large Data Bases (**VLDB**), pages 7(5):329–340, 2014.

- P7. Faisal Nawab, Kamran Jamshaid, Basem Shihada, and Pin-Han Ho
Fair Packet Scheduling in Wireless Mesh Networks
Elsevier Journal of **Ad Hoc Networks**, pages 13:414–427, 2014.
- P6. Faisal Nawab, Divyakant Agrawal, Amr El Abbadi
Message Futures: Fast Commitment of Transactions in Multi-Datacenter Environments
Biennial Conference on Innovative Data Systems Research (**CIDR**), pages 1–10, 2013.
- P5. Hatem Mahmoud, Faisal Nawab, Alexander Pucher, Divyakant Agrawal, Amr El Abbadi
Low-latency Multi-datacenter Databases Using Replicated Commit
International Conference on Very Large Data Bases (**VLDB**), pages 6(9):661–672, 2013.
- P4. Divyakant Agrawal, Amr El Abbadi, Hatem Mahmoud, Faisal Nawab, Ken Salem
Managing Geo-replicated Data in Multi-datacenters
Springer’s Databases in Networked Information Systems (**DNIS**), pages 23–43, 2013.
- P3. Stacy Patterson, Aaron Elmore, Faisal Nawab, Divyakant Agrawal, Amr El Abbadi
Serializability, not Serial: Concurrency Control and Availability in Multi-Datacenter Datastores
International Conference on Very Large Data Bases (**VLDB**), pages 5(11):1459–1470, 2012.
- P2. Faisal Nawab, Kamran Jamshaid, Basem Shihada, and Pin-Han Ho
MAC-Layer Protocol for TCP Fairness in Wireless Mesh Networks
IEEE International Conference on Communications in China (**ICCC**), pages 448–453, 2012.
- P1. Faisal Nawab, Kamran Jamshaid, Basem Shihada, and Pin-Han Ho
TMAC: Timestamp-ordered MAC for CSMA/CA Wireless Mesh Networks
International Conference on Computer Communications and Networks (**ICCCN**), pages 1–6, 2011.

Theses

- T2. *Global-Scale Data Management With Strong Transactional Guarantees*
Ph.D. Dissertation, UC Santa Barbara, *in preparation*, expected in 2017.
- T1. *TMAC: Timestamp-Ordered MAC Protocol for Wireless Mesh Networks*
M.S. Thesis, KAUST (2011).

Presentations

- *The Challenges of Global-scale Data Management*, 3-hour tutorial.
IEEE International Conference on Data Engineering (ICDE), 2017.
- *COP: Planning Conflicts for Faster Parallel Transactional Machine Learning*.
International Conference on Extending Database Technology (EDBT), 2017.
- *The Challenges of Global-scale Data Management*, 3-hour tutorial.
ACM International Conference on Management of Data (SIGMOD), 2016.
- *Minimizing commit latency in of transactions in geo-replicated data stores*.
ACM International Conference on Management of Data (SIGMOD), 2015.
- *High Performance Temporal Indexing on Modern Hardware*.
IEEE International Conference on Data Engineering (ICDE), 2015.
- *Chariots: A Scalable Shared Log for Data Management in Multi-datacenter Environments*.
International Conference on Extending Database Technology (EDBT), 2015.
- *Procrastination Beats Prevention: Timely Sufficient Persistence for Efficient Crash Resilience*.
International Conference on Extending Database Technology (EDBT), 2015.
- *Geo-Replication: A Journey From The Simple to The Optimal*
Big Data UCLA seminar, 2015.
- *MaaT: Effective and Scalable Coordination of Distributed Transactions in the Cloud*.
International Conference on Very Large Data Bases (VLDB), 2014.
- *Message Futures: Fast Commitment of Transactions in Multi-Datacenter Environments*.
Biennial Conference on Innovative Data Systems Research (CIDR), 2013.
- *Low-latency Multi-datacenter Databases Using Replicated Commit*.
International Conference on Very Large Data Bases (VLDB), 2013.
- *TMAC: Timestamp-ordered MAC for CSMA/CA Wireless Mesh Networks*.
International Conference on Computer Communications and Networks (ICCCN), 2011.

Service

Reviewer	IEEE Transactions on Knowledge and Data Engineering (2016) IEEE Transactions on Computers (2016) IEEE/ACM Transactions on Networking (2016) Elsevier Journal of Ad Hoc Networks (2016) Springer's Knowledge and Information Systems (2016) Springer's Cluster Computing (2016) Elsevier Information Processing Letters (2016) Springer's Distributed and Parallel Databases (2015) ACM Transactions on Database Systems (2015)
PC member	ACM International Conference on Information and Knowledge Management (CIKM) (2017) International Conference on Informatics, Health, and Technology (2017) Workshop on Principles and Practice of Consistency for Distributed Data (2016)
External reviewer	International Conference on Extending Database Technology (EDBT), industrial track, 2017. International Conference on Database Systems for Advanced Applications (DASFAA), 2017. International Conference on Big Data and Smart Computing (BigComp), 2017. ACM Symposium on Principles Of Database Systems (PODS), 2016. Brazilian Symposium on Databases (SBBD), 2016. IEEE International Conference on Data Engineering (ICDE), 2015. IEEE Transactions on Computers (TC), 2015. ACM International Conference on Management of Data (SIGMOD), 2015. The International Conference on Networked Systems (NETYS), 2015. ACM International Conference on Management of Data (SIGMOD), 2014. International Conference on Extending Database Technology (EDBT), 2014. The International Conference on Networked Systems (NETYS), 2014. ACM/IFIP/USENIX Middleware conference (MIDDLEWARE), 2012.

Honors and awards

- The paper entitled: "High Performance Temporal Indexing on Modern Hardware" was selected as one of the best papers in IEEE International Conference on Data Engineering (ICDE), 2015.
- Outstanding publication award from UC Santa Barbara's Computer Science department, 2015.
- Travel grants from SIGMOD, ICDE, and the UC Santa Barbara senate, 2015.
- Best presentation award at UC Santa Barbara's graduate student workshop, 2014.
- Outstanding TA award by UC Santa Barbara's Computer Science department, 2013.
- Scholarship from the Saudi Arabian Cultural Mission, 2012-now.
- Academic excellence award from KAUST, 2010.
- Graduate fellowship from KAUST, 2009-2011.
- Discovery scholarship from KAUST, 2008-2009.

References

Divyakant Agrawal
Professor of Computer Science
University of California, Santa Barbara
Santa Barbara, CA 93106
E-mail: agrawal@cs.ucsb.edu

David B. Lomet
Principal Researcher and Research Manager
Database Group
Microsoft Research
Redmond, WA 98052
E-mail: lomet@microsoft.com

Amr El Abbadi
Professor of Computer Science
University of California, Santa Barbara
Santa Barbara, CA 93106
E-mail: amr@cs.ucsb.edu

Terence P. Kelly
Principal Researcher
Systems Lab at Hewlett Packard Labs
Palo Alto, CA 94304
E-mail: terence.p.kelly@hpe.com
E-mail: tpkelly@eecs.umich.edu