

Ashraf Aboulnaga

Qatar Computing Research Institute
PO Box 5825
Doha, Qatar

aaboulnaga@qf.org.qa
<http://www.qcri.org/>
Phone: +974 4454 0630

Education

- **Ph.D. in Computer Science**, University of Wisconsin - Madison, 2002.
Thesis: *Cost Estimation Techniques for Database Systems*.
Advisor: Professor Jeffrey Naughton.
- **M.S. in Computer Science**, University of Wisconsin - Madison, 1999.
- **M.S. in Computer Science**, Faculty of Engineering, Alexandria University, Egypt, 1996.
Thesis: *An Algorithm and a Cost Model for Window Queries in Spatial Databases*.
Advisor: Professor Walid Aref.
- **B.S. in Computer Science**, Faculty of Engineering, Alexandria University, Egypt, 1993.

Professional Experience

- **Qatar Computing Research Institute**, Doha, Qatar.
Principal Scientist and Manager of Distributed Systems Research Center, 2013–present.
- **University of Waterloo**, Waterloo, Ontario, Canada.
Cheriton School of Computer Science.
Associate Professor, 2009–present.
Assistant Professor, 2004–2009.
- **Google**, Waterloo, Ontario, Canada.
Visiting Research Scientist, December 2009–June 2010.
- **IBM Almaden Research Center**, San Jose, California, USA.
Research Staff Member, 2002–2004.
- **University of Wisconsin**, Madison, Wisconsin, USA.
Research Assistant, 1997–2002.
Teaching Assistant, 1996.
- **Microsoft Research**, Redmond, Washington, USA.
Researcher Intern, May–August 1999 and May–August 1998.
- **Alexandria University**, Alexandria, Egypt.
Lecturer, 1993–1996.

Awards

- Best Paper Award at the International Conference on Very Large Data Bases (VLDB), 2011.
- Ontario Early Researcher Award, 2009.
- Google Research Award, 2008.
- IBM Centre for Advanced Studies (CAS) Faculty Fellow, 2006–present.
- IBM CAS Project of the Year Award (2008) and Innovation Impact of the Year Award (2009) for the “XML Design Advisor” project.
- Microsoft Research Graduate Fellowship, 1998–2000.
- B.S. with First Degree of Honor (Summa Cum Laude), 1993.

Professional Memberships

- Senior member of IEEE and IEEE Computer Society.
- Senior member of ACM and ACM SIGMOD.

Publications

Refereed Journal Papers

- [J1] Rebecca Taft, Essam Mansour, Marco Serafini, Jennie Duggan, Aaron J. Elmore, Ashraf Abounaga, Andrew Pavlo, and Michael Stonebraker. “E-Store: Fine-Grained Elastic Partitioning for Distributed Transaction Processing,” *Proceedings of the VLDB Endowment (PVLDB)*, volume 8, number 3, pp. 245–256, November 2014. Presented at the *International Conference on Very Large Data Bases (VLDB ’15)*, Kohala Coast, Hawaii, USA, August 2015.
- [J2] Marco Serafini, Essam Mansour, Ashraf Abounaga, Kenneth Salem, Taha Rafiq, and Umar Farooq Minhas. “Accordion: Elastic Scalability for Database Systems Supporting Distributed Transactions,” *Proceedings of the VLDB Endowment (PVLDB)*, volume 7, number 12, pp. 1035–1046, August 2014. Presented at the *International Conference on Very Large Data Bases (VLDB ’14)*, Hangzhou, China, September 2014.
- [J3] Essam Mansour, Ahmed El-Roby, Panos Kalnis, Aron Ahmadia, and Ashraf Abounaga. “RACE: A Scalable and Elastic Parallel System for Discovering Repeats in Very Long Sequences,” volume 6, number 10, pp. 865–876, August 2013. Presented at the *International Conference on Very Large Data Bases (VLDB ’13)*, Riva del Garda, Italy, August 2013.
- [J4] Iman Elghandour, Ashraf Abounaga, Daniel C. Zilio, and Calisto Zuzarte. “Recommending XML Physical Designs for XML Databases,” *The VLDB Journal (VLDBJ)*, volume 22, number 4, pp. 447–470, August 2013. (extended version of [C23] and [C30])
- [J5] Rui Liu, Ashraf Abounaga, and Kenneth Salem. “DAX: A Widely Distributed Multi-tenant Storage Service for DBMS Hosting,” *Proceedings of the VLDB Endowment (PVLDB)*, volume 6, number 4, pp. 253–264, February 2013. Presented at the *International Conference on Very Large Data Bases (VLDB ’13)*, Riva del Garda, Italy, August 2013.
- [J6] Umar Farooq Minhas, Shriram Rajagopalan, Brendan Cully, Ashraf Abounaga, Kenneth Salem, and Andrew Warfield. “RemusDB: Transparent High Availability for Database Systems,” *The VLDB Journal (VLDBJ)*, volume 22, number 1, pp. 29–45, February 2013. (extended version of [J8], special issue for best papers of VLDB 2011)
- [J7] Iman Elghandour and Ashraf Abounaga. “ReStore: Reusing Results of MapReduce Jobs,” *Proceedings of the VLDB Endowment (PVLDB)*, volume 5, number 6, pp. 586–597, February 2012. Presented at the *International Conference on Very Large Data Bases (VLDB ’12)*, Istanbul, Turkey, August 2012.
- [J8] Umar Farooq Minhas, Shriram Rajagopalan, Brendan Cully, Ashraf Abounaga, Kenneth Salem, and Andrew Warfield. “RemusDB: Transparent High Availability for Database Systems,” *Proceedings of the VLDB Endowment (PVLDB)*, volume 4, number 11, pp. 738–748, August 2011. Presented at the *International Conference on Very Large Data Bases (VLDB ’11)*, Seattle, Washington, USA, August 2011. (**Best Paper Award**)
- [J9] Mumtaz Ahmad, Ashraf Abounaga, Shivnath Babu, and Kamesh Munagala. “Interaction-Aware Scheduling of Report Generation Workloads,” *The VLDB Journal (VLDBJ)*, volume 20, number 4, pp. 589–615, August 2011. (extended version of [C27] and [C31])
- [J10] Ahmed A. Soror, Umar Farooq Minhas, Ashraf Abounaga, Kenneth Salem, Peter Kokosielis, and Sunil Kamath. “Automatic Virtual Machine Configuration for Database Workloads,” *ACM Transactions on Database Systems (TODS)*, volume 35, number 1, article 7 (47 pages), February 2010. (extended version of [C29], special issue for best papers of SIGMOD 2008)
- [J11] Ashraf Abounaga and Walid G. Aref. “Window Query Processing in Linear Quadrees,” *Distributed and Parallel Databases*, volume 10, number 2, pp. 111–126, September 2001.

