

VIJAY GADEPALLY

Curriculum Vitae

updated March 21, 2023

Lincoln Laboratory
Massachusetts Institute of Technology
☎ +1-781-981-8298
✉ vijayg@mit.edu
🌐 <https://vijayg.mit.edu>

Education

- 2013 **PhD, Electrical & Computer Engineering**, *The Ohio State University*.
Thesis: "Estimation of Driver Behavior for Autonomous Vehicle Applications"
Committee: Dr. Ashok Krishnamurthy (advisor), Dr. Umit Ozguner, Dr. Giorgio Rizzoni
- 2010 **M.Sc, Electrical & Computer Engineering**, *The Ohio State University*.
Signal and Image Processing, Parallel Computing
- 2006 **B.Tech, Electrical Engineering**, *Indian Institute of Technology (IIT), Kanpur*.
Power Electronics

Professional Experience

Massachusetts Institute of Technology

- 2021 – **Visiting Scientist**, *Connection Sciences*.
present Leading efforts related to artificial intelligence and education for national security applications
- 2013 – **Senior Member of the Technical Staff**, *Lincoln Laboratory*.
present Vijay leads numerous research efforts that span artificial intelligence, high performance computing, education and national security applications. He established and continues to lead the Supercomputing group comprised of 10+ researchers at MIT and participates in a number of multi-University collaborations. The team's research has been featured in media articles, has received numerous awards at conferences, and has been transitioned to a variety of commercial and government partners. Contributions include:
- Research lead of Lincoln Laboratory Supercomputing Center. Responsible for setting research vision, program development, recruiting and technical execution of \$4M+/year of external research funding. (2017–present)
 - Leading (as PI, Co-I or Team Lead) research efforts supported by organizations such as Department of the Air Force, Undersecretary of Defense for Research and Engineering, Department of the Navy, National Science Foundation, Intel Corporation, DARPA, and Lincoln Laboratory Internal funding (2014 – Present)
 - Developing strategic partnerships with organizations outside of MIT such as JPL, Oak Ridge, Intel, Northeastern, BU, OSU, MGHPCC, NVIDIA, Amazon, Meta and other renowned institutions (2013–present)
 - Extensive publication record in top-tier conferences and journals (2013–present)
 - Contributing to and leading open-source software projects such as BigDAWG, Computing on Masked Data, Graphulo, D4M (2013–present)
 - Helping define the national roadmap through involvement in studies on application of advanced techniques such as data architectures, AI and recommender systems for national security applications (2015–present)
 - Educating and mentoring the next generation of technical leaders from MIT and the US Government

Radium Cloud

- 2022 – **Technical Advisor**.
present Member of technical advisory board to Radium Cloud - a startup company focused on cloud computing technologies for AI and Data Science applications.

The Ohio State University

- 2006 – 2013 **Graduate Research Associate**, *Ohio Supercomputer Center*.
Research to develop mathematical models for driver behavior for autonomous vehicle applications, image processing and high performance computing systems. Accomplishments include development of new Hidden Markov Model technique to estimate driver behavior, design and execution of human-subject research, and development of an education program around cyber-physical systems. Apart from thesis research, additional work in developing high performance computing tools for signal and image processing
- 2012–2013 **President**, *Student Commercialization Board*, Technology Transfer Office.
Established first ever student commercialization board for the Ohio State University Technology Commercialization and Knowledge Transfer Office in order to improve quality of commercialized technology.
- 2011 – 2012 **President**, *Council of Graduate Students*.
Elected to represent nearly 15,000 graduate students at The Ohio State University for the nation's oldest graduate student government. As a part of this role, my job was to preside over all meetings of the Council and serve as the primary liaison between the Council, Graduate School, University Administration, and Ohio State Board of Trustees. In addition to appointing a 13 person cabinet, I managed nearly 75 elected delegates and hundreds of volunteers.

Raytheon Company

- 2007 – 2007 **Post Graduate Technical Intern**, *Surveillance and Sensor Center*.
Developed a publish-subscribe operation framework for an Internal Research and Development (IRAD) project. Innovations include a new mechanism for inter-process communication across heterogeneous hardware.

Rensselaer Polytechnic Institute

- 2005 – 2005 **Visiting Scholar**, *Electrical, Computer and Systems Engineering*.
Developed a high efficiency boost converter for DC-DC applications and prototyped a coupled inductor design for mobile micro fuel cell applications. Completed Printed Circuit Board (PCB) design of prototype.

Indian Institute of Technology, Bombay

- 2003 – 2003 **Student Intern**, *Department of Electrical Engineering*.
Programmed a simulation of mixed-signal devices using the SPICE software and contributed to the development of SEQUEL - a circuit simulation software for mixed-signal devices.

