

Jignesh M. Patel

WORK ADDRESS

5000 Forbes Ave.
Pittsburgh, PA 15213
USA

Phone: (412) 268-1453
E-mail: jignesh@cmu.edu
jigneshpatel.org

CURRENT RESEARCH INTERESTS

Data management.

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 and Engineering (with honors), May 1991.
Institute of Technology-Benaras Hindu University (IIT-Varanasi), India.

PROFESSIONAL EXPERIENCE

Aug 2023- Professor, Computer Science Department, Carnegie Mellon University.

**June 2010-
Aug 2023** Professor, Computer Sciences Department, University of Wisconsin-Madison.
Also affiliated with the Department of Biostatistics and Medical Informatics,
University of Wisconsin-Madison.
Co-founder and co-Chair of the UW Creative Destruction Lab (CDL-WI) from 2020-
2022.

**June 2020-
present** DataChat Inc., Co-founder, CEO and President.

**Sept. 2019-
present** Wisconsin Economic Development Corp. (WEDC) Entrepreneurship and Innovation
Committee, Member.

**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 Inc. (a Madison, WI healthcare integration startup).
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.
2009-present	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.
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

- Fellow, American Association for the Advancement of Science (AAAS), 2022.
- Vilas Associates Award, U. of Wisconsin, 2019.
- Fellow, Institute of Electrical and Electronic Engineers (IEEE), 2018.
- Fellow, Association for Computing Machinery (ACM), 2014.
- SACM Student Choice Professor of the Year Award, CS Department, U. Wisconsin, 2014.
- Senior Member, Institute of Electrical and Electronic Engineers (IEEE), 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 in VLDB 2012.
- Best papers in ICDE 2012.

- Distinguished Scientist, Association for Computing Machinery (ACM), 2011.
- Best papers in SIGMOD 2011.
- Best papers in ICDE 2010.
- Best paper award, DaMoN 2010.
- Best paper award, DaMoN 2009.
- 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

- "SHF: Medium: A hardware-software co-design approach for high-performance in-memory analytic data processing," National Science Foundation (NSF), CCF- 2312739, \$1,200,000. (CMU share \$400,000, other co-PIs at Cornell and UVA). Role: PI.
- "Elements: Software: Towards Efficient Embedded Data Processing," National Science Foundation (NSF), OAC- 1835446, 2019-23, \$ 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.
- Microsoft Jim Gray Research Lab Collaboration Grant, ~\$125K/year, 2012-present. Role: co-PI.
- Google/YouTube, gift donation, \$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, \$10M. 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, 2013, \$500,000. 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.
- "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.
- "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

- 2024: SIGMOD (Associate Editor).
- 2024: CIDR (PC co-chair).
- 2022: CIDR.
- 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: BIODDD, CIDR, ICDE, SIGMOD, VLDB.
- 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.

Abbreviations

BIDM	International Workshop on Biological Data Management
BIODDD	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

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: BIODDD.
- 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

- AAAS
- 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, Distinguished Lecture Series, and Invited Talks

- "Revenge of the database machines?" Distinguished Lecture Series, Columbia University, September 19, 2022.
- "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

- CMU 15-445/645, Database Systems: F23 (in progress)
- 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. Yunjia Zhang, Yannis Chronis, Jignesh M. Patel, Theodoros Rekatsinas: Simple Adaptive Query Processing vs. Learned Query Optimizers: Observations and Analysis. VLDB 2023.
2. Ling Zhang, Shaleen Deep, Avrilia Floratou, Anja Gruenheid, Jignesh M. Patel, Yiwen Zhu: Exploiting Structure in Regular Expression Queries. Proc. ACM Manag. Data 1(2): 152:1-152:28 (2023)
3. Yunjia Zhang, Avrilia Floratou, Joyce Cahoon, Subru Krishnan, Andreas C. Müller, Dalitso Banda, Fotis Psallidas, Jignesh M. Patel: Schema Matching using Pre-Trained Language Models. ICDE 2023: 1558-1571
4. Rogers Jeffrey Leo John, Dylan Bacon, Junda Chen, Ushmal Ramesh, Jiatong Li, Deepan Das, Robert V. Claus, Amos Kendall, Jignesh M. Patel: DataChat: An Intuitive and Collaborative Data Analytics Platform. SIGMOD Conference Companion 2023: 203-215
5. Helena Caminal, Yannis Chronis, Tianshu Wu, Jignesh M. Patel, José F. Martínez: Accelerating