Refereed Conference and Workshop Papers

- [C12] Jingen Xiang, Huangdong Meng, and Ashraf Abounnaga. “Scalable Matrix Inversion Using MapReduce,” In *Proceedings of the International Symposium on High-Performance Parallel and Distributed Computing (HPDC '14)*, Vancouver, Canada, pp. 177–190, June 2014.
- [C13] Mostafa Ead, Herodotos Herodotou, Ashraf Abounnaga, and Shivnath Babu. “PStorM: Profile Storage and Matching for Feedback-Based Tuning of MapReduce Jobs,” In *Proceedings of the International Conference on Extending Database Technology (EDBT '14)*, Athens, Greece, pp. 1–12, March 2014.
- [C14] Farhana H. Zulkernine, Patrick Martin, Ying Zou, Michael Bauer, Femida Gwadry-Sridhar, and Ashraf Abounnaga. “Towards Cloud-Based Analytics-as-a-Service (CLAAaaS) for Big Data Analytics in the Cloud,” In *Proceedings of the IEEE International Congress on Big Data*, Santa Clara, California, USA, pp. 62–69, June 2013.
- [C15] Jingen Xiang, Cong Guo, and Ashraf Abounnaga. “Scalable Maximum Clique Computation Using MapReduce,” In *Proceedings of the IEEE International Conference on Data Engineering (ICDE '13)*, Brisbane, Australia, pp. 74–85, April 2013.
- [C16] Umar Farooq Minhas, Rui Liu, Ashraf Abounnaga, Kenneth Salem, Jonathan Ng, and Sean Robertson. “Elastic Scale-out for Partition-Based Database Systems,” In *Proceedings of the International Workshop on Self-Managing Database Systems (SMDB '12)*, Washington, DC, USA, April 2012.
- [C17] Muhammad Bilal Sheikh, Umar Farooq Minhas, Omar Zia Khan, Ashraf Abounnaga, Pascal Poupart, and David J. Taylor. “A Bayesian Approach to Online Performance Modeling for Database Appliances using Gaussian Models,” In *Proceedings of the International Conference on Autonomic Computing (ICAC '11)*, Karlsruhe, Germany, pp. 121–130, June 2011.
- [C18] Mumtaz Ahmad, Songyun Duan, Ashraf Abounnaga, and Shivnath Babu. “Predicting Completion Times of Batch Query Workloads Using Interaction-aware Models and Simulation,” In *Proceedings of the International Conference on Extending Database Technology (EDBT '11)*, Uppsala, Sweden, pp. 449–460, March 2011. (*extended version of [C20]*)
- [C19] Hatem A. Mahmoud and Ashraf Abounnaga. “Schema Clustering and Retrieval for Multi-domain Pay-As-You-Go Data Integration Systems,” In *Proceedings of the ACM SIGMOD International Conference on Management of Data (SIGMOD '10)*, Indianapolis, Indiana, USA, pp. 411–422, June 2010.
- [C20] Mumtaz Ahmad, Songyun Duan, Ashraf Abounnaga, and Shivnath Babu. “Interaction-Aware Prediction of Business Intelligence Workload Completion Times,” In *Proceedings of the IEEE International Conference on Data Engineering (ICDE '10)*, Long Beach, California, USA, pp. 413–416, March 2010. (*short paper*)
- [C21] Sean Tozer, Tim Brecht, and Ashraf Abounnaga. “Q-Cop: Avoiding Bad Query Mixes to Minimize Client Timeouts Under Heavy Loads,” In *Proceedings of the IEEE International Conference on Data Engineering (ICDE '10)*, Long Beach, California, USA, pp. 397–408, March 2010.
- [C22] Ashraf Abounnaga, Ziyu Wang, and Zi Ye Zhang. “Packing the Most Onto Your Cloud,” In *Proceedings of the International Workshop on Cloud Data Management (CloudDB '09)*, Hong Kong, China, pp. 25–28, November 2009. (*short paper*)
- [C23] Iman Elghandour, Ashraf Abounnaga, Daniel C. Zilio, and Calisto Zuzarte. “Recommending XMLTable Views for XQuery Workloads,” In *Proceedings of the International XML Database Symposium (XSym '09)*, Lyon, France, pp. 129–144, August 2009.
- [C24] Chen Zhang, Hans De Sterck, Ashraf Abounnaga, Haig Djambazian, and Robert Sladek. “Case Study of Scientific Data Processing on a Cloud Using Hadoop,” In *Proceedings of the International Symposium on High Performance Computing Systems and Applications (HPSC '09)*, Kingston, Ontario, Canada, pp. 400–415, June 2009.
- [C25] Mumtaz Ahmad, Ashraf Abounnaga, and Shivnath Babu. “Query Interactions in Database Workloads,” In *Proceedings of the International Workshop on Testing Database Systems (DBTest '09)*, Providence, Rhode Island, USA, article 11 (6 pages), June 2009.
- [C26] Xin Liu, Ashraf Abounnaga, Kenneth Salem, and Xuhui Li. “CLIC: CLient-Informed Caching for Storage Servers,” In *Proceedings of the USENIX Conference on File and Storage Technologies (FAST '09)*, San Francisco, California, USA, pp. 297–310, February 2009.