Honors & Awards

- 2022 **Outstanding Paper** - Benchmarking Resource Usage for Efficient Distributed Deep Learning, IEEE High Performance Extreme Computing Conference
- 2021 **Outstanding Student Paper** - Serving Machine Learning Inference Using Heterogeneous Hardware, IEEE High Performance Extreme Computing Conference
- 2020 **Outstanding Paper** - GPU Concurrency of Deep Residual Neural Networks Via Nonlinear Multigrid, IEEE High Performance Extreme Computing Conference
- 2019 **Outstanding Student Paper** - A Survey of Attacks and Defenses of Edge-Deployed Neural Networks, IEEE High Performance Extreme Computing Conference
- 2019 **IEEE Senior Member** - Elevated to IEEE Senior Member
- 2018 **Outstanding Paper Award** - Measuring the Impact of Spectre and Meltdown, IEEE High Performance Extreme Computing Conference
- 2017 **'40 under 40'** - Named to the inaugural 40 under 40 list of the Armed Forces Communication and Electronics Association (AFCEA)
- 2017 **R&D 100 Award (finalist)** - BigDAWG system led by Vijay Gadepally was a finalist for the R&D 100 Award
- 2017 **MIT Lincoln Laboratory Early Career Technical Achievement Award** - Given to two staff under the age of 35 annually
- 2017 **Team Award** - Presented as a part of the MIT Lincoln Laboratory Team Awards
- 2017 **Outstanding Student Paper** - Cross-Engine Query Execution in Federated Database Systems, IEEE High Performance Extreme Computing Conference
- 2017 **Outstanding Student Paper** - A Cloud-based Large-scale Brain Connectivity Analysis Using Accumulo and D4M, IEEE High Performance Extreme Computing Conference
- 2016 **Best Paper Award** - Julia Implementation of the Dynamic Distributed Dimensional Data Model, IEEE High Performance Extreme Computing Conference
- 2016 **Outstanding Student Paper** - From NoSQL Accumulo to NewSQL Graphulo: Design and Utility of Graph Algorithms inside a BigTable Database, IEEE High Performance Extreme Computing Conference
- 2016 **Best Paper (finalist)** - Benchmarking SciDB Data Import on HPC Systems,, IEEE High Performance Extreme Computing Conference
- 2015 **Best Student Paper (finalist)** - Improving Big Data Visual Analytics with Interactive Virtual Reality, IEEE High Performance Extreme Computing Conference
- 2015 **SuperComputing Early Career Program** - Selected for SuperComputing (SC) 2015 Early Career Program
- 2014 **Team Award** - Lincoln Laboratory Team Award, 2014
- 2012 **Introducing President Barack Obama** - during White House visit to The Ohio State University on energy independence
- 2012 **NASA Future Forum Panel** - Selected to serve on panel with Senator John Glenn, and NASA administrator Charles Bolden
- 2012 **Ray Travel Award** - Awarded by The Ohio State University
- 2011 **Outstanding Graduate Student Award** - Awarded by The Ohio State University

Selected Media and Press Coverage

- 2022 "Taking a magnifying glass to data center operations":
 - MIT News, August 24, 2022
 - Tech Xplore, August 24, 2022
 - MIT Lincoln Laboratory News, August 8, 2022

- 2022 "Curbing the Growing Power Needs of Machine Learning":
 - Dataconomy, June 3, 2022
 - Unite.AI, December 2022

- 2021 "Artificial Intelligence Leaders to Convene in New Lecture Series", LSU Media Center, January 29, 2021

- 2020 "New algorithm uses supercomputing to combat cyberattacks", MIT Lincoln Laboratory News, August 2020

- 2019 "Supercomputers analyze web traffic across entire internet", MIT News, October 2019

- 2019 "Supercomputers can spot cyber threats", MIT News, February 2019

- 2018 "Spectre/Meltdown fixes in HPC: Want the bad news or the bad news? It's slower, say boffins", The Register, July 26, 2018.

- 2016 "Capitalizing on machine learning – from life sciences to financial services", HPC Wire, December 26, 2016

- 2015 "Databases Unlock Big Data", Paradigm4 Use Case, January 9, 2015

- 2014 "How 3D Printing is Revolutionizing the Display of Big Data":
 - MIT Technology Review, October 8, 2014
 - The Daily Dot, October 9, 2014
 - Motherboard Magazine, October 1, 2014
 - BigData4Analytics.com, October 2014

- 2014 "A Workshop: Unlocking the Power of Medical Big Data", Featured on ISTC for Big Data Blog, 2014.

- 2013 "Drive Simulation Lab" television program, Feature in Big10 Television Network Special on The Ohio State University Driving Simulator.

- 2012 "President Obama at Ohio State" (among various news articles):
 - The Ohio State University Feature, March 2012
 - Fox 8 Cleveland News, March 2012
 - The Lantern, March 2012
 - The Columbus Dispatch, March 2012

- 2012 "Task force formed in response to hate crimes", The Lantern, April 2012

- 2012 "Student Government Combats OSU's low ranking in graduate stipends", The Lantern, February 2012.