- database analytic query workloads using an associative processor. ISCA 2022: 623-637
6. Elena Milkai, Yannis Chronis, Kevin P. Gaffney, Zhihan Guo, Jignesh M. Patel, Xiangyao Yu: How Good is My HTAP System? SIGMOD Conference 2022: 1810-1824
 7. Martin Prammer, Suryadev Sahadevan Rajesh, Junda Chen, Jignesh M. Patel: Introducing a Query Acceleration Path for Analytics in SQLite3. CIDR 2022
 8. Yunjia Zhang, Yannis Chronis, Jignesh M. Patel, Theodoros Rekatsinas: Can Transfer Learning be used to build a Query Optimizer? CIDR 2022
 9. Harshad Deshmukh, Bruhathi Sundarmurthy, Jignesh M. Patel: On inter-operator data transfers in query processing. ICDE 2022: 820-832
 10. Zubeyr F. Eryilmaz, Aarati Kakaraparthi, Jignesh M. Patel, Rathijit Sen, Kwanghyun Park: FPGA for Aggregate Processing: The Good, The Bad, and The Ugly. ICDE 2021: 1044-1055
 11. Adalbert Gerald Soosai Raj, Pan Gu, Eda Zhang, Arockia Xavier Annie Rayan, Jim Williams, Richard Halverson, Jignesh M. Patel: Live-coding vs Static Code Examples: Which is better with respect to Student Learning and Cognitive Load? ACE 2020: 152-159
 12. Zuyu Zhang, Harshad Deshmukh, Jignesh M. Patel: Data Partitioning for In-Memory Systems: Myths, Challenges, and Opportunities. CIDR 2019
 13. Adalbert Gerald Soosai Raj, Eda Zhang, Saswati Mukherjee, Jim Williams, Richard Halverson, Jignesh M. Patel: Effect of Native Language on Student Learning and Classroom Interaction in an Operating Systems Course. ITiCSE 2019: 499-505
 14. Adalbert Gerald Soosai Raj, Hanqi Zhang, Viren Abhyankar, Saswati Mukherjee, Eda Zhang, Jim Williams, Richard Halverson, Jignesh M. Patel: Impact of Bilingual CS Education on Student Learning and Engagement in a Data Structures Course. Koli Calling 2019: 24:1-24:10
 15. Peter Van Sandt, Yannis Chronis, Jignesh M. Patel: Efficiently Searching In-Memory Sorted Arrays: Revenge of the Interpolation Search? SIGMOD Conference 2019: 36-53
 16. Fengan Li, Lingjiao Chen, Yijing Zeng, Arun Kumar, Xi Wu, Jeffrey F. Naughton, Jignesh M. Patel: Tuple-oriented Compression for Large-scale Mini-batch Stochastic Gradient Descent. SIGMOD Conference 2019: 1517-1534
 17. Adalbert Gerald Soosai Raj, Varun Naik, Jignesh M. Patel, Richard Halverson: How to teach "modern C++" to someone who already knows programming? ACE 2018: 97-104
 18. Adalbert Gerald Soosai Raj, Jignesh M. Patel, Richard Halverson, Erica Rosenfeld Halverson: Role of Live-coding in Learning Introductory Programming. Koli Calling 2018: 13:1-13:8
 19. Adalbert Gerald Soosai Raj, Jignesh M. Patel, Richard Halverson: Is More Active Always Better for Teaching Introductory Programming? LaTICE 2018: 103-109
 20. Rogers Jeffrey Leo John, Jignesh M. Patel, Andrew L. Alexander, Vikas Singh, Nagesh Adluru: A Natural Language Interface for Dissemination of Reproducible Biomedical Data Science. MICCAI (4) 2018: 197-205
 21. Adalbert Gerald Soosai Raj, Kasama Ketsuriyonk, Jignesh M. Patel, Richard Halverson: Does Native Language Play a Role in Learning a Programming Language? SIGCSE 2018: 417-422
 22. Harshad Deshmukh, Hakan Memisoglu, Jignesh M. Patel: Adaptive Concurrent Query Execution Framework for an Analytical In-Memory Database System. BigData Congress 2017: 23-30
 23. Rogers Jeffrey Leo John, Navneet Potti, Jignesh M. Patel: Ava: From Data to Insights Through Conversations. CIDR 2017
 24. Lalitha Viswanathan, Bikash Chandra, Willis Lang, Karthik Ramachandra, Jignesh M. Patel, Ajay