- [C27] Mumtaz Ahmad, Ashraf Abounaga, Shivnath Babu, and Kamesh Munagala. “Modeling and Exploiting Query Interactions in Database Systems,” In *Proceedings of the ACM International Conference on Information and Knowledge Management (CIKM '08)*, Napa, California, USA, pp. 183–192, October 2008.
- [C28] Ahmed Atallah, Ashraf Abounaga, and Frank Wm. Tompa. “Records Retention in Relational Database Systems,” In *Proceedings of the ACM International Conference on Information and Knowledge Management (CIKM '08)*, Napa, California, USA, pp. 873–882, October 2008.
- [C29] Ahmed A. Soror, Umar Farooq Minhas, Ashraf Abounaga, Kenneth Salem, Peter Kokosielis, and Sunil Kamath. “Automatic Virtual Machine Configuration for Database Workloads,” In *Proceedings of the ACM SIGMOD International Conference on Management of Data (SIGMOD '08)*, Vancouver, Canada, pp. 953–966, June 2008.
- [C30] Iman Elghandour, Ashraf Abounaga, Daniel C. Zilio, Fei Chiang, Andrey Balmin, Kevin Beyer, and Calisto Zuzarte. “XML Index Recommendation with Tight Optimizer Coupling,” In *Proceedings of the IEEE International Conference on Data Engineering (ICDE '08)*, Cancun, Mexico, pp. 833–842, April 2008.
- [C31] Mumtaz Ahmad, Ashraf Abounaga, Shivnath Babu, and Kamesh Munagala. “QShuffler: Getting the Query Mix Right,” In *Proceedings of the IEEE International Conference on Data Engineering (ICDE '08)*, Cancun, Mexico, pp. 1415–1417, April 2008. (*poster*)
- [C32] Umar Farooq Minhas, Jitendra Yadav, Ashraf Abounaga, and Kenneth Salem. “Database Systems on Virtual Machines: How Much do You Lose?,” In *Proceedings of the International Workshop on Self-Managing Database Systems (SMDB '08)*, Cancun, Mexico, pp. 35–41, April 2008.
- [C33] Kareem El Gebaly and Ashraf Abounaga. “Robustness in Automatic Physical Database Design,” In *Proceedings of the International Conference on Extending Database Technology (EDBT '08)*, Nantes, France, pp. 145–156, March 2008.
- [C34] Ashraf Abounaga and Kareem El Gebaly. “ μ BE: User Guided Source Selection and Schema Mediation for Internet Scale Data Integration,” In *Proceedings of the IEEE International Conference on Data Engineering (ICDE '07)*, Istanbul, Turkey, pp. 186–195, April 2007.
- [C35] Ahmed A. Soror, Ashraf Abounaga, and Kenneth Salem. “Database Virtualization: A New Frontier for Database Tuning and Physical Design,” In *Proceedings of the International Workshop on Self-Managing Database Systems (SMDB '07)*, Istanbul, Turkey, pp. 388–394, April 2007.
- [C36] Ning Zhang, M. Tamer Özsu, Ihab F. Ilyas, and Ashraf Abounaga. “FIX: Feature-based Indexing Techniques for XML Documents,” In *Proceedings of the International Conference on Very Large Data Bases (VLDB '06)*, Seoul, Korea, pp. 259–270, September 2006.
- [C37] Ning Zhang, M. Tamer Özsu, Ashraf Abounaga, and Ihab F. Ilyas. “XSeed: Accurate and Fast Cardinality Estimation for XPath Queries,” In *Proceedings of the IEEE International Conference on Data Engineering (ICDE '06)*, Atlanta, Georgia, USA, pp. 61–72, April 2006.
- [C38] Xuhui Li, Ashraf Abounaga, Kenneth Salem, Aamer Sachedina, and Shaobo Gao. “Second-Tier Cache Management Using Write Hints,” In *Proceedings of the USENIX Conference on File and Storage Technologies (FAST '05)*, San Francisco, California, USA, pp. 115–128, December 2005.
- [C39] Ashraf Abounaga, Peter J. Haas, Mokhtar Kandil, Sam Lightstone, Guy M. Lohman, Volker Markl, Ivan Popivanov, and Vijayshankar Raman. “Automated Statistics Collection in DB2 UDB,” In *Proceedings of the International Conference on Very Large Data Bases (VLDB '04)*, Toronto, Canada, pp. 1146–1157, August 2004.
- [C40] Ihab F. Ilyas, Peter J. Haas, Volker Markl, Paul Brown, and Ashraf Abounaga. “CORDS: Automatic Discovery of Correlations and Soft Functional Dependencies,” In *Proceedings of the ACM SIGMOD International Conference on Management of Data (SIGMOD '04)*, Paris, France, pp. 647–658, June 2004.
- [C41] Ihab F. Ilyas, Peter J. Haas, Volker Markl, Paul Brown, and Ashraf Abounaga. “Automatic Relationship Discovery in Self-Managing Database Systems,” In *Proceedings of the International Conference on Autonomic Computing (ICAC '04)*, New York, New York, USA, pp. 340–341, May 2004. (*poster*)
- [C42] Ashraf Abounaga and Jeffrey F. Naughton. “Building XML Statistics for the Hidden Web,” In *Proceedings of the ACM International Conference on Information and Knowledge Management (CIKM '03)*, New Orleans, Louisiana, USA, pp. 358–365, November 2003.

