Jignesh M. Patel

Work Address

1210 W. Dayton St.
University of Wisconsin
Madison, WI 53706-1685
USA

Phone: 608.263.7308
Fax: 608.262.9777
E-mail: jignesh@cs.wisc.edu
http://www.cs.wisc.edu/~jignesh

Current Research Interests

Data management and analytics.

Education

Ph.D.	Computer Sciences, May 1998. University of Wisconsin—Madison. Dissertation: Efficient Database Support for Spatial Applications. Advisor: David J. DeWitt.
M.S.	Computer Sciences, May 1993. University of Wisconsin—Madison.
B.Tech.	Bachelor of Technology in Computer Science & Engineering (with honors), May 1991. Institute of Technology-Benaras Hindu University (IIT-Varanasi), India.

Professional Experience

June 2010-present	Professor, Computer Sciences Department, University of Wisconsin-Madison. Also affiliated with the Department of Biostatistics and Medical Informatics, University of Wisconsin-Madison. Since 2020 also co-Chair of the UW Creative Destruction Lab (CDL-WI).
June 2017-present	DataChat Inc., Co-founder, CEO and President.
Sept. 2019-present	Wisconsin Economic Development Corp. (WEDC) Entrepreneurship and Innovation Committee, Member. Committee created by WEDC at the request of the Wisconsin Governor to assist entrepreneurs in the state.
April 2014-present	Board Member/Director, Lands' End Inc.; Chair of the Technology Committee (2014-17), Member of the Audit Committee (2018-present).

June 2015-June 2016	Chief Scientist, Pivotal Inc.
May 2015-June 2018	Board Member/Director, Redox (a Madison, WI healthcare integration startup).
June 2013-present	JMP Consulting LLC, founder. Providing consulting services on data analytics to companies such as American Family Insurance, Johnson Controls, Samsung, and Twitter.
May 2007-Dec. 2013	Co-founder and Chairman, Locomatix Inc. since inception, CEO since 2010. Locomatix became part of Twitter in the Summer of 2013. As CEO led the company to win numerous awards including "20 Hottest Startups" (MobileBeat in 2010), "Top 16 Innovative Startups" (CTIA E&A in 2010) and selected into the "Innovation Showcase" (CTIA Wireless 2011).
2009-2018	Founder and organizer of the UW CS NEST Contest – a software entrepreneurship contest that encourages creative software startups.
Sept. 2008-June 2010	Associate Professor, Computer Sciences Department, University of Wisconsin – Madison. Also affiliated with the Department of Biostatistics and Medical Informatics.
May 2005-Aug. 2008	Associate Professor, Department of Electrical Engineering and Computer Science, University of Michigan, Ann Arbor, MI. Also affiliated with the Bioinformatics Program at the University of Michigan.
July 2012-July 2013	Advisor, EatStreet.com (a Madison, WI startup for online food ordering).
June 2006-May 2008	Director of Media Analysis (2006-07), Technical Advisory Board (2007-08), Zattoo Inc. (a video streaming startup).
Aug. 1999-May 2005	Assistant Professor, Department of Electrical Engineering and Computer Science, University of Michigan, Ann Arbor, MI.
Feb. 1998-Aug. 1999	Consulting Software Engineer, Advanced Development Ctr., NCR, Madison, WI.
May 1995-Sept. 1995	Summer Intern, IBM T. J. Watson Research Center, Hawthorne, NY.

AWARDS

- Vilas Associates Award, 2019.
- IEEE Fellow, 2018.
- Fellow of the ACM, 2014.
- SACM Student Choice Professor of the Year Award, 2014.
- IEEE Senior Member, 2014.
- Inaugural M-List winner. Named as one of the top technology entrepreneurs in Madison, WI, 2013.
- University of Wisconsin Police Department (UWPD) Chief's award, 2012.
- Best papers of VLDB 2012.
- ACM Distinguished Scientist, 2011.
- Best papers of ICDE 2011.
- Best papers of SIGMOD 2011.
- Best papers of ICDE 2010.
- DaMoN 2010 – Best paper award.
- DaMoN 2009 – Best paper award.
- College of Engineering Education Excellence Award, University of Michigan, 2008.
- Google Faculty Award, 2012
- NSF CAREER Award, 2001.
- IBM Faculty Award, 2001 and 2003.

GRANTS AND RESEARCH GIFTS

- "Elements: Software: Towards Efficient Embedded Data Processing," National Science Foundation (NSF), OAC- 1835446, 2019-2021, $ 599,800. Role: PI.
- "CRISP: Center for Research on Intelligent Storage and Processing-in-memory," Semiconductor Research Council (SRC), 2018-22, $27.5M, UW component: $3,697,805. Role: Theme Leader, and co-PI.
- Collaboration with the Microsoft Gray Research Lab, $328,500, 2016-17. Role: co-PI.
- Collaboration with Google/YouTube, $52,000, 2017-18. Role: PI.
- "TRACE: TRacking and Analysis of Causality at Enterprise-level," DARPA, 2015-19, $5.3M, UW Component: $799,999. Role: co-PI.
- "The Center for Predictive Computational Phenotyping," National Institutes of Health (NIH), 2014-19, $10,048,809. Role: co-PI.

- Google Focus Award, gift donation from Google, 2015-16, $500,000. Role: co-PI.
- Google Focus Award, gift donation from Google, 2014-15, $500,000. Role: co-PI.
- "BIGDATA: Small: DCM: Data Management for Analytics Applications on Modern Architecture," IIS-1250886, 2013-16, $680,916. Role: PI.
- "QuickLog – Dynamic Modeling and Monitoring of Log Data," gift donation from Oracle, 2013, $100,000. Role: PI.
- Google Focus Award, gift donation from Google, 2013-14, $500,000. Role: co-PI.
- Collaboration with the Microsoft Gray Research Lab, $464,682, 2013-14. Role: co-PI.
- "Comparison of NoSQL and SQL Databases for Petascale Genomics Applications," Exploratory grant from Marshfield Clinic, WI, 2013, $32,510. Role: PI.
- "CSR: Small: Accelerating Towards the Hardware Specialization Era: A Holistic Approach," National Science Foundation (NSF), CSR-1218432, 2012-15, $387,496. Role: co-PI.
- "Rethinking Data Analytics for Modern Multi-core and Main Memory Environments," gift donation from Google, 2012, $50,000. Role: PI.
- Collaboration with the Microsoft Gray Research Lab, $465,821, 2012-13. Role: co-PI.
- "III: Large: Collaborative Research: SciDB – An Array Oriented Data Management System for Massive Scale Scientific Data," IIS-1110948, 2011-15, $370,706 (UW portion). Role: co-PI.
- "Energy-Efficient Data Processing," National Science Foundation (NSF), IIS-0963993, 2010-2014, $783,846. Role: PI.
- Collaboration with the Microsoft Gray Research Lab, $466,464, 2011-12. Role: co-PI.
- "Data processing in Smart SSDs," gift donation from Samsung, $75,000, 2011. Role: PI.
- "Integrated Biological Sequence Data Management," National Science Foundation (NSF), BDI-0543272, 2006-2009, extended to 2010, $575,105. Role: PI.
- "CRI-IAD: Collaborative Research: Enabling Security and Network Management Research for Future Networks," National Science Foundation (NSF), CNS-0751116, 2008-14, $368,773. Role: co-PI.
- "COMET: An Efficient and Scalable Trajectory Data Management System," National Science Foundation (NSF), IIS-0414510, 2005-08, extended to 2010, $270,000. Role: PI.
- "Secure Coordination and Communication in a Crisis Using Hand-Held Devices," Department of Homeland Security (DHS), W911NF-05-1-0415, 2005-08, $1,365,397. Role: co-PI.
- "Nemo LEAPS Integration," Department of Navy, N00024-01-D-7017, 2006-2008, Subcontract amount: $129,253. Role: co-PI.