Educational Activities

Teaching

- 2023 **Mobile Communications** (Course Organizer) - Contributor to course on 5G and Mobile Communications for professional course (30 students)
- 2022 **AI Challenges** (Course Coordinator) - Led development of course on AI challenges for MIT Independent Activities Period. Course took 40 MIT students through a series of challenges being developed through the USAF-MIT AI Accelerator
- 2021 **Trusted Analytics** (Course Coordinator) - Led development of course on Trusted Analytics for 35 professional students. Course focused on robust, resilient AI; human-machine teaming and AI ethics
- 2021 **AI Challenges** (Course Coordinator) - Led development of course on AI challenges for MIT Independent Activities Period. Course took 40 MIT students through a series of challenges being developed through the USAF-MIT AI Accelerator
- 2020 **Network: Graph, Neural and Social** (Course Coordinator and Instructor) - Course Lead for Networks and Analytics professional course (30 students). Course focused on a variety of network related algorithms.
- 2020 **AI Enabling Technologies** (Instructor) - Contributor to MIT Open Courseware series on Mathematics of Big Data (led by Jeremy Kepner). Course presentation has ~500K views on YouTube
- 2018 & 2019 **Practical Machine Learning** (Course Coordinator and Instructor) - Course lead for Practical Machine Learning course for professional audiences as well as at Lincoln Laboratory (100+ students total). Course focused on applying HPC to AI development.
- 2018 **Modern Algorithms** (Course Coordinator and Instructor)- Developed course on Modern Algorithms (in collaboration with others) for professional course (30 participants). Course focused on modern algorithms and performance tuning.
- 2021 **Modern Graph Analysis** (Course Coordinator and Instructor) - Course lead for semester course on Modern Graph Analysis. Course focused on various facets of graph algorithms.
- 2016 **Privacy Preserving Technologies** (Course Coordinator and Instructor) - Developed course on Privacy Preserving Technologies for Databases f
- 2016 **Big Data Tools and Technologies** (Instructor) - MIT Lincoln Laboratory course on Big Data Tools and Technologies
- 2013 & 2014 **Advanced Database Technologies** (Course Coordinator and Instructor) - Course lead for Advanced Database Technologies professional course (70+ participants total)

Post Doctoral Associate Advisees

- Current Dr. Joseph McDonald
- 2022 Dr. Elkindi Rezig (CSAIL with Prof. Mike Stonebraker)(currently Research Scientist at CSAIL/ starting Asst. Prof. at University of Utah)
- 2022 Dr. Matthew Weiss (now full-time staff at MIT Lincoln Laboratory)
- 2022 Dr. Nathan Frey (now at Preseicent Design/Genentech)
- 2021 Dr. Hayden Jananthan (now full-time staff at MIT Lincoln Laboratory)
- 2020 Dr. Andrew Kirby (now at Scientific Simulations, LLC)

Students (Advised, Mentored and/or Supported)

- Current Oscar Moll, Ph.D Student, MIT
- Current Daniel Edelman, MEng Student, MIT

Current Baolin Li, Ph.D Student, Northeastern University
Current Chelsea Conard, M.S. Student, MIT
2022 Mihailo Isakov, Ph.D, Arizona State University
2020 Emily Do, M.Eng, MIT
2020 Matthew Hutchinson, M.Eng, MIT - **Winner of MIT Reinjes Thesis Award**
2018 Jonathan Herrera, Undergraduate Student, MIT (summer)
2018 Nicki Tubacki, M.Eng, MIT
2017 Matthew Mucklo, M.Eng, MIT
2017 Arvind Prasad, University of Michigan (summer)
2016 Catherine Yu, M.Eng, MIT
2016 Ankush Gupta, M.Eng, MIT
2016 Connor Hanlon, Undergraduate Student, MIT (summer)
2015 Shana Hutchinson (year long co-op)
2015 Brendan Gavin, PhD Student, UMass Amherst (summer)
2015 Jaroor Modi, Undergraduate Student, Stevens Institute of Technology (summer)
2014 & 2015 Braden Hancock, Stanford University (summers)
2014 Sherwin Wu, M.Eng, MIT
2014 Andrew Moran, M.Eng, MIT

Leadership and Service

Technical

- 2022 **PC Member**, NeurIPS Dataset and Challenges Track
- 2016– **Deputy Lead** of grant administered for internal research at MIT Lincoln Laboratory
- 2018– **Organizer/Chair**, Poly@VLDB Workshop: Polystores and other methods to manage heterogeneous data, VLDB
- 2015– **Technical Program Committee member** of IEEE High Performance Extreme Computing (HPEC)
- 2018– **Organizer/Chair**, Workshop on Methods to Manage Heterogeneous Data and Polystore Databases, IEEE Big Data
- 2014– Member of Lincoln Laboratory college recruiting team
- 2018 **Technical Program Committee member** of IEEE International Conference on Big Data
- 2017–2019 Technical Program Committee member of IEEE International Conference on Distributed Computing Systems (ICDCS)
- 2017 **Technical Program Committee Member** of Third International Workshop on Data Management and Analytics for Medicine and Healthcare at VLDB, 2017
- 2016 **Organizer** of SIAM Minisymposium on “Big Data and Scientific Applications”, SIAM Annual Meeting
- 2016 **Reviewer**, National Science Foundation
- 2015 **Technical Program Committee member** of Supercomputing (SC'15) conference
- 2015–2016 **Program Committee member** of ACM Parallel Programming or Analytics Applications (PPAA)
- 2014 **Reviewer**, IEEE Transaction on Intelligent Transportation Systems