- [C43] Ashraf Aboulnaga, Alaa R. Alameldeen, and Jeffrey F. Naughton. “Estimating the Selectivity of XML Path Expressions for Internet Scale Applications,” In *Proceedings of the International Conference on Very Large Data Bases (VLDB '01)*, Rome, Italy, pp. 591–600, September 2001.
- [C44] Ashraf Aboulnaga, Jeffrey F. Naughton, and Chun Zhang. “Generating Synthetic Complex-structured XML Data,” In *Proceedings of the International Workshop on the Web and Databases (WebDB '01)*, Santa Barbara, California, USA, pp. 79–84, May 2001.
- [C45] Ashraf Aboulnaga and Jeffrey F. Naughton. “Accurate Estimation of the Cost of Spatial Selections,” In *Proceedings of the IEEE International Conference on Data Engineering (ICDE '00)*, San Diego, California, USA, pp. 123–134, February 2000.
- [C46] Ashraf Aboulnaga and Surajit Chaudhuri. “Self-tuning Histograms: Building Histograms Without Looking at Data,” In *Proceedings of the ACM SIGMOD International Conference on Management of Data (SIGMOD '99)*, Philadelphia, Pennsylvania, USA, pp. 181–192, June 1999.

System Demonstrations

- [D47] Andrei Sambra, Sandro Hawke, Tim Berners-Lee, Lalana Kagal, and Ashraf Aboulnaga. “CIMBA - Client-Integrated MicroBlogging Architecture,” Demonstration at the *International Semantic Web Conference (ISWC '14)*, Riva del Garda, Italy, pp. 57–60, October 2014.
- [D48] Iman Elghandour and Ashraf Aboulnaga. “ReStore: Reusing Results of MapReduce Jobs in Pig,” Demonstration at the *ACM SIGMOD International Conference on Management of Data (SIGMOD '12)*, Scottsdale, Arizona, USA, pp. 701–704, May 2012. (*demonstration of the work in [J7]*)
- [D49] Iman Elghandour, Ashraf Aboulnaga, Daniel C. Zilio, Fei Chiang, Andrey Balmin, Kevin Beyer, and Calisto Zuzarte. “An XML Index Advisor for DB2,” Demonstration at the *ACM SIGMOD International Conference on Management of Data (SIGMOD '08)*, Vancouver, Canada, pp. 1267–1270, June 2008. (*demonstration of the work in [C30]*)
- [D50] Ashraf Aboulnaga, Kareem El Gebaly, and Daniel Wong. “ μ BE: User Guided Source Selection and Schema Mediation for Internet Scale Data Integration,” Demonstration at the *IEEE International Conference on Data Engineering (ICDE '07)*, Istanbul, Turkey, pp. 1513–1514, April 2007. (*demonstration of the work in [C34]*)
- [D51] Ihab F. Ilyas, Volker Markl, Peter J. Haas, Paul G. Brown, and Ashraf Aboulnaga. “CORDS: Automatic Generation of Correlation Statistics in DB2,” Demonstration at the *International Conference on Very Large Data Bases (VLDB '04)*, Toronto, Canada, pp. 1341–1344, August 2004. (*demonstration of the work in [C40]*)

Tutorials

- [T52] Ashraf Aboulnaga and Shivnath Babu. “Workload Management for Big Data Analytics,” Tutorial at the *ACM SIGMOD International Conference on Management of Data (SIGMOD '13)*, New York, New York, USA, pp. 929–932, June 2013.
- [T53] Ashraf Aboulnaga and Shivnath Babu. “Workload Management for Big Data Analytics,” Advanced technology seminar at the *IEEE International Conference on Data Engineering (ICDE '13)*, Brisbane, Australia, p. 1249, April 2013.
- [T54] Ashraf Aboulnaga, Cristiana Amza, and Kenneth Salem. “Virtualization and Databases: State of the Art and Research Challenges,” Tutorial at the *International Conference on Extending Database Technology (EDBT '08)*, Nantes, France, pp. 746–747, March 2008.
- [T55] Ashraf Aboulnaga, Cristiana Amza, and Kenneth Salem. “Virtualization and Databases: State of the Art and Research Challenges,” Advanced technology seminar at the *IEEE International Conference on Data Engineering (ICDE '07)*, Istanbul, Turkey, April 2007.