- "Proteomics Alliance for Cancer Research," Michigan Economic Development Corporation (MEDC), 2005-08, $2,363,605. Role: co-I.
- "National Center for Integrative Biomedical Informatics", National Institutes of Health (NIH), U54-DA021519-01A1, 2006-10, $18,780,596. Role: co-I.
- "Antidepressants, Concurrent Treatments, and Completed Suicide in VA Registry Data," National Institutes of Health (NIH), 1R01MH078698, 2006-09, Year 1 subcontract: $9,9527. Role: co-I.
- "Quickstep: An Architecture-Conscious DBMS," National Science Foundation (NSF), NSF CAREER Award IIS-0093059, 2001-06, extended to 2007, $349,172. Role: PI.
- "Virtual Center for Network and Security Data," Department of Homeland Security (DHS), 2004-07. 1st year contract: $1,256,068. Role: co-PI.
- "Integration of Bioengineering & Biocomputing to Advance Michigan Computer-Assisted Surgery Research," Michigan Economic Development Corporation (MEDC), $137,725, 2003-06. Subcontract on a grant from MEDC, lead by Dr. Vipin Kumar at Wayne State University.
- "TIMBER: A Native XML Database System," National Science Foundation (NSF), IIS-0208852, 2002-04, $165,000. Role: co-PI.
- "Graph Querying for Life Sciences Applications," Microsoft Research Grant, 2006, $35,000. Role: PI.
- "Declarative and Efficient Methods for Biological Data Management," Microsoft Research Grant, 2004, $35,000. Role: PI.
- "Declarative Querying of Biological Data," Microsoft Research Grant, 2004, $40,000. Role: PI.
- "Declarative and Efficient Querying on Protein Structures," Microsoft Research Grant, 2003, $35,000. Role: PI.
- "Periscope: A System for Declarative Querying on Protein Data Sets," IBM Faculty Award, 2003, $40,000. Role: PI.
- "Scalable Processing of Location-Based Notifications," IBM Faculty Award, 2001, $40,000. Role: PI.
- "Modeling and Querying Protein Sequences Using an Object-Relational Database Management System," Eli Lilly & Co., 2001, $59,074. Role: PI.
- "Harnessing the Power of Object-Relational Databases for XML Query Processing," NCR Corp., 2000, $121,000. Role: PI

PROFESSIONAL ACTIVITIES

Conference Officer

- 2021: CIDR (PC co-chair), PVLDB.
- 2020: CIDR (PC co-chair), PVLDB, SIGMOD.
- 2019: CIDR (PC co-chair), PVLDB, SIGMOD.
- 2015: VLDB/PVLDB (Associate Editor), SIGMOD (Group Leader), ICDE (Applications Track co-chair).
- 2014: ICDE (Area chair), VLDB (Panel co-chair).
- 2011: VLDB (Demo PC co-chair), SIGMOD (Group Leader), EDBT (Tutorials Chair).
- 2010: DaMoN.
- 2009: Core Database Technology Program Chair, VLDB.
- 2006: KDD (Treasurer).
- 2005: ICDE (Vice Chair).

Program Committee (Conferences)

- 2020: PVLDB.
- 2019: SIGMOD, PVLDB.
- 2018: SIGMOD, PVLDB.
- 2017: SIGMOD, PVLDB.
- 2016: PVLDB, ICDE.
- 2014: SIGMOD, PVLDB, EDBT, IEEE BigData (Senior PC).
- 2013: PVLDB, PODS (External Review Committee).
- 2008: ICDE, VLDB.
- 2007: BIOKDD, CIDR, ICDE, SIGMOD, VLDB.

Abbreviations

BIDM	International Workshop on Biological Data Management
BIOKDD	ACM SIGKDD Workshop on Data Mining in Bioinformatics
CIDR	Conference on Innovative Data Systems Research
CIKM	ACM International Conference on Information and Knowledge Management
COMAD	International Conference on Management of Data
DaMoN	International Workshop on Data Management on New Hardware
DanaC	Data Analytics in the Cloud
DEXA	International Conference on Database and Expert Systems Applications
DILS	Data Integration in the Life Sciences
EDBT	International Conference on Extending Database Technology
GIS	ACM International Symposium on Geographic Information Systems
GRADES	Graph Data Management Experiences and Systems
ICDE	IEEE International Conference on Data Engineering
ISMB	International Conference on Intelligent Systems for Molecular Biology
KDD	ACM International Conference on Knowledge Discovery and Data Mining
MDM	IEEE International Conference on Mobile Data Management
MobiDE	ACM Workshop on Data Engineering for Wireless and Mobile Access
MSP	ACM SIGPLAN Workshop on Memory System Performance
PODS	Symposium on Principles of Database Systems
SIGMOD	ACM International Conference on Management of Data
SSD	International Symposium on Spatial Databases
SSTD	International Symposium on Spatial and Temporal Databases
STDBM	Workshop on Spatio-Temporal Database Management
VLDB/ PVLDB	International Conference on Very Large Data Bases
WebDB	Workshop on Web and Databases

- 2006: ICDE, KDD, MDM, VLDB.
- 2005: COMAD, DEXA, ISMB, MDM, SIGMOD, VLDB.
- 2004: CIKM, DEXA, ICDE, MDM, VLDB.
- 2003: SIGMOD, ICDE, MDM.
- 2002: EDBT, ICDE, VLDB.

Program Committee (Workshops)

- 2015: DaMoN.
- 2014: DaMoN.
- 2013: WebDB, DanaC, GRADES.
- 2012: DanaC, DaMoN.
- 2006: DILS, GIS, STDBM.
- 2005: BIDM, DaMoN, DILS, GIS, MobiDE, SSTD.
- 2004: BIOKDD.
- 2003: MobiDE
- 2002: MSP.
- 2001: GIS.
- 1999: SSD.

Editorial Board

- Editorial Board, VLDB Journal, July 2013-2017.
- Editorial Board, Foundations and Trends in Databases, 2006-2017.
- Editorial Board, Distributed and Parallel Databases, 2008-2013.
- Associate Editor (Systems and Prototypes), ACM SIGMOD Record, June 1999-2005.
- Associate Editor, IEEE Data Engineering Bulletin, 2004-2005.

Society Membership

- ACM (SIGMOD)
- IEEE (Computer)

Journal Refereeing

- ACM Transactions on Database Systems
- Bioinformatics
- BMC Bioinformatics

- IEEE Transactions on Knowledge and Data Engineering
- IEEE Transactions on Software Engineering
- Information Systems
- Nucleic Acids Research
- VLDB Journal

Panel Discussions, Keynote, and Distinguished Lecture Series, and Invited Talks

- "Towards Hardware-Software Co-Design for Data Processing." Invited CS Colloquium, Northwestern University, October 25, 2017.
- "Towards Hardware-Software Co-Design for Data Processing: A Plea and a Proposal." Invited CS Colloquium, Harvard University, October 8, 2015.
- "Rethinking the Utility of Benchmarks: Towards Benchmark as a Service." Keynote at the Seventh TPC Technology Conference on Performance Evaluation & Benchmarking (TPCTC 2015), Kohala Coast, Hawai'i, August 31, 2015.
- "Towards hardware-software co-design for data analytics: A plea and a proposal." Keynote at the Third International Workshop on In-Memory Data Management and Analytics (IMDM), VLDB Workshop, Kohala Coast, Hawai'i, August 31, 2015.
- "From Data to Insights @ Bare Metal Speed." Keynote at The ACM SIGMOD Conference on Data Management, Melbourne, Victoria, Australia. June 2, 2015.
- "Towards hardware-software co-design for data analytics: A plea and a proposal." Invited talk by SAP and U. Waterloo, Canada. October 15, 2014.
- "Listening to your startup." Keynote at Forward Tech Conference at the Forward Festival. August 27, 2014.
- "Data @ Bare Metal Speed." CSE Colloquium and Distinguished Lecture Series, 2013-2014, University of San Diego, November 25, 2013.
- "To startup or not to startup? Academics/Entrepreneurs share their experiences." In the 39th International Conference on Very Large Data Bases (VLDB), August 29, 2013, Trento, Italy.
- "How to have a successful research career." In the PhD workshop at the 39th International Conference on Very Large Data Bases (VLDB), August 30, 2013, Trento, Italy.
- "Data @ Bare Metal Speed." Invited talk at the Microsoft Research Faculty Summit 2013, July 15, 2013, Redmond, WA.
- "Data @ Bare Metal Speed." Invited talk at the IBM Workshop on Big Data Analytics, June 28, 2013, Thomas J. Watson Research Center, Yorktown Heights, New York.