Student Leadership

- 2012 **Director of Employment Concerns** for National Association for Graduate and Professional Students
- 2011 **Midwest Regional Director of Employment Concerns** for National Association for Graduate and Professional Students
- 2010–2011 **Treasurer**, Council of Graduate Students (CGS), The Ohio State University
- 2010–2011 **Member** Council on Academic Affairs (COAM), The Ohio State University
- 2010–2011 **Member**, Dean of Engineering Search Committee, The Ohio State University
- 2004–2005 **Secretary**, Student's Placement Office at the Indian Institute of Technology, Kanpur

Publications

Book and Book Chapters

- 2022 (Expected) Lauren Milechin, Shana Hutchinson, Hayden Jananthan, Jeremy Kepner, Benjamin Miller, Andrew Prout, Siddharth Samsi, Chuck Yee, and Vijay Gadepally. *Graphulo: Linear Algebra Graph Kernels*. Massive Graph Algorithms. Taylor and Francis, 2022 (Expected).
- 2009 Ashok Krishnamurthy, Siddharth Samsi, and Vijay Gadepally. *Parallel MATLAB Techniques*. Image Processing. InTech, 2009.
- 2021 Vijay Gadepally, Timothy G Mattson, Michael Stonebraker, Tim Kraska, Fusheng Wang, Gang Luo, Jun Kong, and Alevtina Dubovitskaya. *Heterogeneous Data Management, Polystores, and Analytics for Healthcare: VLDB Workshops, Poly 2020 and DMAH 2020, Virtual Event, August 31 and September 4, 2020, Revised Selected Papers*, volume 12633. Springer Nature, 2021.
- 2019 Vijay Gadepally, Timothy Mattson, Michael Stonebraker, Fusheng Wang, Gang Luo, and George Teodoro. *Heterogeneous Data Management, Polystores, and Analytics for Healthcare: VLDB 2018 Workshops, Poly and DMAH, Rio de Janeiro, Brazil, August 31, 2018, Revised Selected Papers*, volume 11470. Springer, 2019.
- 2019 Vijay Gadepally, Timothy Mattson, Michael Stonebraker, Fusheng Wang, Gang Luo, Yanhui Laing, and Alevtina Dubovitskaya. *Heterogeneous Data Management, Polystores, and Analytics for Healthcare: VLDB 2019 Workshops, Poly and DMAH, Los Angeles, CA, USA, August 30, 2019, Revised Selected Papers*, volume 11721. Springer Nature, 2019.
- 2022 (Expected) Vijay Gadepally, Sam Madden, and Michael Stonebraker. *Polystore Databases and Methods to Manage Heterogenous Data*. MIT Press (as a part of the MIT Lincoln Laboratory Book Series), 2022 (Expected).
- 2016 Vijay Gadepally, Jeremy Kepner, and Albert Reuther. *Storage and database management for big data*, volume 30 of *Big Data: Storage, Sharing and Security*. CRC Press, 2016.
- 2018 Vijay Gadepally. *The BigDAWG codeline*, pages 367–376. Making Databases Work: the Pragmatic Wisdom of Michael Stonebraker. 2018.

Journal Articles

- 2021 (accepted) Matthew Hutchinson and Vijay Gadepally. Video action understanding. *Ieee Access*. IEEE, 2021 (accepted).
- 2019 Yasar Khan, Antoine Zimmermann, Alokumar Jha, Vijay Gadepally, Mathieu d'Aquin, and Ratnesh Sahay. One size does not fit all: querying web polystores. *Ieee Access*, volume 7, pages 9598–9617. IEEE, 2019.
- 2017 Julia Mullen, Chansup Byun, Vijay Gadepally, Siddharth Samsi, Albert Reuther, and Jeremy Kepner. Learning by doing, high performance computing education in the mooc era. *Journal of Parallel and Distributed Computing*, volume 105, pages 105–115. Academic Press, 2017.
- 2017 Vijay Gadepally, Ashok Krishnamurthy, and Ümit Özgüner. A framework for estimating long term driver behavior. *Journal of advanced transportation*, volume 2017. Hindawi, 2017.
- 2017 Benjamin Fuller, Mayank Varia, Arkady Yerukhimovich, Emily Shen, Ariel Hamlin, Vijay Gadepally, Richard Shay, John Darby Mitchell, and Robert K Cunningham. Sok: Cryptographically protected database search. In *2017 IEEE Symposium on Security and Privacy (SP)*, pages 172–191. IEEE, 2017.
- 2016 Vijay N Gadepally, Braden J Hancock, Kara B Greenfield, Joseph P Campbell, William M Campbell, and Albert I Reuther. Recommender systems for the department of defense and intelligence community. *Lincoln Laboratory Journal*, volume 22, 2016.