Other Papers

- [O56] Ashraf Abounnaga, Beng Chin Ooi, and Patrick Valduriez. “Guest Editorial for Special Section on Data-Intensive Cloud Infrastructure,” *The VLDB Journal (VLDBJ)*, volume 23, number 6, p. 843, December 2014.
- [O57] Ashraf Abounnaga and Carlo Curino. “Report on the Second International Workshop on Data Management in the Cloud (DMC 2013),” *IEEE Data Engineering Bulletin*, volume 36, number 2, pp. 2–4, June 2013. (*invited*)
- [O58] Ashraf Abounnaga and Chee Yong Chan. “Letter from the Guest Editors,” *Proceedings of the VLDB Endowment (PVLDB)*, volume 6, number 4, p. x, February 2013.
- [O59] Ashraf Abounnaga and Kenneth Salem. “Report on the Fourth International Workshop on Self-Managing Database Systems (SMDB 2009),” *IEEE Data Engineering Bulletin*, volume 32, number 4, pp. 2–5, December 2009. (*invited*)
- [O60] Ashraf Abounnaga, Kenneth Salem, Ahmed A. Soror, Umar Farooq Minhas, Peter Kokosielis, and Sunil Kamath. “Deploying Database Appliances in the Cloud,” *IEEE Data Engineering Bulletin*, volume 32, number 1, pp. 13–20, March 2009. (*invited*)
- [O61] Jeffrey F. Naughton, David J. DeWitt, David Maier, Ashraf Abounnaga, Jianjun Chen, Leonidas Galanis, Jaewoo Kang, Rajasekar Krishnamurthy, Qiong Luo, Naveen Prakash, Ravishankar Ramamurthy, Jayavel Shanmugasundaram, Feng Tian, Kristin Tufte, Stratis Viglas, Yuan Wang, Chun Zhang, Bruce Jackson, Anurag Gupta, and Rushan Chen. “The Niagara Internet Query System,” *IEEE Data Engineering Bulletin*, volume 24, number 2, pp. 27–33, June 2001. (*invited*)

Technical Reports

- [R62] Muhammad Bilal Sheikh, Umar Farooq Minhas, Omar Zia Khan, Ashraf Abounnaga, Pascal Poupart, and David J. Taylor. “A Bayesian Approach to Online Performance Modeling for Database Appliances using Gaussian Models,” Technical Report CS-2011-13, Cheriton School of Computer Science, University of Waterloo, 2011. (*extended version of [C17]*)
- [R63] Mumtaz Ahmad, Ashraf Abounnaga, Shivnath Babu, and Kamesh Munagala. “Interaction-Aware Scheduling of Report Generation Workloads,” Technical Report CS-2010-22, Cheriton School of Computer Science, University of Waterloo, December 2010. (*preprint of [J9]*)
- [R64] Iman Elghandour, Ashraf Abounnaga, Daniel C. Zilio, Fei Chiang, Andrey Balmin, Kevin Beyer, and Calisto Zuzarte. “XML Index Recommendation with Tight Optimizer Coupling,” Technical Report CS-2007-22, Cheriton School of Computer Science, University of Waterloo, July 2007. (*extended version of [C30]*)
- [R65] Ashraf Abounnaga and Jeffrey F. Naughton. “Towards Building XML Statistics for the Hidden Web,” Technical Report TR-CS-03-1477, Computer Sciences Department, University of Wisconsin - Madison, April 2003. (*extended version of [C42]*)
- [R66] Ashraf Abounnaga and Surajit Chaudhuri. “Accurate Query Optimization by Sub-plan Memoization,” Technical Report MSR-TR-1999-102, Microsoft Research, December 1999. (*not otherwise published*)

Patents

- [P67] Ashraf Abounnaga, Andrey Balmin, Kevin Beyer, Iman Elghandour, Fei Chiang, Daniel C. Zilio, and Calisto Zuzarte. “Index Selection for XML Database Systems,” US Patent No. 8,229,920, filed August 2007, issued July 2012.
- [P68] Volker Markl, Peter J. Haas, Ashraf Abounnaga, Vijayashankar Raman, and Felix Endres. “System and Method for Updating Database Statistics According to Query Feedback,” US Patent No. 7,831,592, filed October 2004, issued November 2010.
- [P69] Ashraf Abounnaga, Peter J. Haas, Sam Lightstone, Volker Markl, Ivan Popivanov, and Vijayshankar Raman. “Method, System and Program for Prioritizing Maintenance of Database Tables,” US Patent No. 7,363,324, filed December 2004, issued April 2008.

[P70] Surajit Chaudhuri and Ashraf Abounaga. “Query Optimization by Sub-plan Memoization,” US Patent No. 6,850,925, filed May 2001, issued February 2005.

[P71] Ashraf Abounaga and Surajit Chaudhuri. “Self-tuning Histogram and Database Modeling,” US Patent No. 6,460,045, filed March 1999, issued October 2002.

Research Funding

All amounts are given as totals for the entire grant duration, in Canadian dollars. For grants with multiple investigators and/or where overhead was charged, my share of the grant, net of overhead, is given in parentheses.