- Kalhan, David J. DeWitt, Alan Halverson: Predictive Provisioning: Efficiently Anticipating Usage in Azure SQL Database. ICDE 2017: 1111-1116
25. Adalbert Gerald Soosai Raj, Kasama Ketsuriyonk, Jignesh M. Patel, Richard Halverson: What Do Students Feel about Learning Programming Using Both English and Their Native Language? LaTiCE 2017: 1-8
 26. Arun Kumar, Jeffrey F. Naughton, Jignesh M. Patel, Xiaojin Zhu: To Join or Not to Join?: Thinking Twice about Joins before Feature Selection. SIGMOD Conference 2016: 19-34
 27. Avriilia Floratou, Jignesh M. Patel: Replica Placement in Multi-tenant Database Environments. BigData Congress 2015: 246-253 (Best paper award.)
 28. Jing Fan, Adalbert Gerald Soosai Raj, Jignesh M. Patel: The Case Against Specialized Graph Analytics Engines. CIDR 2015
 29. Sanjeev Kulkarni, Nikunj Bhagat, Maosong Fu, Vikas Kedigehalli, Christopher Kellogg, Sailesh Mittal, Jignesh M. Patel, Karthik Ramasamy, Siddarth Taneja: Twitter Heron: Stream Processing at Scale. SIGMOD Conference 2015: 239-250
 30. Yinan Li, Craig Chasseur, Jignesh M. Patel: A Padded Encoding Scheme to Accelerate Scans by Leveraging Skew. SIGMOD Conference 2015: 1509-1524
 31. Arun Kumar, Jeffrey F. Naughton, Jignesh M. Patel: Learning Generalized Linear Models Over Normalized Data. SIGMOD Conference 2015: 1969-1984
 32. Ankit Toshniwal, Siddarth Taneja, Amit Shukla, Karthikeyan Ramasamy, Jignesh M. Patel, Sanjeev Kulkarni, Jason Jackson, Krishna Gade, Maosong Fu, Jake Donham, Nikunj Bhagat, Sailesh Mittal, Dmitry V. Ryaboy: Storm@twitter. SIGMOD Conference 2014: 147-156
 33. Ning Zhang, Jun'ichi Tatemura, Jignesh M. Patel, Hakan Hacigümüs: Re-evaluating designs for multi-tenant OLTP workloads on SSD-based I/O subsystems. SIGMOD Conference 2014: 1383-1394
 34. Spyros Blanas, Jignesh M. Patel: Memory footprint matters: efficient equi-join algorithms for main memory data processing. SoCC 2013: 19:1-19:16
 35. Jaeyoung Do, Donghui Zhang, Jignesh M. Patel, David J. DeWitt: Fast peak-to-peak behavior with SSD buffer pool. ICDE 2013: 1129-1140
 36. Yinan Li, Jignesh M. Patel: BitWeaving: fast scans for main memory data processing. SIGMOD Conference 2013: 289-300
 37. Jaeyoung Do, Yang-Suk Kee, Jignesh M. Patel, Chanik Park, Kwanghyun Park, David J. DeWitt: Query processing on smart SSDs: opportunities and challenges. SIGMOD Conference 2013: 1221-1230
 38. Willis Lang, Srinath Shankar, Jignesh M. Patel, Ajay Kalhan: Towards Multi-tenant Performance SLOs. ICDE 2012: 702-713 (Selected as the best papers in the conference, ICDE 2012.)
 39. Spyros Blanas, Yinan Li, Jignesh M. Patel: Design and evaluation of main memory hash join algorithms for multi-core CPUs. SIGMOD Conference 2011: 37-48
 40. Yinan Li, Allison Terrell, Jignesh M. Patel: WHAM: a high-throughput sequence alignment method. SIGMOD Conference 2011: 445-456 (Selected as the best papers in the conference, SIGMOD 2011.)
 41. Wen Jin, Jignesh M. Patel: Efficient and generic evaluation of ranked queries. SIGMOD Conference 2011: 601-612
 42. Jaeyoung Do, Donghui Zhang, Jignesh M. Patel, David J. DeWitt, Jeffrey F. Naughton, Alan

- Halverson: Turbocharging DBMS buffer pool using SSDs. SIGMOD Conference 2011: 1113-1124
43. Wen Jin, Michael D. Morse, Jignesh M. Patel, Martin Ester, Zengjian Hu: Evaluating skylines in the presence of equijoins. ICDE 2010: 249-260
 44. Avriilia Floratou, Sandeep Tata, Jignesh M. Patel: Efficient and accurate discovery of patterns in sequence datasets. ICDE 2010: 461-472 (Selected as the best papers in the conference, ICDE 2010.)
 45. Ning Zhang, Yuanyuan Tian, Jignesh M. Patel: Discovery-driven graph summarization. ICDE 2010: 880-891
 46. Wen Jin, Kristen LeFevre, Jignesh M. Patel: An online framework for publishing privacy-sensitive location traces. MobiDE 2010: 1-8
 47. Spyros Blanas, Jignesh M. Patel, Vuk Ercegovic, Jun Rao, Eugene J. Shekita, Yuanyuan Tian: A comparison of join algorithms for log processing in MapReduce. SIGMOD Conference 2010: 975-986
 48. Willis Lang, Jignesh M. Patel: Towards Eco-friendly Database Management Systems. CIDR 2009
 49. Yun Chen, Jignesh M. Patel: Design and evaluation of trajectory join algorithms. GIS 2009: 266-275
 50. Yuanyuan Tian, Jignesh M. Patel: TALE: A Tool for Approximate Large Graph Matching. ICDE 2008: 963-972
 51. Mark A. Iwen, Willis Lang, Jignesh M. Patel: Scalable Rule-Based Gene Expression Data Classification. ICDE 2008: 1062-1071
 52. Sandeep Tata, Jignesh M. Patel: FLAME: Shedding Light on Hidden Frequent Patterns in Sequence Datasets. ICDE 2008: 1343-1345
 53. Yuanyuan Tian, Richard A. Hankins, Jignesh M. Patel: Efficient aggregation for graph summarization. SIGMOD Conference 2008: 567-580
 54. You Jung Kim, Jignesh M. Patel: Rethinking Choices for Multi-dimensional Point Indexing: Making the Case for the Often Ignored Quadtree. CIDR 2007: 281-291
 55. Yun Chen, Jignesh M. Patel: Efficient Evaluation of All-Nearest-Neighbor Queries. ICDE 2007: 1056-1065
 56. Stelios Pappas, Jignesh M. Patel, H. V. Jagadish: SIGOPT: Using Schema to Optimize XML Query Processing. ICDE 2007: 1456-1460
 57. Michael D. Morse, Jignesh M. Patel: An efficient and accurate method for evaluating time series similarity. SIGMOD Conference 2007: 569-580
 58. Michael D. Morse, Jignesh M. Patel, H. V. Jagadish: Efficient Skyline Computation over Low-Cardinality Domains. VLDB 2007: 267-278
 59. Sandeep Tata, Willis Lang, Jignesh M. Patel: Periscope/SQ: Interactive Exploration of Biological Sequence Databases. VLDB 2007: 1406-1409
 60. Sandeep Tata, Jignesh M. Patel, James S. Friedman, Anand Swaroop: Declarative Querying for Biological Sequences. ICDE 2006: 87
 61. Michael D. Morse, Jignesh M. Patel, William I. Grosky: Efficient Continuous Skyline Computation. ICDE 2006: 108
 62. Sushant Sinha, Farnam Jahanian, Jignesh M. Patel: WIND: Workload-Aware INtrusion Detection. RAID 2006: 290-310