- 2015 Nabil A Schear, Arkady B Yerukhimovich, Vijay N Gadepally, Thomas M Moyer, Patrick II T Cable, and Robert K Cunningham. Secure and resilient cloud computing for the department of defense. *Lincoln Laboratory Journal*, 2015.
- 2015 Jeremy Kepner, Vijay Gadepally, and Pete Michaleas. Percolation model of insider threats to assess the optimum number of rules. *Environment Systems and Decisions*, volume 35, pages 504–510. Springer US, 2015.
- 2013 Vijay Gadepally, Ashok Krishnamurthy, and Umit Ozguner. A framework for estimating driver decisions near intersections. *IEEE Transactions on Intelligent Transportation Systems*, volume 15, pages 637–646. IEEE, 2013.
- 2012 Vijay Gadepally, Ashok Krishnamurthy, and Umit Ozguner. A hands-on education program on cyber physical systems for high school students. *Journal of Computational Science Education*, volume 3, pages 1–17. Shodor, 2012.
- 2010 Siddharth Samsi, Vijay Gadepally, and Ashok Krishnamurthy. Matlab for signal processing on multiprocessors and multicores. *IEEE Signal Processing Magazine*, volume 27, pages 40–49. IEEE, 2010.
- 2009 David E Hudak, Neil Ludban, Ashok Krishnamurthy, Vijay Gadepally, Siddharth Samsi, and John Nehrbass. A computational science ide for hpc systems: design and applications. *International journal of parallel programming*, volume 37, pages 91–105. Springer US, 2009.

Peer-reviewed Conference & Workshop Publications

- 2023 Oscar Moll, Sam Madden, and Vijay Gadepally. Ad-hoc searches on image databases. In *Heterogeneous Data Management, Polystores, and Analytics for Healthcare: VLDB Workshops, Poly 2022 and DMAH 2022, Virtual Event, September 9, 2022, Revised Selected Papers*, pages 3–9. Springer Nature Switzerland Cham, 2023.
- 2023 Andrew Bowne, Lindsey McEvoy, Dhruv Gupta, Cameron Brown, Vijay Gadepally, and El Kindi Rezig. A survey of data challenges across a modernizing bureaucracy: A new perspective on examining old government problems. In *Heterogeneous Data Management, Polystores, and Analytics for Healthcare: VLDB Workshops, Poly 2022 and DMAH 2022, Virtual Event, September 9, 2022, Revised Selected Papers*, pages 10–23. Springer Nature Switzerland Cham, 2023.
- 2022 Dan Zhao, Nathan C Frey, Joseph McDonald, Matthew Hubbell, David Bestor, Michael Jones, Andrew Prout, Vijay Gadepally, and Siddharth Samsi. A green (er) world for ai. In *2022 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW)*, pages 742–750. IEEE, 2022.
- 2022 Dan Zhao, Nathan C Frey, Vijay Gadepally, and Siddharth Samsi. Loss curve approximations for fast neural architecture ranking & training elasticity estimation. In *2022 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW)*, pages 715–723. IEEE, 2022.
- 2022 Matthew L Weiss, Joseph McDonald, David Bestor, Charles Yee, Daniel Edelman, Michael Jones, Andrew Prout, Andrew Bowne, Lindsey McEvoy, Vijay Gadepally, et al. An evaluation of low overhead time series preprocessing techniques for downstream machine learning. In *2022 IEEE High Performance Extreme Computing Conference (HPEC)*, pages 1–6. IEEE, 2022.
- 2022 Tyler Trigg, Chad Meiners, Sandeep Pisharody, Hayden Jananthan, Michael Jones, Adam Michaleas, Timothy Davis, Erik Welch, William Arcand, David Bestor, et al. Hypersparse network flow analysis of packets with graphblas. In *2022 IEEE High Performance Extreme Computing Conference (HPEC)*, pages 1–7. IEEE, 2022.