- \$50,000. Natural Sciences and Engineering Research Council of Canada (NSERC) Discovery Grant, “Systems and Algorithms for Big Data Analytics,” 2013–2014.
- \$318,316. Equipment Grant from Canada Foundation for Innovation (CFI) Leaders Opportunity Fund and Ontario Research Fund - Research Infrastructure (ORF-RI), “Research Infrastructure for Computational Analytics on Big Data in the Cloud,” (Hans De Sterck co-PI), 2013–2014.
- \$25,000. Natural Sciences and Engineering Research Council of Canada (NSERC) Strategic Network on Business Intelligence, “Feedback-driven Tuning of Hadoop Performance Parameters,” PI: Frank Tompa (with 1 co-PI), 2012–2013.
- \$25,000. Natural Sciences and Engineering Research Council of Canada (NSERC) Strategic Network on Business Intelligence, “Focusing on the Data in Web-Scale Data Integration,” PI: M. Tamer Özsü (with 1 co-PI), 2012–2013.
- \$3,347,385 (net: \$127,695). Ontario Research Fund for Research Excellence, “Delivering Ultra-large-scale Services,” PI: Patrick Martin, Queen’s University (with 10 co-PIs), 2011–2016.
- \$10,000. Amazon Web Services Research Grant (AWS credits in support of research), “Database Services on the Cloud,” 2011.
- \$190,000 (net: \$150,000). Ontario Early Researcher Award (ERA) from the Ontario Ministry of Research and Innovation with matching from the University of Waterloo, “Improving Database Management in Cloud Computing Environments,” 2009–2014.
- \$93,250 (net: \$71,200). Natural Sciences and Engineering Research Council of Canada (NSERC) Strategic Network on Business Intelligence, “Best-Effort Data Integration,” PI: Frank Tompa (with 2 co-PIs), 2009–2012.
- \$40,000 (net: \$30,691). Google Research Award, “Structured Mashups on Demand,” 2008.
- \$135,000. Natural Sciences and Engineering Research Council of Canada (NSERC) Discovery Grant, “Robustness in Database Management,” 2008–2013.
- \$426,951 (net: \$142,317). Natural Sciences and Engineering Research Council of Canada (NSERC) Collaborative Research and Development Grant, “Dynamic Resource Allocation for Virtual Database Appliances,” PI: Kenneth Salem (with 2 co-PIs), 2008–2011.
- \$1,724,802 (net: \$178,666). Ontario Centres of Excellence (OCE), “Centre of Excellence for Research in Adaptive Systems (CERAS),” 11 co-PIs, 2007–2010.
- \$85,000. IBM Centre for Advanced Studies (CAS), Toronto, “Database Virtualization,” 2007–2009.
- \$146,296 (net: \$29,259). IBM Shared University Research (SUR) Equipment Grant, “Autonomic Computing Lab at the University of Waterloo,” PI: Johnny Wong (with 4 co-PIs), 2006.
- \$164,000. Natural Sciences and Engineering Research Council of Canada (NSERC) Collaborative Research and Development Grant, “Database Design and Tuning for XML Workloads,” 2006–2010.
- \$85,000. IBM Centre for Advanced Studies (CAS), Toronto, “XML Design Advisor,” 2006–2008.
- \$962,325 (net: \$180,970). Equipment Grant from Canada Foundation for Innovation (CFI) New Opportunities Fund and Ontario Innovation Trust (OIT), “Resources For Computational Research Innovation,” PI: Kevin Hare (with 13 co-PIs), 2006–2008.
- \$66,000. Natural Sciences and Engineering Research Council of Canada (NSERC) Discovery Grant, “Advanced Querying of the Web,” 2005–2008.

- \$58,000. Software Telecommunications Group, University of Waterloo, “XML View Caching for the Efficient Execution of Web Queries,” 2005–2007.
- \$80,000. School of Computer Science, University of Waterloo, Startup Funds, 2004–2008.

Students and Postdoctoral Fellows Supervised

Current Graduate Students and Postdoctoral Fellows

Ph.D. Students

- Kareem El Gebaly.
- Ahmed El-Roby.

Former Students and Postdoctoral Fellows

Postdoctoral Fellows

- Rui Liu, 2010–2013. (*co-supervised with Kenneth Salem*)
First Employment: Senior Software Engineer at LinkedIn, Mountain View, California.
- Iman Elghandour, 2011–2012.
First Employment: Assistant Professor at Alexandria University, Egypt.

Ph.D.

- Umar Farooq Minhas, December 2012.
Thesis: Scalable and Highly Available Database Systems in the Cloud.
First Employment: Research Staff Member at the IBM Almaden Research Center, San Jose, California.
- Mumtaz Ahmad, November 2012.
Thesis: Query Interactions in Database Systems.
First Employment: Software Developer at Auvik Networks, Waterloo.
- Iman Elghandour, October 2010.
Thesis: Automatic Physical Design for XML Databases.
First Employment: Postdoctoral Fellow at the University of Waterloo.

Master’s (M.Math.)

- Huangdong Meng, January 2014.
Thesis: High Availability for Database Systems in Geographically Distributed Cloud Computing Environments.
First Employment: Software Development Engineer at Amazon, Toronto.
- Jingen Xiang, September 2013.
Thesis: Scalable Scientific Computing Algorithms Using MapReduce.
First Employment: Software Developer at SAP, Waterloo.
- Taha Rafiq, May 2013.
Thesis: Elasca: Workload-Aware Elastic Scalability for Partition Based Database Systems.
First Employment: Software Development Engineer at Amazon, Toronto.
- Mostafa Ead, December 2012.
Thesis: PStorM: Profile Storage and Matching for Feedback-Based Tuning of MapReduce Jobs.
First Employment: Software Development Engineer at Amazon, Vancouver.
- Hatem Mahmoud, June 2010.
Thesis: Schema Clustering and Retrieval for Multi-domain Pay-As-You-Go Data Integration Systems.
First Employment: Ph.D. student at the University of California - Santa Barbara.

- Sean Tozer, August 2009. (*co-supervised with Tim Brecht*)
Thesis: Avoiding Bad Query Mixes to Minimize Unsuccessful Client Requests Under Heavy Loads.
First Employment: Senior Developer at the Centre for Phenogenomics, Mount Sinai Hospital, Toronto.
- Ahmed Ataullah, May 2008. (*co-supervised with Frank Tompa*)
Thesis: A Framework for Records Management in Relational Database Systems.
First Employment: Ph.D. student at the University of Waterloo.
- Nickolay Tchervenski, February 2008.
Essay: An Evaluation of Storage Load Balancing Alternatives for Database Systems.
First Employment: Software Developer at the IBM Toronto Laboratory.
- Umar Farooq Minhas, December 2007.
Thesis: A Performance Evaluation of Database Systems on Virtual Machines.
First Employment: Ph.D. student at the University of Waterloo.
- Kareem El Gebaly, August 2007.
Thesis: Robustness in Automatic Physical Database Design.
First Employment: Software Developer at Ericsson Inc., Egypt.