- "The Inevitable Inflection Point for Big Data: Toward Energy-Conscious Main Memory Data Processing." In coordination with the 2012 US and UK Big Data Week at the Adobe Big Data Analytics meetup, June 20, 2012.
- "Are DBMSs simply dumb data stores for modern applications?" In the First ACM SIGMOD Workshop on Database and Social Networks, June 12, 2011, Athens, Greece.
- "Cloud Databases: What's New?" In the 36th International Conference on Very Large Data Bases (VLDB), Sept 15, 2010, Singapore.
- "Indexing for Success: Effective and Efficient Analysis of Biological Data." The Bioinformatics Distinguished Lecture Series, May 6, 2009, Ohio University, Athens, OH.
- "Green Database Computing: Should We Care?" Panel at International Workshop on Data Management on New Hardware (DaMoN), June 28 2009, Providence, RI.
- "Indexing for Success: Scaling to Handle the Increasing Volume and Complexity of Biological Databases", Keynote Speaker at Ohio Collaborative Conference on Bioinformatics (OCCBIO), 2008.
- "Scientific Data Management: An Orphan in the Database Community?" Panel at the 24th International Conference on Data Engineering (ICDE), April 7-12, 2008, Cancún, México, 2008.
- "Indexing for Success: Effective Index-Based Methods for Querying Biological Sequences," Panel at the 39th Symposium on the Interface: Computer Science and Statistics. Theme: Systems Biology, May 25, 2007, Philadelphia, PA.
- "Architecture-conscious databases: sub-optimization or the next big leap?" Panel at International Workshop on Data Management on New Hardware (DaMoN), June 12 2005, Baltimore, Maryland.

Scientific Review Panel

- National Science Foundation (NSF), 2003, 2006.
- National Institutes of Health (NIH), 2003-2004, Member, ZRG1 SSS-H (Computational Biology) Study Section.
- NIH, October 2006, Member, ZLM1 ZH-R (J2) (Special Emphasis Panel) Study Section.
- NIH 2005 – 2008, Member, NIH BDMA (Biodata Management and Analysis) Study Section.
- NIH/NLM 2014, Member, National Library of Medicine Special Emphasis Panel.

TEACHING EXPERIENCE

- CS 564, Database Management Systems: Design and Implementation: S09, S10, S11, S12, S13, F13, S17 (U. Wisconsin)

- CS 764, Topics in Database Management Systems: F08, F09, F10, F11, F12, S14, F16, F17, F19 (U. Wisconsin)
- EECS 281, Data Structures and Algorithms: W03 (U. Michigan)
- EECS 484, Database Management Systems: W00, F00, F01, F03. F05, F07 (U. Michigan)
- EECS 485, Web Database and Information Systems, W05 (U. Michigan)
- EECS 496, Major Design Experience Professionalism, W06 (U. Michigan)
- EECS 584, Advanced Database Management Systems: F99, F02, F04, F07 (U. Michigan)
- EECS 684, Current Topics in Database Management Systems: W01, W02, W04 (U. Michigan)

PUBLICATIONS

Conference Papers

1. A. G. S. Raj, P. Gu, E. Zhang, A. X. Annie R., J. Williams, R. Halverson, J. M. Patel: Live-coding vs Static Code Examples: Which is better with respect to Student Learning and Cognitive Load? ACE 2020: 152-159
2. P. V. Sandt, Y. Chronis, J. M. Patel: Efficiently Searching In-Memory Sorted Arrays: Revenge of the Interpolation Search? SIGMOD Conference 2019
3. A. Kakaraparthy, J. M. Patel, K. Park, B. Kroth: Optimizing Databases by Learning Hidden Parameters of Solid State Drives. PVLDB 13(4): 519-532 (2019)
4. F. Li, L. Chen, Y. Zeng, A. Kumar, X. Wu, J. Naughton, J. M. Patel: Tuple-oriented Compression for Large-scale Mini-batch Stochastic Gradient Descent. SIGMOD Conference 2019
5. Z. Zhang, H. Deshmukh, J. M. Patel: Data Partitioning for In-Memory Systems: Myths, Challenges, and Opportunities. CIDR 2019
6. A. G. S. Raj, E. Zhang, S. Mukherjee, J. Williams, R. Halverson, J. M. Patel: Effect of Native Language on Student Learning and Classroom Interaction in an Operating Systems Course. ITiCSE 2019: 499-505
7. J. M. Patel, H. Deshmukh, J. Zhu, N. Potti, Z. Zhang, M. Spehlmann, H. Memisoglu, S. Saurabh: Quickstep: A Data Platform Based on the Scaling-Up Approach. PVLDB 11(6): 663-676 (2018)
8. A. G. S. Raj, V. Naik, J. M. Patel, R. Halverson: How to teach "modern C++" to someone who already knows programming? ACE 2018: 97-104
9. A. G. S. Raj, K. Ketsuriyonk, J. M. Patel, R. Halverson: Does Native Language Play a Role in Learning a Programming Language? SIGCSE 2018: 417-422
10. J. Zhu, N. Potti, S. Saurabh, J. M. Patel: Looking Ahead Makes Query Plans Robust. PVLDB

10(8): 889-900 (2017)

11. L. Chen, A. Kumar, J. F. Naughton, J. M. Patel: Towards Linear Algebra over Normalized Data. PVLDB 10(11): 1214-1225 (2017)
12. H. Deshmukh, H. Memisoglu, J. M. Patel: Adaptive Concurrent Query Execution Framework for an Analytical In-Memory Database System. BigData Congress 2017: 23-30.
13. R. J. Leo John, N. Potti, and J. M. Patel: Ava: From Data to Insights Through Conversation. CIDR 2017
14. L. Viswanathan, B. Chandra, W. Lang, K. Ramachandra, J. M. Patel, A. Kalhan, D. J. DeWitt, A. Halverson: Predictive Provisioning: Efficiently Anticipating Usage in Azure SQL Database. ICDE 2017: 1111-1116
15. A. G. S. Raj, K. Ketsuriyonk, J. M. Patel, R. Halverson: What Do Students Feel about Learning Programming Using Both English and Their Native Language? LaTiCE 2017: 1-8
16. A. Kumar, J. F. Naughton, J. M. Patel, X. Zhu: To Join or Not to Join?: Thinking Twice about Joins before Feature Selection. SIGMOD Conference 2016: 19-34
17. A. Floratou and J. M. Patel: Replica Placement in Multi-tenant Database Environments. BigData Congress 2015: 246-253. **Best paper award.**
18. S. Kulkarni, N. Bhagat, M. Fu, V. Kedigehalli, C. Kellogg, S. Mittal, J. M. Patel, K. Ramasamy, and S. Taneja: Twitter Heron: Stream Processing at Scale. *In Proc. of the ACM SIGMOD Conference on Data Management*, June 22-27, 2015, Melbourne, Victoria, Australia.
19. Y. Li, C. Chasseur, and J. M. Patel: A Padded Encoding Scheme to Accelerate Scans by Leveraging Skew. *In Proc. of the ACM SIGMOD Conference on Data Management*, June 22-27, 2015, Melbourne, Victoria, Australia.
20. A. Kumar, J. Naughton, and J. M. Patel: Learning Generalized Linear Models Over Normalized Data. *In Proc. of the ACM SIGMOD Conference on Data Management*, June 22-27, 2015, Melbourne, Victoria, Australia.
21. N. Potti and J. M. Patel: DAQ: A New Paradigm for Approximate Query Processing. VLDB 2015/PVLDB 8(9): 898-909 (2015)
22. Arun Kumar, Mona Jalal, Boqun Yan, Jeffrey F. Naughton, and Jignesh M. Patel: Demonstration of Santoku: Optimizing Machine Learning over Normalized Data. PVLDB 8(12): 1864-1875 (2015) (demo paper).
23. J. Fan, A. G. Soosai Raj, and J. M. Patel: The Case Against Specialized Graph Analytics Engines. CIDR 2015.
24. Q. Zeng, J. M. Patel, and D. Page: QuickFOIL: Scalable Inductive Logic Programming. VLDB 2015/PVLDB 8(3): 197-208 (2014)
25. S. Sridharan and J. M. Patel: Profiling R on a Contemporary Processor. VLDB 2015 /PVLDB 8(2): 173-184 (2014)

26. Y. Li and J. M. Patel: WideTable: An Accelerator for Analytical Data Processing. *VLDB 2014/PVLDB 7(10)*: 907-918 (2014).