- 2022 Benny J Tang, Qiqi Chen, Matthew L Weiss, Nathan C Frey, Joseph McDonald, David Bestor, Charles Yee, William Arcand, William Bergeron, Chansup Byun, et al. The mit supercloud workload classification challenge. In *2022 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW)*, pages 708–714. IEEE, 2022.
- 2022 Rohan Basu Roy, Tirthak Patel, Vijay Gadepally, and Devesh Tiwari. Mashup: making serverless computing useful for hpc workflows via hybrid execution. In *Proceedings of the 27th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*, pages 46–60, 2022.
- 2022 Albert Reuther, Peter Michaleas, Michael Jones, Vijay Gadepally, Siddharth Samsi, and Jeremy Kepner. Ai and ml accelerator survey and trends. In *2022 IEEE High Performance Extreme Computing Conference (HPEC)*, pages 1–10. IEEE, 2022.
- 2022 Oscar Moll, Favyen Bastani, Sam Madden, Mike Stonebraker, Vijay Gadepally, and Tim Kraska. Exsample: Efficient searches on video repositories through adaptive sampling. In *2022 IEEE 38th International Conference on Data Engineering (ICDE)*, pages 2956–2968. IEEE, 2022.
- 2022 Joseph McDonald, Baolin Li, Nathan Frey, Devesh Tiwari, Vijay Gadepally, and Siddharth Samsi. Great power, great responsibility: Recommendations for reducing energy for training language models. *arXiv preprint arXiv:2205.09646*, 2022.
- 2022 Joseph McDonald, James M Kurdzo, Phillip M Stepanian, Mark Veillette, David Bestor, Michael Jones, Vijay Gadepally, and Siddharth Samsi. Performance estimation for efficient image segmentation training of weather radar algorithms. In *2022 IEEE High Performance Extreme Computing Conference (HPEC)*, pages 1–7. IEEE, 2022.
- 2022 Baolin Li, Tirthak Patel, Siddharth Samsi, Vijay Gadepally, and Devesh Tiwari. Miso: exploiting multi-instance gpu capability on multi-tenant gpu clusters. In *Proceedings of the 13th Symposium on Cloud Computing*, pages 173–189, 2022.
- 2022 Baolin Li, Tirthak Patel, Siddharth Samsi, Vijay Gadepally, and Devesh Tiwari. Using multi-instance gpu for efficient operation of multi-tenant gpu clusters. *arXiv preprint arXiv:2207.11428*, 2022.
- 2022 Baolin Li, Tirthak Patel, Vijay Gadepally, Karen Gettings, Siddharth Samsi, and Devesh Tiwari. Dash: Scheduling deep learning workloads on multi-generational gpu-accelerated clusters. In *2022 IEEE High Performance Extreme Computing Conference (HPEC)*, pages 1–7. IEEE, 2022.
- 2022 Baolin Li, Rohin Arora, Siddharth Samsi, Tirthak Patel, William Arcand, David Bestor, Chansup Byun, Rohan Basu Roy, Bill Bergeron, John Holodnak, et al. Ai-enabling workloads on large-scale gpu-accelerated system: Characterization, opportunities, and implications. In *2022 IEEE International Symposium on High-Performance Computer Architecture (HPCA)*, pages 1224–1237. IEEE, 2022.
- 2022 Jeremy Kepner, Michael Jones, Daniel Andersen, Aydin Buluc, Chansup Byun, K Claffy, Timothy Davis, William Arcand, Jonathan Bernays, David Bestor, et al. Temporal correlation of internet observatories and outposts. In *2022 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW)*, pages 247–254. IEEE, 2022.
- 2022 Michael Jones, Jeremy Kepner, Daniel Andersen, Aydin Buluç, Chansup Byun, K Claffy, Timothy Davis, William Arcand, Jonathan Bernays, David Bestor, et al. Graphblas on the edge: Anonymized high performance streaming of network traffic. In *2022 IEEE High Performance Extreme Computing Conference (HPEC)*, pages 1–8. IEEE, 2022.
- 2022 Vijay Gadepally, Gregory Angelides, Andrei Barbu, Andrew Bowne, Laura J Brattain, Tamara Broderick, Armando Cabrera, Glenn Carl, Ronisha Carter, Miriam Cha, et al. Developing a series of ai challenges for the united states department of the air force. In *2022 IEEE High Performance Extreme Computing Conference (HPEC)*, pages 1–7. IEEE, 2022.