63. Jignesh M. Patel, Yun Chen, V. Prasad Chakka: STRIPES: An Efficient Index for Predicted Trajectories. SIGMOD Conference 2004: 637-646
64. Sandeep Tata, Richard A. Hankins, Jignesh M. Patel: Practical Suffix Tree Construction. VLDB 2004: 36-47
65. V. Prasad Chakka, Adam Everspaugh, Jignesh M. Patel: Indexing Large Trajectory Data Sets With SETI. CIDR 2003
66. Yuqing Wu, Jignesh M. Patel, H. V. Jagadish: Structural Join Order Selection for XML Query Optimization. ICDE 2003: 443-454
67. Richard A. Hankins, Jignesh M. Patel: Effect of node size on the performance of cache-conscious B+-trees. SIGMETRICS 2003: 283-294
68. Stelios Paparizos, Shurug Al-Khalifa, Adriane Chapman, H. V. Jagadish, Laks V. S. Lakshmanan, Andrew Nierman, Jignesh M. Patel, Divesh Srivastava, Nuwee Wiwatwattana, Yuqing Wu, Cong Yu: TIMBER: A Native System for Querying XML. SIGMOD Conference 2003: 672
69. Sandeep Tata, Jignesh M. Patel: PiQA: An Algebra for Querying Protein Data Sets. SSDBM 2003: 141-150
70. Richard A. Hankins, Jignesh M. Patel: Data Morphing: An Adaptive, Cache-Conscious Storage Technique. VLDB 2003: 417-428
71. Colin Meek, Jignesh M. Patel, Shruti Kasetty: OASIS: An Online and Accurate Technique for Local-alignment Searches on Biological Sequences. VLDB 2003: 910-921
72. Yuqing Wu, Jignesh M. Patel, H. V. Jagadish: Estimating Answer Sizes for XML Queries. EDBT 2002: 590-608
73. Shurug Al-Khalifa, H. V. Jagadish, Jignesh M. Patel, Yuqing Wu, Nick Koudas, Divesh Srivastava: Structural Joins: A Primitive for Efficient XML Query Pattern Matching. ICDE 2002: 141-152
74. Laurie Hammel, Jignesh M. Patel: Searching on the Secondary Structure of Protein Sequences. VLDB 2002: 634-645
75. Murali Annavaram, Jignesh M. Patel, Edward S. Davidson: Call Graph Prefetching for Database Applications. HPCA 2001: 281-290
76. Murali Annavaram, Jignesh M. Patel, Edward S. Davidson: Data prefetching by dependence graph precomputation. ISCA 2001: 52-61
77. Jignesh M. Patel, David J. DeWitt: Clone join and shadow join: two parallel spatial join algorithms. ACM-GIS 2000: 54-61
78. Karthikeyan Ramasamy, Jignesh M. Patel, Jeffrey F. Naughton, Raghav Kaushik: Set Containment Joins: The Good, The Bad and The Ugly. VLDB 2000: 351-362
79. Jignesh M. Patel, Jie-Bing Yu, Navin Kabra, Kristin Tufte, Biswadeep Nag, Josef Burger, Nancy E. Hall, Karthikeyan Ramasamy, Roger Lueder, Curt J. Ellmann, Jim Kupsch, Shelly Guo, David J. DeWitt, Jeffrey F. Naughton: Building a Scaleable Geo-Spatial DBMS: Technology, Implementation, and Evaluation. SIGMOD Conference 1997: 336-347
80. Jignesh M. Patel, David J. DeWitt: Partition Based Spatial-Merge Join. SIGMOD Conference 1996: 259-270
81. Jignesh M. Patel, Michael J. Carey, Mary K. Vernon: Accurate Modeling of the Hybrid Hash Join Algorithm. SIGMETRICS 1994: 56-66
82. David J. DeWitt, Navin Kabra, Jun Luo, Jignesh M. Patel, Jie-Bing Yu: Client-Server Paradise. VLDB 1994: 558-569