27. A. Floratou, F. Bertsch, J. M. Patel, and G. Laskaris: Towards Building Wind Tunnels for Data Center Design. *VLDB 2014/PVLDB 7(9)*: 781-784 (2014).

28. V. R. Gankidi, N. Teletia, J. M. Patel, A. Halverson, and D. J. DeWitt: Indexing HDFS Data in PDW: Splitting the data from the index. VLDB 2014/PVLDB 7(13): 1520-1528 (2014).

29. A. Toshniwal, S. Taneja, A. Shukla, K. Ramasamy, J. M. Patel, S. Kulkarni, J. Jackson, K. Gade, M. Fu, J. Donham, N. Bhagat, S. Mittal, and D. V. Ryaboy: Storm@Twitter. *In Proc. of the ACM SIGMOD Conference on Data Management*, June 22-27, 2014, Snowbird, UT, USA.

30. N. Zhang, J. Tatemura, J. M. Patel, and H. Hacigümüs. Re-evaluating Designs for Multi-Tenant OLTP Workloads on SSD-based I/O Subsystems, *In Proc. of the ACM SIGMOD Conference on Data Management*, June 22-27, 2014, Snowbird, UT, USA.

31. C. Chasseur and J. M. Patel. Design and Evaluation of Storage Organizations For Read-Optimized Main-Memory Databases, PVLDB 2013/VLDB 2014.

32. Y. Li and J. M. Patel. BitWeaving: Fast Scans for Main Memory Data Processing, In *Proc. of the ACM SIGMOD Conference on Data Management*, June 22-27 2013, New York, NY, USA.

33. J. Do, Y. Kee, J. M. Patel, C. Park, K. Park, and D. J. DeWitt. Query Processing on Smart SSDs: Opportunities and Challenges, In *Proc. of the ACM SIGMOD Conference on Data Management*, June 22-27 2013, New York, NY, USA.

34. W. Lang, S. Shankar, and J. M. Patel. "Towards Multi-tenant Performance SLOs." In Proc. of the 28th IEEE International Conference on Data Engineering (ICDE), April 1-5, 2012, Washington DC, USA, pp. 702-713. **Selected as the best papers in ICDE 2012.**

35. A. Floratou, N. Teletia, D. J. DeWitt, J. M. Patel, and D. Zhang. "Can the elephants handle the NoSQL onslaught?" In *Proceedings of the VLDB Endowment*, Volume 5, Number 12, 2012, pp. 1712-1723, and presented at the 38th International Conference on Very Large Databases (VLDB), August 27-31, 2012, Istanbul, Turkey.

36. W. Lang, S, Harizopoulos, J. M. Patel, M. A. Shah, and D. Tsirogiannis. "Towards Energy-Efficient Database Cluster Design." In *Proceedings of the VLDB Endowment*, Volume 5, Number 12, 2012, pp. 1684-1695, and presented at the 38th International Conference on Very Large Databases (VLDB), August 27-31, 2012, Istanbul, Turkey.

37. N. Zhang, J. Tatemura, J. M. Patel, and H. Hacigümüs. "Towards Cost-Effective Storage Provisioning for DBMSs." In *Proceedings of the VLDB Endowment*, Volume 5, Number 4, 2011, pp. 274-285, and presented at the 38th International Conference on Very Large Databases (VLDB), August 27-31, 2012, Istanbul, Turkey. **Selected as the best papers in VLDB 2012.**

38. P. Larson, S. Blanas, C. Diaconu, C. Freedman, J. M. Patel, and M. Zwilling: High-Performance Concurrency Control Mechanisms for Main-Memory Databases. In

Proceedings of the VLDB Endowment, Volume 5, Number 4, 2011, pp. 298-309, and presented at the 38th International Conference on Very Large Databases (VLDB), August 27-31, 2012, Istanbul, Turkey.

39. A. Floratou, J. M. Patel, W. Lang, and A. Halverson. "When free is not really free: What does it cost to run a database workload in the cloud?" In the third International Conference on Performance Evaluation and Benchmarking (TPCTC), Seattle, USA, August 29, 2011.

40. S. Blanas, Y. Li, and J. M. Patel. "Design and Evaluation of Main Memory Hash Join Algorithms for Multi-core CPUs." In *Proc. of the ACM SIGMOD Conference on Data Management*, June 12-16 2011, Athens, Greece, pp. 37-48.

41. Y. Li, A. Terrel, and J. M. Patel. "WHAM: A High-throughput Sequence Alignment Method." In *Proc. of the ACM SIGMOD Conference on Data Management*, June 12-16 2011, Athens, Greece, pp.445-456. **Selected as one the best papers in SIGMOD 2011.**

42. W. Jin and J. M. Patel. "Efficient and Generic Evaluation of Ranked Queries." In *Proc. of the ACM SIGMOD Conference on Data Management*, June 12-16 2011, Athens, Greece, pp. 601-612.

43. J. Do, D. Zhang, J. M. Patel, D. J. DeWitt, J. F. Naughton, and A. Halverson. "Turbocharging DBMS Buffer Pool Using SSDs." In *Proc. of the ACM SIGMOD Conference on Data Management*, June 12-16, 2011, Athens, Greece, pp. 1113-1124.

44. A. Floratou, J. M. Patel, E. J. Shekita, and S. Tata. "Column-Oriented Storage Techniques for MapReduce." In *Proceedings of the VLDB Endowment*, Volume 4, Number 7, 2011, pp. 419-429, and presented at the 37th International Conference on Very Large Databases (VLDB), August 29-Sept 3, 2011, Seattle, WA, USA.

45. W. Lang and J. M. Patel. "Energy Management for MapReduce Clusters." In *Proc. of the 36th International Conference on Very Large Data Bases (VLDB)*, Sept 13-17, 2010, Singapore, pp. 129-139.

46. R. Kandhan, N. Teletia, and J. M. Patel. "SigMatch: Fast and Scalable Multi-Pattern Matching." In *Proc. of the 36th International Conference on Very Large Data Bases (VLDB)*, Sept 13-17, 2010, Singapore, pp. 1173-1184.

47. N. Zhang, Y. Tian, and J. M. Patel. "Discovery-Driven Graph Summarization." In *Proc. of the 26th IEEE International Conference on Data Engineering (ICDE)*, March 1-6, 2010, Long Beach, California, USA, pp. 880-891.

48. A. Floratou, S. Tata, and J. M. Patel. "Efficient and Accurate Discovery of Patterns in Sequence Datasets." In *Proc. of the 26th IEEE International Conference on Data Engineering (ICDE)*, March 1-6, 2010, Long Beach, California, USA, pp. 461-472. **Selected as the best papers in ICDE 2010.**

49. W. Jin, M. Morse, J. M. Patel, M. Ester, and Z. Hu. "Evaluating Skylines in the Presence of Equijoins." In *Proc. of the 26th IEEE International Conference on Data Engineering (ICDE)*,

March 1-6, 2010, Long Beach, California, USA, pp. 249-260.

50. J. Chen and J, M. Patel. "Trajectory Joins and Their Application in Privacy Preservation." In *Proc. of the 17th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems (ACM GIS)*, November 4-6, 2009, Seattle, WA, pp. 266-275.