Undergraduate Interns and Co-op Students Supervised

- Fábio Albuquerque (Mitacs Globalink intern), *Universidade Federal da Paraíba*, 2013.
- Aniruddha Gangopadhyay (Intern), *Birla Institute of Technology and Science (BITS)*, 2012.
- Sean Robertson (Co-op student, NSERC USRA), *University of Waterloo*, 2011.
- Jonathan Ng (Co-op student, NSERC USRA), *University of Waterloo*, 2011.
- Ziyu Wang (Co-op student, NSERC USRA), *University of Waterloo*, 2008.
- Zi Ye Zhang (Co-op student, NSERC USRA), *University of Waterloo*, 2008.
- Harikrishna Bikmal (Intern), *Indian Institute of Technology, Bombay*, 2008.
- Jitendra Yadav (Intern), *Indian Institute of Technology, Kanpur*, 2007.
- Hidayath Ansari (Intern), *Indian Institute of Technology, Bombay*, 2006.
- Mayank Singhal (Intern), *Indian Institute of Technology, Bombay*, 2006.
- Daniel Wong (Co-op student), *University of Waterloo*, 2006.
- Daniel Brotherston (Co-op student), *University of Waterloo*, 2006.

Membership on Thesis Committees

- Ph.D. external examiner for:
Aysan Rasooli, McMaster University, 2013.
Madalin Mihailescu, University of Toronto, 2012.
Jun Hu, University of Ottawa, 2011.
Samir Mohammad, Queen's University, 2011.
- Member of the Ph.D. thesis committee of:
Khalid Alsubhi, current.
Xin Liu, 2014.
Patrick Kling, 2012.
Xuhui Li, 2011.
Chen Zhang, 2011.
Rolando Blanco, 2010.
Loubna Mekouar, 2010.
David DeHaan, 2009.

- Master thesis reader for:
Alexander Szlavik, 2013.
Alexey Karyakin, 2011.
Prashant Gaharwar, 2010.
Tim Benke, 2009.
Ye Hu, 2009.
Mosharaf Chowdhury, 2009.
Weihan Wang, 2007.
Mohammed Abouzour, 2007.
E. Cem Sözgen, 2006.
Jie Zhang, 2005.
- Master essay reader for:
Yasemin Ugur Ozekinci, 2007.

Selected Recent Talks

- Elasticity and High Availability for Databases in the Cloud.
Tsinghua University, September 2014.
- Self-Tuning Hadoop.
Qatar University, May 2014.
Carnegie Mellon University - Qatar, February 2014.
University of Toronto, June 2012.
- Workload Management for Big Data Analytics.
IBM Workshop on Big Data Analytics, June 2013.
- High Availability for Database Systems in Cloud Computing Environments.
SAP Waterloo, December 2012.
Workshop on Ultra-large Scale Services, CASCON Conference, November 2012.
Keynote at the International Conference on Informatics and Systems (INFOS), Cairo, May 2012.
Qatar Computing Research Institute, May 2012.
King Abdullah University of Science and Technology, April 2012.
University of Washington, November 2011.
Microsoft Research, November 2011.
Amazon, November 2011.
McMaster University, October 2011.
Egyptian Japanese University of Science and Technology (E-JUST), July 2011.
- Multi-domain Pay-as-you-go Data Integration Systems.
IBM Cogons, Ottawa, December 2010.
- Optimizing Complex Analytics Workloads on the Cloud.
Workshop on Web Data Processing on the Cloud, CASCON Conference, November 2010.
- Application Informed Tuning of Virtualized Environments.
Hong Kong University of Science and Technology, November 2009.
Nile University, Cairo, July 2009.
IBM Cairo Technical Development Center, July 2009.
Cairo Microsoft Innovation Center, July 2009.
Queen's University, June 2009.
Google Waterloo, April 2009.
University of Alberta, November 2008.

Teaching Activities

Courses Taught

Undergraduate

- CS 338 - Computer Applications in Business: Databases.
Winter 2005 (1 section, 81 students).
- CS 348 - Introduction to Database Management.
Winter 2013 (1 section, 94 students), Fall 2007 (1 section, 50 students), Winter 2007 (2 sections, 156 students).
- CS 350 - Operating Systems.
Fall 2012 (2 sections, 113 students), Fall 2011 (2 sections, 119 students), Winter 2011 (2 sections, 129 students),
Winter 2009 (2 sections, 121 students), Spring 2006 (2 sections, 128 students).

Graduate

- CS 848 - Advanced Topics in Data Bases: Databases in Cloud Computing Environments.
Winter 2013 (14 students), Winter 2012 (20 students), Fall 2008 (16 students).
- CS 848 - Advanced Topics in Data Bases: Self-managing Database Systems.
Fall 2007 (11 students), Winter 2007 (11 students), Winter 2006 (6 students), Spring 2005 (13 students).

Scientific Activities

Advisory Board Membership

- ClevrU, 2011–present.

Journal Editing

- Member of the Editorial Board of Distributed and Parallel Databases (DAPD), 2014–present.
- Associate Editor of the IEEE Transactions on Knowledge and Data Engineering (TKDE), 2012–present.
- Guest Co-editor (with Beng Chin Ooi and Patrick Valduriez) of the VLDB Journal Special Section on Data-intensive Cloud Infrastructure, December 2014.