- 2022 Nathan C Frey, Dan Zhao, Simon Axelrod, Michael Jones, David Bestor, Vijay Gadepally, Rafael Gómez-Bombarelli, and Siddharth Samsi. Energy-aware neural architecture selection and hyperparameter optimization. In *2022 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW)*, pages 732–741. IEEE, 2022.
- 2022 Nathan C Frey, Baolin Li, Joseph McDonald, Dan Zhao, Michael Jones, David Bestor, Devesh Tiwari, Vijay Gadepally, and Siddharth Samsi. Benchmarking resource usage for efficient distributed deep learning. In *2022 IEEE High Performance Extreme Computing Conference (HPEC)*, pages 1–8. IEEE, 2022.
- 2022 Nathan Frey, Ryan Soklaski, Simon Axelrod, Siddharth Samsi, Rafael Gomez-Bombarelli, Connor Coley, and Vijay Gadepally. Neural scaling of deep chemical models. 2022.
- 2022 Rohan Basu Roy, Tirthak Patel, Vijay Gadepally, and Devesh Tiwari. Mashup: Making serverless computing useful for hpc workflows via hybrid execution. *ACM| 27th ACM SIGPLAN Symposium on Principles and Practice of Parallel . . .*, 2022.
- 2021 Kaira Samuel, Vijay Gadepally, David Jacobs, Michael Jones, Kyle McAlpin, Kyle Palko, Ben Paulk, Sid Samsi, Ho Chit Siu, Charles Yee, et al. Maneuver identification challenge. In *2021 IEEE High Performance Extreme Computing Conference (HPEC)*, pages 1–7. IEEE, 2021.
- 2021 Rohan Basu Roy, Tirthak Patel, Vijay Gadepally, and Devesh Tiwari. Bliss: auto-tuning complex applications using a pool of diverse lightweight learning models. In *Proceedings of the 42nd ACM SIGPLAN International Conference on Programming Language Design and Implementation (PLDI)*, pages 1280–1295, 2021.
- 2021 El Kindi Rezig, Allan Vanterpool, Vijay Gadepally, Benjamin Price, Michael Cafarella, and Michael Stonebraker. Towards data discovery by example. In *Heterogeneous Data Management, Polystores, and Analytics for Healthcare: VLDB Workshops, Poly 2020 and DMAH 2020, Virtual Event, August 31 and September 4, 2020, Revised Selected Papers 6*, pages 66–71. Springer International Publishing, 2021.
- 2021 El Kindi Rezig, Anshul Bhandari, Anna Fariha, Benjamin Price, Allan Vanterpool, Vijay Gadepally, and Michael Stonebraker. Dice: data discovery by example. *Proceedings of the VLDB Endowment*, volume 14, pages 2819–2822. VLDB Endowment, 2021.
- 2021 El Kindi Rezig, Anshul Bhandari, Anna Fariha, Benjamin Price, Allan Vanterpool, Vijay Gadepally, and Michael Stonebraker. Dice: Data discovery by example. In *Proceedings of the VLDB Endowment*, 2021.
- 2021 Albert Reuther, Peter Michaleas, Michael Jones, Vijay Gadepally, Siddharth Samsi, and Jeremy Kepner. Ai accelerator survey and trends. In *2021 IEEE High Performance Extreme Computing Conference (HPEC)*, pages 1–9. IEEE, 2021.
- 2021 Joseph McDonald, Siddharth Samsi, Daniel Edelman, Chansup Byun, Jeremy Kepner, and Vijay Gadepally. Improved compression for word embeddings by scaling principal components. In *2021 IEEE High Performance Extreme Computing Conference (HPEC)*, pages 1–7. IEEE, 2021.
- 2021 Baolin Li, Rohan Basu Roy, Tirthak Patel, Vijay Gadepally, Karen Gettings, and Devesh Tiwari. Ribbon: cost-effective and qos-aware deep learning model inference using a diverse pool of cloud computing instances. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, pages 1–13, 2021.
- 2021 Baolin Li, Rohan Roy, Tirthak Patel, Vijay Gadepally, Karen Gettings, and Devesh Tewari. Ribbon: cost-effective and qos-aware deep learning model inference using a diverse pool of cloud computing instances. In *IEEE/ACM Confernece on Supercomputing*, 2021.
- 2021 Baolin Li, Vijay Gadepally, Siddharth Samsi, Mark Veillette, and Devesh Tiwari. Serving machine learning inference using heterogeneous hardware. In *2021 IEEE High Performance Extreme Computing Conference (HPEC)*, pages 1–8. IEEE, 2021.