51. W. Lang and J. M. Patel. "Towards Eco-friendly Database Management Systems." In *Proc. of the Fourth Biennial Conference on Innovative Data Systems Research (CIDR)*, January 4-7, 2009, Asilomar, CA.

52. P. Cudré-Mauroux, H. Kimura, K-T Lim, J. Rogers, R. Simakov, E. Soroush, P. Velikhov, D. Wang, M. Balazinska, J. Becla, D. J. DeWitt, B. Heath, D. Maier, S. Madden, J. M. Patel, M. Stonebraker, and S. B. Zdonik. "A Demonstration of SciDB: A Science-Oriented DBMS." *In the 35th International Conference on Very Large Data Bases (VLDB)*, August 24-28, 2009, Lyon, France, published as PVLDB 2(2), pp. 1534-1537 (demo paper).

53. Y. Tian, R. A. Hankins, and J. M. Patel. "Efficient Aggregation for Graph Summarization." In *Proc. of the ACM SIGMOD Conference on Data Management*, June 9-12 2008, Vancover, Canada, pp. 567-580.

54. Y. Tian, J. M. Patel, V. Nair, S. Martini, and M. Kretzler. "Periscope/GQ: A Graph Querying Toolkit." *In 34th International Conference on Very Large Data Bases (VLDB)*, August 23-28, 2008, Auckland, New Zealand, published as PVLDB 1(2), pp. 1404-1407 (demo paper).

55. Y. Tian and J. M. Patel. "TALE: A Tool for Approximate Large Graph Matching." In *Proc. of the 24th International Conference on Data Engineering (ICDE)*, April 7-12 2008, Cancun, Mexico, pp. 963-972. Full paper, long presentation.

56. M. A. Iwen, W. Lang, and J. M. Patel. "Scalable Rule-Based Gene Expression Data Classification." In *Proc. of the 24th International Conference on Data Engineering (ICDE)*, April 7-12 2008, Cancun, Mexico, pp. 1062-1071. Full paper, long presentation.

57. M. Morse, J. M. Patel, and H. V. Jagadish. "Efficient Skyline Computation over Low-Cardinality Domains." *In Proc. of the 33rd International Conference on Very Large Data Bases (VLDB)*, Sept. 23-28 2007, Vienna, Austria, pp. 267-278.

58. S. Tata, W. Lang, and J. M. Patel. "Periscope/SQ: Interactive Exploration of Biological Sequence Databases." *In Proc. of the 33rd International Conference on Very Large Data Bases (VLDB)*, Sept. 23-28 2007, Vienna, Austria, pp. 1406-1409 (demo paper).

59. M. Morse and J. M. Patel. "An Efficient and Accurate Method for Evaluating Time Series Similarity." In *Proc. of the ACM SIGMOD Conference on Data Management*, June 11-14, 2007, Bejing, China, pp. 569-580.

60. Y. Chen and J. M. Patel. "Efficient Evaluation of All-Nearest-Neighbors Queries." In *Proc. of the 23rd International Conference on Data Engineering (ICDE),* April 16-20, 2007, Istanbul, Turkey.

61. Y. J. Kim and J. M. Patel. "Rethinking Choices for Multi-dimensional Point Indexing:

Making the Case for the Often Ignored Quadtree." In *the Third Biennial Conference on Innovative Data Systems Research (CIDR)*, January 7-10, 2007, Asilomar, CA, 281-291.

62. S. Sinha, F. Jahanian, and Jignesh M. Patel. "WIND: Workload-Aware Intrusion Detection." In *Proc. of the 9th International Symposium on Recent Advances in Intrusion Detection (RAID)*, September 20-22, 2006, Hamburg, Germany, pp. 290-310.

63. S. Tata, J. M. Patel, J. S. Friedman, and A. Swaroop. "Declarative Querying for Biological Sequence Databases." In *Proc. of the 22nd International Conference on Data Engineering (ICDE),* April 3-7, Atlanta, GA, 2006, pp. 87-98.

64. S. Tata, R. A. Hankins, and J. M. Patel. "Practical Suffix Tree Construction." In *Proc. of the 30th International Conference on Very Large Data Bases (VLDB)*, August 31-Sept 3, 2004, Toronto, Canada, pp. 36-47.

65. J. M. Patel, Y. Chen, and V. P. Chakka. "STRIPES: An Efficient Index for Predicted Trajectories." In *Proc. of the ACM SIGMOD Conference on Data Management*, June 13-18, 2004, Maison de la Chimie, Paris, France, pp. 637-646.

66. C. Meek, J. M. Patel, and S. Kasetty. "OASIS: An Online and Accurate Technique for Local-alignment Searches on Biological Sequences." In *Proc. of the 29th International Conference on Very Large Data Bases (VLDB)*, September 9-12, 2003, Berlin, Germany, pp. 910-921.

67. R. A. Hankins and J. M. Patel. "Data Morphing: An Adaptive, Cache-Conscious Storage Technique." In *Proc. of the 29th International Conference on Very Large Data Bases (VLDB)*, September 9-12, 2003, Berlin, Germany, pp. 417-428.

68. S. Tata and J. M. Patel. "PiQA: An Algebra for Querying Protein Data Sets." In Proceedings of the 15th International Conference on Scientific and Statistical Database Management (SSDBM 2003), July 9-11 2003, Cambridge, MA, pp. 141-150.

69. R. A. Hankins and J. M. Patel. "Effect of Node Size on the Performance of Cache-Conscious B+-trees." In *Proc. of the ACM SIGMETRICS Conference on Measurement and Modeling of Computer Systems,* June 10-14, 2003, San Diego, California, pp. 283-294.

70. V. P. Chakka, A. C. Everspaugh, and J. M. Patel. "Indexing Large Trajectory Data Sets with SETI." In *the First Biennial Conference on Innovative Data Systems Research (CIDR)*, January 5-8, 2003, Asilomar, CA, pp. 164-175.

71. Y. Wu, J. M. Patel, and H. V. Jagadish. "Structural Join Order Selection for XML Query Optimization." In *Proc. of the 19th International Conference on Data Engineering (ICDE),* March 5-8, 2003, Banglore, India, pp. 141-152.

72. L. Hammel and J. M. Patel. "Searching on the Secondary Structure of Protein Sequences." In *Proc. of the 28th International Conference on Very Large Data Bases (VLDB)*, August 20-23, 2002, Hong Kong, China, pp. 634-645.

73. Y. Wu, J. M. Patel, and H. V. Jagadish. "Estimating Answer Sizes for XML Queries." In *Proc of the 8th International Conference on Extending Database Technology (EDBT) 2002*, March 24-28 2002, Prague, Czechoslovakia, pp. 590-608.

74. S. Al-Khalifa, H. V. Jagadish, N. Koudas, J. M. Patel, D. Srivastava, and Y. Wu. "Structural Joins: A Primitive for Efficient XML Query Pattern Matching." In *Proc. of the 18th International Conference on Data Engineering (ICDE)*, February 26-March 1, 2002, San Jose, CA, pp. 141-152.

75. M. Annavaram, J. M. Patel, and E. S. Davidson. "Data Prefetching by Dependence Graph Precomputation." In *Proc. of the 28th International Symposium on Computer Architecture (ISCA)*, June 30-July 4, 2001, Goteborg, Sweden, pp. 52-61.

76. M. Annavaram, J. M. Patel, and E. S. Davidson. "Call Graph Prefetching for Database Applications." In *Proc. of the 7th International Symposium on High Performance Computer Architecture (HPCA)*, January 20-24, 2001, Monterrey, Mexico, pp. 281-290.

77. J. M. Patel and D. J. DeWitt. "Clone Join and Shadow Join: Two Parallel Spatial Join Algorithms." In *Proc. of the Eighth ACM Symposium on Advances in Geographic Information Systems (ACM-GIS)*, November 9-10, 2000, Washington D.C., pp.54-61.