Journal Papers

1. Daniel Abadi, Anastasia Ailamaki, David G. Andersen, Peter Bailis, Magdalena Balazinska, Philip A. Bernstein, Peter A. Boncz, Surajit Chaudhuri, Alvin Cheung, AnHai Doan, Luna Dong, Michael J. Franklin, Juliana Freire, Alon Y. Halevy, Joseph M. Hellerstein, Stratos Idreos, Donald Kossmann, Tim Kraska, Saitesh Krishnamurthy, Volker Markl, Sergey Melnik, Tova Milo, C. Mohan, Thomas Neumann, Beng Chin Ooi, Fatma Ozcan, Jignesh M. Patel, Andrew Pavlo, Raluca A. Popa, Raghu Ramakrishnan, Christopher Ré, Michael Stonebraker, Dan Suciu: The Seattle report on database research. *Commun. ACM* 65(8): 72-79 (2022)
2. Aarati Kakaraparthi, Jignesh M. Patel, Brian Kroth, Kwanghyun Park: VIP Hashing - Adapting to Skew in Popularity of Data on the Fly. *Proc. VLDB Endow.* 15(10): 1978-1990 (2022)
3. Kevin P. Gaffney, Martin Prammer, Laurence C. Brasfield, D. Richard Hipp, Dan R. Kennedy, Jignesh M. Patel: SQLite: Past, Present, and Future. *Proc. VLDB Endow.* 15(12): 3535-3547 (2022)
4. Ashish Gehani, Raza Ahmad, Hassaan Irshad, Jianqiao Zhu, Jignesh M. Patel: Digging into big provenance (with SPADE). *Commun. ACM* 64(12): 48-56 (2021)
5. Tim Kraska, Umar Farooq Minhas, Thomas Neumann, Olga Papaemmanouil, Jignesh M. Patel, Christopher Ré, Michael Stonebraker: ML-In-Databases: Assessment and Prognosis. *IEEE Data Eng. Bull.* 44(1): 3-10 (2021)
6. Kevin P. Gaffney, Robert K. Claus, Jignesh M. Patel: Database Isolation by Scheduling. *Proc. VLDB Endow.* 14(9): 1467-1480 (2021)
7. Ashish Gehani, Raza Ahmad, Hassan Irshad, Jianqiao Zhu, Jignesh M. Patel: Digging into Big Provenance (with SPADE): A user interface for querying provenance. *ACM Queue* 19(3): 77-106 (2021)
8. Hassaan Irshad, Gabriela F. Ciocarlie, Ashish Gehani, Vinod Yegneswaran, Kyu Hyung Lee, Jignesh M. Patel, Somesh Jha, Yonghwi Kwon, Dongyan Xu, Xiangyu Zhang: TRACE: Enterprise-Wide Provenance Tracking for Real-Time APT Detection. *IEEE Trans. Inf. Forensics Secur.* 16: 4363-4376 (2021)
9. Zhiwei Fan, Jianqiao Zhu, Zuyu Zhang, Aws Albarghouthi, Paraschos Koutris, Jignesh M. Patel: Scaling-Up In-Memory Datalog Processing: Observations and Techniques. *Proc. VLDB Endow.* 12(6): 695-708 (2019)
10. Aarati Kakaraparthi, Jignesh M. Patel, Kwanghyun Park, Brian Kroth: Optimizing Databases by Learning Hidden Parameters of Solid State Drives. *Proc. VLDB Endow.* 13(4): 519-532 (2019)
11. Jignesh M. Patel, Harshad Deshmukh, Jianqiao Zhu, Navneet Potti, Zuyu Zhang, Marc Spehlmann, Hakan Memisoglu, Saket Saurabh: Quickstep: A Data Platform Based on the Scaling-Up Approach. *Proc. VLDB Endow.* 11(6): 663-676 (2018)
12. Jianqiao Zhu, Navneet Potti, Saket Saurabh, Jignesh M. Patel: Looking Ahead Makes Query Plans Robust. *Proc. VLDB Endow.* 10(8): 889-900 (2017)
13. Lingjiao Chen, Arun Kumar, Jeffrey F. Naughton, Jignesh M. Patel: Towards Linear Algebra over Normalized Data. *Proc. VLDB Endow.* 10(11): 1214-1225 (2017)
14. Jignesh M. Patel: Operational NoSQL Systems: What's New and What's Next? *Computer* 49(4): 23-30 (2016)
15. Navneet Potti, Jignesh M. Patel: DAQ: A New Paradigm for Approximate Query Processing. *Proc. VLDB Endow.* 8(9): 898-909 (2015)
16. Arun Kumar, Mona Jalal, Boqun Yan, Jeffrey F. Naughton, Jignesh M. Patel: Demonstration of