Conference and Workshop Organization

- Associate Editor of the Proceedings of the VLDB Endowment (PVLDB), volume 6, 2012–2013.
Papers published in this volume of PVLDB were presented at the International Conference on Very Large Databases (VLDB), Riva del Garda, Italy, August 2013.
- Co-chair (with Carlo Curino) of the International Workshop on Data Management in the Cloud (DMC), 2013.
Held in conjunction with the IEEE International Conference on Data Engineering (ICDE), Brisbane, Australia, April 2013.
- Senior Program Committee Member of the ACM Conference on Information and Knowledge Management (CIKM), 2012.
- Co-chair (with Kenneth Salem) of the International Workshop on Self-Managing Database Systems (SMDB), 2009.
Held in conjunction with the IEEE International Conference on Data Engineering (ICDE), Shanghai, China, March 2009.

Program Committee Membership

- ACM SIGMOD International Conference on Management of Data, 2006, 2007 (Demonstration Committee), 2009, 2011–2014 (all years).
- International Conference on Very Large Databases (VLDB), 2007 (Demonstration Committee), 2009–2012 (all years), 2015, 2016.
- IEEE International Conference on Data Engineering (ICDE), 2005–2009 (all years), 2009 (Demonstration Committee), 2011, 2013, 2014, 2015.
- Conference on Extending Database Technology (EDBT), 2006, 2015.
- ACM Conference on Information and Knowledge Management (CIKM), 2006, 2008, 2011, 2013, 2014.
- ACM Symposium on Cloud Computing (SoCC), 2013, 2014.
- ACM SIGKDD Knowledge Discovery and Data Mining Conference (KDD), 2013.
- IEEE International Conference on Distributed Computing Systems (ICDCS), 2007, 2009.
- International Conference on Web Information System Engineering (WISE), 2012.
- ACM International Workshop on Web Information and Data Management (WIDM), 2004, 2005, 2007, 2008.
- International Conference on Scientific and Statistical Database Management (SSDBM), 2007, 2013.
- International XML Database Symposium (XSym), 2005, 2006, 2007, 2009.
- USENIX Conference on File and Storage Technologies (FAST), 2007.
- International Workshop on Self-Managing Database Systems (SMDB), 2011, 2012, 2014.
- International Workshop on Reliable Data Services and Systems (RDSS), 2014.
- International Workshop on End-to-end Management of Big Data (BigData), 2012.
- Conference of the Centre for Advanced Studies on Collaborative Research (CASCON), 2009, 2011.
- SIGMOD/PODS Ph.D. Symposium, 2012.
- VLDB Ph.D. Workshop, 2007, 2008.

Grant Reviewing for Government Funding Agencies

- Natural Sciences and Engineering Research Council of Canada (NSERC):
Reviewer for Discovery Grants, 2005–2014 (all years).
Reviewer for Strategic Project Grants, 2013.
Reviewer for Collaborative Research and Development Grants, 2005, 2013.
- National Science Foundation, USA:
Panelist on Division of Information and Intelligent Systems (IIS) Grant Review Panel, 2003 (two panels), 2011.
External reviewer for Division of Information and Intelligent Systems (IIS), 2010.
- Research Network on Mathematics of Information Technology and Complex Systems (MITACS), Canada:
Reviewer for Accelerate Program, 2010, 2012, 2013, 2014.
Reviewer for Elevate Program, 2013.
- Qatar National Research Fund:
Reviewer for National Priorities Research Program, 2009.
Reviewer for Undergraduate Research Experience Program, 2007–2009 (all years).
- Research Grant Council (RGC) of Hong Kong:
Reviewer for General Research Fund, 2013, 2014.
Reviewer for Collaborative Research Fund, 2013.

Journal Reviewing

- ACM Transactions on Database Systems, 2001–2007 (all years), 2013.
- The VLDB Journal, 2004, 2007, 2008, 2009, 2010, 2013, 2014.
- IEEE Transactions on Knowledge and Data Engineering, 2003–2006 (all years), 2009, 2010, 2011.
- ACM Computing Surveys, 2008.
- IEEE Computer, 2011.
- Information Systems, 2011, 2012, 2014.
- Distributed and Parallel Databases, 2005, 2008, 2010, 2011, 2012, 2014.
- SIGMOD Record, 2012.
- Data and Knowledge Engineering, 2001, 2004, 2005.
- Egyptian Informatics Journal, 2013.
- Kuwait Journal of Science and Engineering, 2011.

Other Scientific Activities

- Member of the SIGMOD Jim Gray Doctoral Dissertation Award Committee, 2013–present.
- Member of the Executive Committee of the IEEE Computer Society Data Engineering Workgroup on Self-Managing Database Systems, 2010–present.
- Member of the Editorial Advisory Board of the book “Large Scale and Big Data: Processing and Management,” CRC Press, 2014.

University Service at the University of Waterloo

Faculty and University Level

- Member of the Software Engineering Board, 2011–2013.
- Committee member:
Faculty of Mathematics Web Design Committee, 2011–2012.

Cheriton School of Computer Science

- Committee chair:
Web Committee, 2011–2013.
- Committee member:
Graduate Academic Plans Committee, 2010–2013.
Tenure and Promotion Committee, 2011–2012, 2012–2013.
Graduate Recruiting Committee, 2010–2012.
Undergraduate Recruiting Committee, 2005–2009.
External Relations Committee, 2004–2005.

November 24, 2014