78. K. Ramasamy, J. M. Patel, J. F. Naughton, and R. Kaushik. "Optimal Set Containment Joins: The Good, The Bad and The Ugly." In *Proc. of the 26th Int'l Conference on Very Large Databases (VLDB)*, September 10-14, 2000, Cairo, Egypt. pp. 351-362.

79. J. M. Patel, J.-B. Yu, N. Kabra, K. Tufte, B. Nag, J. Burger, N. E. Hall, K. Ramasamy, R. Lueder, C. Ellman, J. Kupsch, S. Guo, D. J. DeWitt, and J. F. Naughton. "Building a Scaleable Geo-Spatial DBMS: Technology, Implementation, and Evaluation." In *Proc. of the ACM SIGMOD Conference on Data Management*, May 13-15, 1997, Tucson, Arizona, pp. 336-347.

80. J. M. Patel and D. J. DeWitt. "Partition Based Spatial-Merge Join." In *Proc. of the ACM SIGMOD Conference on Data Management*, June 4-6, 1996, Montreal, Canada, pp. 259-270.

81. D. J. DeWitt, N. Kabra, J. Luo, J. M. Patel, and J.-B. Yu. "Client-Server Paradise." In *Proc. of the 19th Int'l Conference on Very Large Databases (VLDB)*, September 12-15, 1994, Santiago, Chile, pp. 558-569.

82. J. M. Patel, M. J. Carey, and M. K. Vernon. "Accurate Modeling of the Hybrid Hash Join Algorithm." In *Proc. of the ACM SIGMETRICS Conference on Measurement and Modeling of Computer Systems*, May 16-20, 1994, Nashville, Tennessee, pp. 56-66.

Journal Papers

1. J. M. Patel: Operational NoSQL Systems: What's New and What's Next? IEEE Computer 49(4): 23-30 (2016)

2. J. Power, Y. Li, M. D. Hill, J. M. Patel, and D. A. Wood: Implications of Emerging 3D GPU Architecture on the Scan Primitive. SIGMOD Record 44(1): 18-23 (2015)

3. H. V. Jagadish, J. Gehrke, A. Labrinidis, Y. Papakonstantinou, J. M. Patel, R. Ramakrishnan, and C. Shahabi: Big data and its technical challenges. Commun. ACM 57(7): 86-94 (2014).

4. K. Park, Y-S. Kee, J. M. Patel, J. Do, C. Park, and D. J. DeWitt: Query Processing on Smart

SSDs. IEEE Data Eng. Bull. 37(2): 19-26 (2014).

5. W. Lang, S. Shankar, J. M. Patel, and A. Kalhan: Towards Multi-Tenant Performance SLOs. IEEE Trans. Knowl. Data Eng. 26(6): 1447-1463 (2014).
6. N. Zhang, J. Tatemura, J. M. Patel, and H. Hacigümüs: Toward cost-effective storage provisioning for DBMSs. VLDB J. 23(2): 329-354 (2014).
7. J. B. Hodgin, V. Nair, H. Zhang, A. Randolph, R. C. Harris, R. G. Nelson, E. J. Weil, J. D. Cavalcoli, J. M. Patel, F. C. Brosius, and M. Kretzler: Identification of cross-species shared transcriptional networks of diabetic nephropathy in human and mouse glomeruli. Diabetes. 2013 Jan; 62(1): 299-308.
8. Y. Li, J. M. Patel, and A. Terrell. WHAM: A High-Throughput Sequence Alignment Method. ACM Trans. Database Syst. 37(4): 28 (2012).
9. W. Lang, R. Kandhan, and J. M. Patel. "Rethinking Query Processing for Energy Efficiency: Slowing Down to Win the Race." IEEE Data Eng. Bull. 34(1): 12-23 (2011).
10. A. Floratou, S. Tata, and J. M. Patel. "Efficient and Accurate Discovery of Patterns in Sequence Data Sets." IEEE Trans. Knowl. Data Eng. (TKDE) 23(8): 1154-1168 (2011).
11. W. Lang, M. D. Morse, and J. M. Patel: Dictionary-Based Compression for Long Time-Series Similarity. IEEE Trans. Knowl. Data Eng. (TKDE) 22(11): 1609-1622 (2010).
12. Y. J. Kim and J. M. Patel: Performance Comparison of the R*-Tree and the Quadtree for kNN and Distance Join Queries. IEEE Trans. Knowl. Data Eng. 22(7): 1014-1027 (2010).
13. W. Lang, J. M. Patel, and J. F. Naughton. "On energy management, load balancing and replication." SIGMOD Record 38(4): 35-42 (2009).
14. M. A. Ilgen, K. Downing, K. Zivin, K. J. Hoggatt, H. M. Kim, D. Ganoczy, K. L. Austin, J. McCarthy, J. M. Patel, and M. Valenstein. "Identifying Subgroups of Patients with Depression Who Are at High Risk for Suicide." Journal of Clinical Psychiatry, 2009 Nov; 70(11):1495-500.
15. Y. J. Kim, N. Teletia, V. Ruotti, C. A. Maher, A. M. Chinnaiyan, R. Stewart, J. A. Thomson, and J. M. Patel. "ProbeMatch: Rapid Alignment of Oligonucleotides to Genome Allowing Both Gaps and Mismatches." Bioinformatics 25(11): 1424-1425 (2009).
16. V. G. Tarcea, T. Weymouth T, A. Ade, A. Bookvich, J. Gao, V. Mahavisno, Z. Wright, A. Chapman, M. Jayapandian, A. Ozgür, Y. Tian, J. Cavalcoli, B. Mirel, J. Patel, D. Radev, B. Athey, D. States, and H. V. Jagadish. "Michigan Molecular Interactions r2: From Interacting Proteins to Pathways." *Nucleic Acids Research*, 2009, Vol. 37, Database issue D642-D646, 2009.
17. M. G. Parsons, H. Chung, E. Nick, A. Daniels, S. Liu, and J. M. Patel. "Intelligent Ship Arrangements: A New Approach to General Arrangement." Naval Engineers Journal, 120(3), December 2008, E51-E65.
18. Y. Tian, R. C. McEachin, C. Santos, D. J. States, and J. M. Patel. "SAGA: A Subgraph Matching

Tool for Biological Graphs." *Bioinformatics*, 23(2), 2007, 232-239.

19. M. Morse, J. M. Patel, and W. I. Grosky. "Efficient Continuous Skyline Computation." *Information Systems*, 177(17), September 2007, 3411-3437.
20. M. Morse, J. M. Patel, and W. I. Grosky. "Efficient Evaluation of Radial Queries Using the Target Tree." *International Journal of Bioinformatics Research and Applications*, 177(17), September 2007, 3411-3437.
21. S. Tata and J. M. Patel: Estimating the selectivity of tf-idf based cosine similarity predicates. SIGMOD Record 36(2): 7-12 (2007).
22. Y. J. Kim and J. M. Patel. "A Framework for Protein Structure Classification and Identification of Novel Protein Structures." *BMC Bioinformatics*, 2006, 7:456.
23. K. Runapongsa, J. M. Patel, H. V. Jagadish, Y. Chen, and S. Al-Khalifa. "The Michigan Benchmark: Towards XML Query Performance Diagnostics." *Information Systems,* 31(2), 2006, pp. 73-97.
24. Y. Tian, S. Tata, R. A. Hankins, and J. M. Patel. "Practical Methods for Constructing Suffix Trees." *VLDB Journal,* 14(3), 2005, pp. 281-299.
25. Y. J. Kim, A. Boyd, B. D. Athey, and J. M. Patel. "miBLAST: Scalable Evaluation of a Batch of Nucleotide Sequence Queries with BLAST." *Nucleic Acids Research*, 33(13), 2005, pp. 4335-4344.
26. M. Annavaram, J. M. Patel, and E. S. Davidson. "Call Graph Prefetching for Database Applications." *ACM Trans. on Computing Systems*, 21(4), 2003, pp. 412-444.
27. Y. Wu, J. M. Patel, and H. V. Jagadish. "Using Histograms to Estimate Answer Sizes for XML Queries." *Information Systems,* 28(1-2), 2003, pp. 33-59.
28. H. V. Jagadish, S. Al-Khalifa, A. Chapman, L. Lakshmanan, A. Nierman, S. Paparizos, J. M. Patel, D. Srivastava, Y. Wu, and C. Yu. "TIMBER: A Native XML Database." *VLDB Journal,* 11(4), 2002, pp. 274-291.