- Santoku: Optimizing Machine Learning over Normalized Data. *Proc. VLDB Endow.* 8(12): 1864-1867 (2015)
17. Jason Power, Yinan Li, Mark D. Hill, Jignesh M. Patel, David A. Wood: Implications of Emerging 3D GPU Architecture on the Scan Primitive. *SIGMOD Rec.* 44(1): 18-23 (2015)
 18. Arun Kumar, Robert McCann, Jeffrey F. Naughton, Jignesh M. Patel: Model Selection Management Systems: The Next Frontier of Advanced Analytics. *SIGMOD Rec.* 44(4): 17-22 (2015)
 19. H. V. Jagadish, Johannes Gehrke, Alexandros Labrinidis, Yannis Papakonstantinou, Jignesh M. Patel, Raghu Ramakrishnan, Cyrus Shahabi: Big data and its technical challenges. *Commun. ACM* 57(7): 86-94 (2014)
 20. Kwanghyun Park, Yang-Suk Kee, Jignesh M. Patel, Jaeyoung Do, Chanik Park, David J. DeWitt: Query Processing on Smart SSDs. *IEEE Data Eng. Bull.* 37(2): 19-26 (2014)
 21. Avrielia Floratou, Frank Bertsch, Jignesh M. Patel, Georgios Laskaris: Towards Building Wind Tunnels for Data Center Design. *Proc. VLDB Endow.* 7(9): 781-784 (2014)
 22. Yinan Li, Jignesh M. Patel: WideTable: An Accelerator for Analytical Data Processing. *Proc. VLDB Endow.* 7(10): 907-918 (2014)
 23. Vinitha Reddy Gankidi, Nikhil Teletia, Jignesh M. Patel, Alan Halverson, David J. DeWitt: Indexing HDFS Data in PDW: Splitting the data from the index. *Proc. VLDB Endow.* 7(13): 1520-1528 (2014)
 24. Shriram Sridharan, Jignesh M. Patel: Profiling R on a Contemporary Processor. *Proc. VLDB Endow.* 8(2): 173-184 (2014)
 25. Qiang Zeng, Jignesh M. Patel, David Page: QuickFOIL: Scalable Inductive Logic Programming. *Proc. VLDB Endow.* 8(3): 197-208 (2014)
 26. Willis Lang, Srinath Shankar, Jignesh M. Patel, Ajay Kalhan: Towards Multi-Tenant Performance SLOs. *IEEE Trans. Knowl. Data Eng.* 26(6): 1447-1463 (2014)
 27. Ning Zhang, Jun'ichi Tatemura, Jignesh M. Patel, Hakan Hacigümüs: Toward cost-effective storage provisioning for DBMSs. *VLDB J.* 23(2): 329-354 (2014)
 28. Craig Chasseur, Jignesh M. Patel: Design and Evaluation of Storage Organizations for Read-Optimized Main Memory Databases. *Proc. VLDB Endow.* 6(13): 1474-1485 (2013)
 29. J. B. Hodgins, 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.
 30. Willis Lang, Stavros Harizopoulos, Jignesh M. Patel, Mehul A. Shah, Dimitris Tsirogiannis: Towards Energy-Efficient Database Cluster Design. *Proc. VLDB Endow.* 5(11): 1684-1695 (2012)
 31. Avrielia Floratou, Nikhil Teletia, David J. DeWitt, Jignesh M. Patel, Donghui Zhang: Can the Elephants Handle the NoSQL Onslaught? *Proc. VLDB Endow.* 5(12): 1712-1723 (2012)
 32. Yinan Li, Jignesh M. Patel, Allison Terrell: WHAM: A High-Throughput Sequence Alignment Method. *ACM Trans. Database Syst.* 37(4): 28:1-28:39 (2012)
 33. Willis Lang, Ramakrishnan Kandhan, Jignesh M. Patel: Rethinking Query Processing for Energy Efficiency: Slowing Down to Win the Race. *IEEE Data Eng. Bull.* 34(1): 12-23 (2011)
 34. Avrielia Floratou, Jignesh M. Patel, Eugene J. Shekita, Sandeep Tata: Column-Oriented Storage Techniques for MapReduce. *Proc. VLDB Endow.* 4(7): 419-429 (2011)
 35. Ning Zhang, Jun'ichi Tatemura, Jignesh M. Patel, Hakan Hacigümüs: Towards Cost-Effective Storage Provisioning for DBMSs. *Proc. VLDB Endow.* 5(4): 274-285 (2011) (Selected as best

papers in the conference, VLDB 2012.)