Workshop Papers and Whitepapers

1. K. Park, J. Do, N. Teletia, and J. M. Patel: Aggressive buffer pool warm-up after restart in SQL Server. ICDE Workshops 2016: 31-38
2. J. M. Patel and G. Hofheimer: The State of Data Technology in Credit Unions: The Sink-or-Swim Crossroad Ahead. Whitepaper developed for Filene.org. Released: February 16, 2016.
3. J. Power, Y. Li, M. D. Hill, J. M. Patel, and D. A. Wood: Toward GPUs being mainstream in analytic processing: An initial argument using simple scan-aggregate queries. DaMoN 2015: 11
4. W. Lang, J. M. Patel, and S. Shankar. "Wimpy Node Clusters: What About Non-Wimpy Workloads?" In the *Sixth International Workshop on Data Management on New Hardware*

(DaMoN 2010), held in conjunction with SIGMOD 2010, Indianapolis, IN, June 2010. **(Best paper award)**

5. J. Do and J. M. Patel. "Join Processing for Flash SSDs: Remembering Past Lessons." In the *Fifth International Workshop on Data Management on New Hardware (DaMoN 2009)*, held in conjunction with SIGMOD 2009, Providence, RI, June 2009. **(Best paper award)**
6. M. D. Morse, J. M. Patel, and W. I. Grosky. "Efficient Evaluation of Radial Queries using the Target Tree." *In the International Workshop on Biomedical Data Engineering (BMDE2005)*, held in conjunction with ICDE 2005, Tokyo, Japan, April 3-4, 2005.
7. K. Runapongsa, J. M. Patel, R. Bordawekar, and S. Padmanabhan. "XIST: An XML Index Selection Tool." *In the Second International XML Database Symposium (XSym 2004)*, held in conjunction with VLDB 2004 Toronto, Canada, August 29-30 2004, pp. 219-234.
8. J. M. Patel. "The Role of Declarative Querying in Bioinformatics." *Workshop on Data Management for Molecular and Cell Biology,* Feb. 2-3, 2003, Lister Hill Center, NLM, NIH Campus, Bethesda, MD. Also appears in *OMICS: A Journal of Integrative Biology*, 7(1), 2003, pp. 89-91.
9. K. Runapongsa and J. M. Patel. "Storing and Querying XML Data in Object-Relational DBMSs." In *EDBT XML-Based Data Management (XMLDM) Workshop*, March 24, 2002, Prague, Czech Republic, pp. 266-285.
10. J. Han, G. Malan, J. M. Patel, and F. Jahanian. "A Software Environment for Anomaly Detection of Large-Scale IP Networks." In *Information/System Survivability Workshop*, July 3, 2001, Göteborg, Sweden.
11. D. J. DeWitt, J. Luo, J. M. Patel, and J. Yu. "Paradise - A Parallel Geographic Information System." In *Proc. of the First ACM Workshop on Advances in Geographic Information Systems*, November 5, 1993, Arlington, Virginia.

Posters

1. R. Varadarajan, F. Eichinger, J. M. Patel, and M. Kretzler: "Molecular Re-Classification of Renal Disease using Approximate Graph Matching, Clustering and Pattern Mining" (poster paper). In the 18th Annual International Conference on Intelligent Systems for Molecular Biology (ISMB), 2010, Boston.
2. S. Tata and J. M. Patel. "FLAME: Shedding Light on Hidden Frequent Patterns in Sequence Datasets." In *Proc. of the 24th International Conference on Data Engineering (ICDE)*, April 7-12, Cancun, Mexico, 2006.
3. Y. Tian, R. C. McEachin, C. Santos, D. J. States, and J. M. Patel. "SAGA: Subgraph Matching Tool for Biological Graphs." *In the 11th Annual International Conference on Research in Computational Molecular Biology (RECOMB)*, 2007.
4. S. Paparizos, J. M. Patel, and H.V. Jagadish. "SIGOPT: Using Schema to Optimize XML Query Processing." In *Proc. of the 23rd International Conference on Data Engineering (ICDE),* April

16-20, 2007, Istanbul, Turkey, 1456-1460.

5. M. Morse, J. M. Patel, and W. I. Grosky. "Efficient Continuous Skyline Computation," In *Proc. of the 22nd International Conference on Data Engineering (ICDE)*, April 3-7, Atlanta, GA, 2006, pp 108.

Book Chapters

1. Y. Tian and J. M. Patel. “Interactive Graph Summarization.” In *Link Mining: Models, Algorithms, and Applications*, eds. P. S. Yu, J. Han, and C. Faloutsos, Springer, 2011.
2. S. Tata and J. M. Patel. “Query Languages and Evaluation Techniques for Biological Sequence Data.” Encyclopedia of Database Systems 2009: 2261-2264
3. J. M. Patel, D. P. Huddler, and L. Hammel. “Declarative and Efficient Querying on Protein Secondary Structures.” In *Data Mining in Bioinformatics,* eds. J. T. L. Wang, M. J. Zaki, H. Toivonen, and D. Shasha, Springer, 2005, pp. 243-273.
4. J. M. Patel and H. V. Jagadish, "The Michigan Benchmark: A Micro-Benchmark for XML Query Performance Diagnostics.” In *XML Data Management*, eds. A. B. Chaudhri, R. Zicari, and A. Rashid, Eds., 1st ed: Addison-Wesley, 2003.

Software

1. Quickstep team. Quickstep. http://www.quickstep.io, 2018.
2. C. Chasseur and J. M. Patel. The Quickstep Storage Manager. http://research.cs.wisc.edu/quickstep/, first release 2013.
3. Y. Li and J. M. Patel. BitWeaving. http://research.cs.wisc.edu/quickstep/, first release 2013.
4. Y. Li, A. Terrel and J. M. Patel. WHAM: A High-throughput Sequence Alignment Method. http://www.cs.wisc.edu/wham, first release 2011.
5. N. Teletia, Y. J. Kim, and J. M. Patel. ProbeMatch: Rapid short-read sequence aligner. http://pages.cs.wisc.edu/~jignesh/probematch/, first release 2009.
6. S. Tata, R. A. Hankins, and J. M. Patel. TDD Suffix Tree Construction Software. http://www.eecs.umich.edu/tdd, first release: 2005.
7. Y. J. Kim, A. Boyd, B. D. Athey, and J. M. Patel. miBLAST: Scalable Evaluation of a Batch of Nucleotide Sequence Queries with BLAST. http://www.eecs.umich.edu/miblast/, first release: 2005.
8. K. Runapongsa, J. M. Patel, and H. V. Jagadish. The Michigan Benchmark: A Micro-Benchmark for XML Query Processing Systems. http://www.eecs.umich.edu/db/mbench/, first release: 2002.
9. The Timber Team. TIMBER: A Native XML Database. http://www.eecs.umich.edu/db/timber, first release: 2004.
10. J. M. Patel, J.-B. Yu, N. Kabra, K. Tufte, B. Nag, J. Burger, N. E. Hall, K. Ramasamy, R. Lueder,

C. Ellman, J. Kupsch, S. Guo, D. J. DeWitt, and J. F. Naughton. Paradise: A Parallel DBMS for GIS Applications. Developed while a graduate student at the University of Wisconsin. Purchased by NCR Corp. in 1997.

Patents

1. J. M. Patel, N. S. Potti and R. J. L. John. “Conversational programming interface.” US Patent US9959868B1. Issued May 1, 2018.
2. N. S. Potti and J. M. Patel. “Database acceleration through runtime code generation.” US Patent 15,574,103. Issued December 6, 2018.
3. S. R. Kulkarni, N. Bhagat, M. Fu, V. Kedigehalli, C. Kellogg, S. Mittal, J. M. Patel, K. Ramasamy, and S. Taneja. “Stream processing at scale.” US Patent 15,069,893. Issued October 9, 2018.
4. J. M. Patel, N. S. Potti, and R. J. L. John. “Conversational programming interface.” US Patent 9,959,868. Issued May 5, 2018.
5. Y. Li and J. M. Patel. “Database system with highly denormalized database structure.” U. S. Patent US9870401B2. Issued January 16, 2018.
6. Y. Li and J. M. Patel. “Database system with data organization providing improved bit parallel processing.” U. S. Patent 9,002,903. Issued April 7, 2015.
7. J. M. Patel, Y. Chen, and K. Ramasamy. “Location tracking framework.” U. S. Patent 8,447,320. Issued May 21, 2013.
8. J. M. Patel, Y. Chen, and K. Ramasamy. “Location tracking optimizations.” U. S. Patent 8,155,880. Issued April 10, 2012.
9. N. Koudas, D. Srivastava, J. M. Patel, S. Al-Khalifa, H. V. Jagadish, and Y. Wu. “Method of pattern searching.” U. S. Patent 7,451,144. Issued November 11, 2008. Also, related patents: 8,015,179 and 8,117,190.
10. N. Kabra, J. M. Patel, J-B. Yu, B. Nag, and J-J. Chen. “Method and Apparatus for Parallel Execution of Trigger Actions.” U. S. Patent 6,732,084. Issued May 4, 2004.
11. K. Ramasamy, J. M. Patel, and J. F. Naughton. “Set Containment Join Operation in an Object/Relational Database Management System.” U. S. Patent 6,728,694. Issued April 27, 2004.
12. J. M. Patel and N. Kabra. “Method and Apparatus for Evaluating Index Predicates on Complex Data Types Using Virtual Indexed Streams.” U. S. Patent 6,678,686. Issued January 13, 2004.
13. B. Nag, J-B. Yu, J. M. Patel, and Z. Wang. “Method and Apparatus for Fetching Array Based Objects by Direct Delivery and Batching.” U. S. Patent 6,601,064. Issued July 29, 2003.
14. N. Kabra, J. M. Patel, J-B. Yu, B. Nag, and J-J. Chen. “Method and Apparatus for Parallel Execution of SQL from within User Defined Functions.” U. S. Patent 6,594,651. Issued July

15, 2003.

15. N. Kabra, J. M. Patel, J-B. Yu, B. Nag, and J-J. Chen. "Method and Apparatus for Parallel Execution of SQL from Stored Procedures." U. S. Patent 6,507,834. Issued January 14, 2003.
16. A. Singh, J. M. Patel, and N. Kabra. "Method and Apparatus for Using Java as a Stored Procedure Language and as an Embedded Language on a Client." U. S. Patent 6,477,540. Issued November 5, 2002.

GRADUATE STUDENTS AND POST-DOCTORAL STUDENTS ADVISED

Graduated Ph.D. Students

1. Zuyu Zhang, October 2019, Towards Efficient Data Processing Methods for In-memory architectures. First employment: Amazon.
2. Adalbert Gerald Soosai Raj, May 2019, Effect of Bilingual Education and Live-coding on Student Learning and Engagement in Teaching and Learning Computer Science. First employment: University of California, San Diego.
3. Jianqiao Zhu, January 2019, Towards Interactive Methods for Gathering Insights from Data. First employment: Google.
4. Harshad Deshmukh, August 2018, *Efficient Query Scheduling*. First employment: Google.
5. Navneet Potti, August 2018, *Improving the Usability of Data Analytics Systems*. First employment: Google.
6. KwangHyun Park, August 2016, Data Processing Using Flash Storage: Some Opportunities and Limitations. First employment: Microsoft.
7. Arun Kumar, August 2016, *Learning Over Joins*. First employment: University of California, San Diego. (Co-advised with Jeff Naughton.)
8. Craig Chasseur, September 2015, A Bare-Metal Approach to Database Performance on Contemporary Hardware. First employment: Pivotal.
9. Yinan Li, August 2015, *Analytic Query Processing at Bare Metal Speeds*. First employment: Microsoft Research.
10. Spyros Blanas, August 2013, *High-performance main memory database management*. First employment: Assistant Professor, Ohio State University.
11. Avrilia Floratou, August 2013, *High-Performance Cloud Data Management*. First employment: IBM Almaden Research Center, San Jose, California.
12. Ning Zhang, December 2012, *Towards Cost-Effective Resources Provisioning for DBMS and Storage Clouds.* First employment: NEC Laboratories, Cupertino, CA.
13. Jae Young Do, November 2012, *The Role of Flash Memory in Database Management Systems.*

First employment: Microsoft, Redmond, WA.

14. Willis Lang, August 2012, *Cost-Effective Cloud Data Processing.* First employment: Microsoft Jim Gray Systems Lab, Madison, WI.
15. Yuanyuan Tian, August 2008, *Querying Graph Databases*. First employment: IBM Almaden Research Center, San Jose, California.
16. You Jung Kim, August 2008, *Efficient Index-based Methods for Processing Large Biological Databases*. First employment: Oracle, Redwood Shores, California.
17. Yun Chen, December 2007, Efficient Query Processing for Spatio-Temporal Databases. First employment: eBay, Mountain View, California.
18. Sandeep Tata, August 2007, *Declarative Querying for Biological Sequences*. First employment: IBM Almaden Research Center, San Jose, California.
19. Michael Morse (co-advised), August 2007, *Efficient Algorithms for Similarity and Skyline Summary on Multidimensional Datasets*. First employment: MITRE Corp., McLean, VA.
20. Richard A. Hankins, May 2004, *Architecture-Conscious Storage Management*. First employment: Microprocessor Research Lab, Intel Corporation, Santa Clara, California.
21. Kanda Runapongsa, July 2003, *Methods for Efficient Storage and Indexing in XML Databases*. First employment: Dept. of Computer Engineering, Khon Kaen University, Thailand.

Current Ph.D. Students

1. Dylan Bacon
2. Junda Chen
3. Yannis Chronis
4. Kevin Gaffney
5. Rogers Jeffrey Leo John
6. Aarti Kakaraparthy
7. Elena Milkai (co-advised with Xiangyao Yu)
8. Kendall Park (co-advised with David Schneider)
9. Matthew Prammer
10. Yunjia Zhang (co-advised with Theodoros Rekatsinas)

Past M.S. Students

1. Anuja Golechha, 2019, first employment: Google.
2. Robert Claus, 2018, first employment: MyMo.
3. Saket Saurabh, 2017, first employment: Nutanix.

4. Lalitha Viswanathan, 2017, first employment: Microsoft.
5. Marc Spehlmann, 2017, first employment: Autonomic Inc.
6. Qiang Zeng, 2015, first employment: Google.
7. Frank Bertsch, 2015, first employment: Innovative Software Engineering.
8. Shoban P. Chandrabose, 2015, first employment: Google.
9. Jing Fan, 2015.
10. James Paton, 2014, first employment: Facebook.
11. Anusha Dasarakothapalli, 2014, first employment: Amazon.
12. Vinitha Gankidi, 2014, first employment: LinkedIn.
13. Shriram Shridharan, 2014, first employment: Amazon
14. Ramakrishnan Kandhan, 2010, first employment: Google.
15. Nikhil Teletia, 2010, first employment: Microsoft.
16. Su Liu, 2008, first employment: Microsoft.
17. Shruti Kasetty, May 2004, first employment: Microsoft.
18. Andrew McClory, May 2003, first employment: FactSet Research.
19. Venkatasiva P. Chakkabala, August 2003, first employment: Informatica.
20. Laurie Hammel, May 2002, first employment: National Security Agency.

Postdoctoral Students

1. Wen Jin, September 2007-July 2009, first employment: Independent Consultant.
2. Ramakrishna Varadarajan**,** June 2009-July 2011, first employment: Vertica/HP.