36. Per-Åke Larson, Spyros Blanas, Cristian Diaconu, Craig Freedman, Jignesh M. Patel, Mike Zwilling: High-Performance Concurrency Control Mechanisms for Main-Memory Databases. *Proc. VLDB Endow.* 5(4): 298-309 (2011)
37. Avrielia Floratou, Sandeep Tata, Jignesh M. Patel: Efficient and Accurate Discovery of Patterns in Sequence Data Sets. *IEEE Trans. Knowl. Data Eng.* 23(8): 1154-1168 (2011)
38. Willis Lang, Jignesh M. Patel: Energy Management for MapReduce Clusters. *Proc. VLDB Endow.* 3(1): 129-139 (2010)
39. Ramakrishnan Kandhan, Nikhil Teletia, Jignesh M. Patel: SigMatch: Fast and Scalable Multi-Pattern Matching. *Proc. VLDB Endow.* 3(1): 1173-1184 (2010)
40. You Jung Kim, Jignesh 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)
41. Willis Lang, Michael D. Morse, Jignesh M. Patel: Dictionary-Based Compression for Long Time-Series Similarity. *IEEE Trans. Knowl. Data Eng.* 22(11): 1609-1622 (2010)
42. You Jung Kim, Nikhil Teletia, Victor Ruotti, Christopher A. Maher, Arul M. Chinnaiyan, Ron M. Stewart, James A. Thomson, Jignesh M. Patel: ProbeMatch: rapid alignment of oligonucleotides to genome allowing both gaps and mismatches. *Bioinform.* 25(11): 1424-1425 (2009)
43. V. Glenn Tarcea, Terry E. Weymouth, Alexander S. Ade, Aaron V. Bookvich, Jing Gao, Vasudeva Mahavisno, Zach Wright, Adriane Chapman, Magesh Jayapandian, Arzucan Özgür, Yuanyuan Tian, James D. Cavalcoli, Barbara Mirel, Jignesh M. Patel, Dragomir R. Radev, Brian D. Athey, David J. States, H. V. Jagadish: Michigan molecular interactions r2: from interacting proteins to pathways. *Nucleic Acids Res.* 37(Database-Issue): 642-646 (2009)
44. Philippe Cudré-Mauroux, Hideaki Kimura, Kian-Tat Lim, Jennie Rogers, Roman Simakov, Emad Soroush, Pavel Velikhov, Daniel L. Wang, Magdalena Balazinska, Jacek Becla, David J. DeWitt, Bobbi Heath, David Maier, Samuel Madden, Jignesh M. Patel, Michael Stonebraker, Stanley B. Zdonik: A Demonstration of SciDB: A Science-Oriented DBMS. *Proc. VLDB Endow.* 2(2): 1534-1537 (2009)
45. Willis Lang, Jignesh M. Patel, Jeffrey F. Naughton: On energy management, load balancing and replication. *SIGMOD Rec.* 38(4): 35-42 (2009)
46. Yuanyuan Tian, Jignesh M. Patel, Viji Nair, Sebastian Martini, Matthias Kretzler: Periscope/GQ: a graph querying toolkit. *Proc. VLDB Endow.* 1(2): 1404-1407 (2008)
47. 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.
48. Yuanyuan Tian, Richard C. McEachin, Carlos Santos, David J. States, Jignesh M. Patel: SAGA: a subgraph matching tool for biological graphs. *Bioinform.* 23(2): 232-239 (2007)
49. Michael D. Morse, Jignesh M. Patel, William I. Grosky: Efficient evaluation of radial queries using the target tree. *Int. J. Bioinform. Res. Appl.* 3(1): 24-41 (2007)
50. Michael D. Morse, Jignesh M. Patel, William I. Grosky: Efficient continuous skyline computation. *Inf. Sci.* 177(17): 3411-3437 (2007)
51. Sandeep Tata, Jignesh M. Patel: Estimating the selectivity of tf-idf based cosine similarity predicates. *SIGMOD Rec.* 36(2): 7-12 (2007)
52. You Jung Kim, Jignesh M. Patel: A framework for protein structure classification and identification

- of novel protein structures. *BMC Bioinform.* 7: 456 (2006)
53. Kanda Runapongsa, Jignesh M. Patel, H. V. Jagadish, Yun Chen, Shurug Al-Khalifa: The Michigan benchmark: towards XML query performance diagnostics. *Inf. Syst.* 31(2): 73-97 (2006)
 54. Yuanyuan Tian, Sandeep Tata, Richard A. Hankins, Jignesh M. Patel: Practical methods for constructing suffix trees. *VLDB J.* 14(3): 281-299 (2005)
 55. 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.
 56. Yuqing Wu, Jignesh M. Patel, H. V. Jagadish: Using histograms to estimate answer sizes for XML queries. *Inf. Syst.* 28(1-2): 33-59 (2003)
 57. Jignesh M. Patel: The Role of Declarative Querying in Bioinformatics. *OMICS* 7(1): 89-92 (2003)
 58. Murali Annavaram, Jignesh M. Patel, Edward S. Davidson: Call graph prefetching for database applications. *ACM Trans. Comput. Syst.* 21(4): 412-444 (2003)
 59. H. V. Jagadish, Shurug Al-Khalifa, Adriane Chapman, Laks V. S. Lakshmanan, Andrew Nierman, Stelios Paparizos, Jignesh M. Patel, Divesh Srivastava, Nuwee Wiwatwattana, Yuqing Wu, Cong Yu: TIMBER: A native XML database. *VLDB J.* 11(4): 274-291 (2002)

Workshop Papers

1. Kwanghyun Park, Jaeyoung Do, Nikhil Teletia, Jignesh M. Patel: Aggressive buffer pool warm-up after restart in SQL Server. *ICDE Workshops 2016*: 31-38
2. Jason Power, Yinan Li, Mark D. Hill, Jignesh M. Patel, David A. Wood: Toward GPUs being mainstream in analytic processing: An initial argument using simple scan-aggregate queries. *DaMoN 2015*: 11:1-11:8
3. Jignesh M. Patel: Rethinking Benchmarking for Data. *TPCTC 2015*: 130-134
4. Craig Chasseur, Yinan Li, Jignesh M. Patel: Enabling JSON Document Stores in Relational Systems. *WebDB 2013*: 1-6
5. Avriela Floratou, Jignesh M. Patel, Willis Lang, Alan Halverson: When Free Is Not Really Free: What Does It Cost to Run a Database Workload in the Cloud? *TPCTC 2011*: 163-179
6. Willis Lang, Jignesh M. Patel, Srinath Shankar: Wimpy node clusters: what about non-wimpy workloads? *DaMoN 2010*: 47-55 (Best paper award.)
7. Jaeyoung Do, Jignesh M. Patel: Join processing for flash SSDs: remembering past lessons. *DaMoN 2009*: 1-8 (Best paper award.)
8. Michael D. Morse, Jignesh M. Patel, William I. Grosky: Efficient Evaluation of Radial Queries using the Target Tree. *ICDE Workshops 2005*: 1168
9. Kanda Runapongsa, Jignesh M. Patel, Rajesh Bordawekar, Sriram Padmanabhan: XIST: An XML Index Selection Tool. *XSym 2004*: 219-234
10. 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.
11. Kanda Runapongsa, Jignesh M. Patel: Storing and Querying XML Data in Object-Relational DBMSs. *EDBT Workshops 2002*: 266-285
12. Kanda Runapongsa, Jignesh M. Patel, H. V. Jagadish, Shurug Al-Khalifa: The Michigan Benchmark: A Microbenchmark for XML Query Processing Systems. *EEXTT 2002*: 160-161

13. 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.
14. 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.

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. Yannis Chronis, June 2022, *Data Processing in the Modern Hardware Landscape*. First employment: Google.
2. Rogers Jeffrey Leo John, May 2021, *From Data to Insights Through Conversations*. First

- employment: DataChat (as co-founder).
3. Zuyu Zhang, October 2019, *Towards Efficient Data Processing Methods for In-memory architectures*. First employment: Amazon.
 4. 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.
 5. Jianqiao Zhu, January 2019, *Towards Interactive Methods for Gathering Insights from Data*. First employment: Google.
 6. Harshad Deshmukh, August 2018, *Efficient Query Scheduling*. First employment: Google.
 7. Navneet Potti, August 2018, *Improving the Usability of Data Analytics Systems*. First employment: Google.
 8. KwangHyun Park, August 2016, *Data Processing Using Flash Storage: Some Opportunities and Limitations*. First employment: Microsoft.
 9. Arun Kumar, August 2016, *Learning Over Joins*. First employment: University of California, San Diego. (Co-advised with Jeff Naughton.)
 10. Craig Chasseur, September 2015, *A Bare-Metal Approach to Database Performance on Contemporary Hardware*. First employment: Pivotal.
 11. Yinan Li, August 2015, *Analytic Query Processing at Bare Metal Speeds*. First employment: Microsoft Research.
 12. Spyros Blanas, August 2013, *High-performance main memory database management*. First employment: Ohio State University.
 13. Avrilia Floratou, August 2013, *High-Performance Cloud Data Management*. First employment: IBM Almaden Research Center.
 14. Ning Zhang, December 2012, *Towards Cost-Effective Resources Provisioning for DBMS and Storage Clouds*. First employment: NEC Laboratories.
 15. Jae Young Do, November 2012, *The Role of Flash Memory in Database Management Systems*. First employment: Microsoft.
 16. Willis Lang, August 2012, *Cost-Effective Cloud Data Processing*. First employment: Microsoft Gray Systems Lab (GSL).
 17. Yuanyuan Tian, August 2008, *Querying Graph Databases*. First employment: IBM Almaden Research Center.
 18. You Jung Kim, August 2008, *Efficient Index-based Methods for Processing Large Biological Databases*. First employment: Oracle.
 19. Yun Chen, December 2007, *Efficient Query Processing for Spatio-Temporal Databases*. First employment: eBay.
 20. Sandeep Tata, August 2007, *Declarative Querying for Biological Sequences*. First employment: IBM Almaden Research Center.
 21. Michael Morse (co-advised), August 2007, *Efficient Algorithms for Similarity and Skyline Summary on Multidimensional Datasets*. First employment: MITRE Corp.
 22. Richard A. Hankins, May 2004, *Architecture-Conscious Storage Management*. First employment: Intel Microprocessor Research Lab.
 23. Kanda Runapongsa, July 2003, *Methods for Efficient Storage and Indexing in XML Databases*. First

employment: Khon Kaen University, Thailand.

Current Ph.D. Students

1. Kevin Gaffney.
2. Aarti Kakaraparthi.
3. Elena Milkai (co-advised with Xiangyao Yu).
4. Matthew Prammer.
5. Ling Zhang.
6. Yunjia Zhang (co-advised with Theodoros Rekatsinas).

Past M.S. Students

1. Yannis Chronis, 2022; first employment: Google.
2. Dylan Bacon, 2022; first employment: DataChat.
3. Junda Chen, 2021; first employment: DataChat.
4. Anuja Golechha, 2019; first employment: Google.
5. Robert Claus, 2018; first employment: MyMo.
6. Saket Saurabh, 2017; first employment: Nutanix.
7. Lalitha Viswanathan, 2017; first employment: Microsoft.
8. Marc Spehlmann, 2017; first employment: Autonomic Inc.
9. Qiang Zeng, 2015; first employment: Google.
10. Frank Bertsch, 2015; first employment: Innovative Software Engineering.
11. Shoban P. Chandrabose, 2015; first employment: Google.
12. Jing Fan, 2015.
13. James Paton, 2014; first employment: Facebook.
14. Anusha Dasarathapalli, 2014; first employment: Amazon.
15. Vinitha Gankidi, 2014; first employment: LinkedIn.
16. Shriram Shridharan, 2014; first employment: Amazon
17. Ramakrishnan Kandhan, 2010; first employment: Google.
18. Nikhil Teletia, 2010; first employment: Microsoft.
19. Su Liu, 2008; first employment: Microsoft.
20. Shruti Kasetty, May 2004; first employment: Microsoft.
21. Andrew McClory, May 2003; first employment: FactSet Research.
22. Venkatasiva P. Chakkabala, August 2003; first employment: Informatica.
23. 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.