- 2021 Michael Cafarella, David DeWitt, Vijay Gadepally, Jeremy Kepner, Christos Kozyrakis, Tim Kraska, Michael Stonebraker, and Matei Zaharia. A polystore based database operating system (dbos). In *Heterogeneous Data Management, Polystores, and Analytics for Healthcare: VLDB Workshops, Poly 2020 and DMAH 2020, Virtual Event, August 31 and September 4, 2020, Revised Selected Papers 6*, pages 3–24. Springer International Publishing, 2021.
- 2021 Chansup Byun, William Arcand, David Bestor, Bill Bergeron, Vijay Gadepally, Michael Houle, Matthew Hubbell, Michael Jones, Anna Klein, Peter Michaleas, et al. Node-based job scheduling for large scale simulations of short running jobs. In *2021 IEEE High Performance Extreme Computing Conference (HPEC)*, pages 1–7. IEEE, 2021.
- 2021 Bill Bergeron, Matthew Hubbell, Dylan Sequeira, Winter Williams, William Arcand, David Bestor, Chansup Byun, Vijay Gadepally, Michael Houle, Michael Jones, et al. 3d real-time supercomputer monitoring. In *2021 IEEE High Performance Extreme Computing Conference (HPEC)*, pages 1–7. IEEE, 2021.
- 2020 Michael Stonebraker, Timothy Mattson, Tim Kraska, and Vijay Gadepally. Poly'19 workshop summary: Gdpr. *ACM SIGMOD Record*, volume 49, pages 55–58. ACM New York, NY, USA, 2020.
- 2020 Albert Reuther, Peter Michaleas, Michael Jones, Vijay Gadepally, Siddharth Samsi, and Jeremy Kepner. Survey of machine learning accelerators. In *2020 IEEE High Performance Extreme Computing Conference (HPEC)*, pages 1–12. IEEE, 2020.
- 2020 Andrew Kirby, Siddharth Samsi, Michael Jones, Albert Reuther, Jeremy Kepner, and Vijay Gadepally. Layer-parallel training with gpu concurrency of deep residual neural networks via nonlinear multigrid. In *2020 IEEE High Performance Extreme Computing Conference (HPEC)*, pages 1–7. IEEE, 2020. **Outstanding Paper Award**.
- 2020 Matthew Hutchinson, Siddharth Samsi, William Arcand, David Bestor, Bill Bergeron, Chansup Byun, Micheal Houle, Matthew Hubbell, Micheal Jones, Jeremy Kepner, et al. Accuracy and performance comparison of video action recognition approaches. In *2020 IEEE High Performance Extreme Computing Conference (HPEC)*, pages 1–8. IEEE, 2020.
- 2020 Vijay Gadepally, Mihailo Isakov, Rashmi Agrawal, Jeremy Kepner, Karen Gettings, and Michel A Kinsy. Homomorphic encryption based secure sensor data processing. In *2020 IEEE High Performance Extreme Computing Conference (HPEC)*, pages 1–7. IEEE, 2020.
- 2020 Emily H Do and Vijay N Gadepally. Classifying anomalies for network security. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 2907–2911. IEEE, 2020.
- 2019 Richard Shay, Uri Blumenthal, Vijay Gadepally, Ariel Hamlin, John Darby Mitchell, and Robert K Cunningham. Don't even ask: Database access control through query control. *ACM SIGMOD Record*, volume 47, pages 17–22. ACM, 2019.
- 2019 Albert Reuther, Peter Michaleas, Michael Jones, Vijay Gadepally, Siddharth Samsi, and Jeremy Kepner. Survey and benchmarking of machine learning accelerators. In *2019 IEEE high performance extreme computing conference (HPEC)*, pages 1–9. IEEE, 2019.
- 2019 Jeremy Kepner, Simon Alford, Vijay Gadepally, Michael Jones, Lauren Milechin, Ryan Robinett, and Sid Samsi. Sparse deep neural network graph challenge. In *2019 IEEE High Performance Extreme Computing Conference (HPEC)*, pages 1–7. IEEE, 2019.
- 2019 Mihailo Isakov, Vijay Gadepally, Karen M Gettings, and Michel A Kinsy. Survey of attacks and defenses on edge-deployed neural networks. In *2019 IEEE High Performance Extreme Computing Conference (HPEC)*, pages 1–8. IEEE, 2019. **Outstanding Student Paper Award (finalist)**.
- 2019 Pieter Cailliau, Tim Davis, Vijay Gadepally, Jeremy Kepner, Roi Lipman, Jeffrey Lovitz, and Keren Ouaknine. Redisgraph graphblas enabled graph database. In *2019 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW)*, pages 285–286. IEEE, 2019.

- 2018 Xiangyao Yu, Vijay Gadepally, Stan Zdonik, Tim Kraska, and Michael Stonebraker. Fastdawg: improving data migration in the bigdawg polystore system. In *Heterogeneous Data Management, Polystores, and Analytics for Healthcare*, pages 3–15. Springer, Cham, 2018.
- 2018 Siddharth Samsi, Vijay Gadepally, Michael Hurley, Michael Jones, Edward Kao, Sanjeev Mohindra, Paul Monticciolo, Albert Reuther, Steven Smith, William Song, et al. Graphchallenge. org: Raising the bar on graph analytic performance. In *2018 IEEE High Performance extreme Computing Conference (HPEC)*, pages 1–7. IEEE, 2018.
- 2018 Albert Reuther, Jeremy Kepner, Chansup Byun, Siddharth Samsi, William Arcand, David Bestor, Bill Bergeron, Vijay Gadepally, Michael Houle, Matthew Hubbell, et al. Interactive supercomputing on 40,000 cores for machine learning and data analysis. In *2018 IEEE High Performance extreme Computing Conference (HPEC)*, pages 1–6. IEEE, 2018.
- 2018 Andrew Prout, William Arcand, David Bestor, Bill Bergeron, Chansup Byun, Vijay Gadepally, Michael Houle, Matthew Hubbell, Michael Jones, Anna Klein, et al. Measuring the impact of spectre and meltdown. In *2018 IEEE High Performance extreme Computing Conference (HPEC)*, pages 1–5. IEEE, 2018.
- 2018 Julia Mullen, Albert Reuther, William Arcand, Bill Bergeron, David Bestor, Chansup Byun, Vijay Gadepally, Michael Houle, Matthew Hubbell, Michael Jones, et al. Lessons learned from a decade of providing interactive, on-demand high performance computing to scientists and engineers. In *International Conference on High Performance Computing*, pages 655–668. Springer, Cham, 2018.
- 2018 Jeremy Kepner, Vikalo Gadepally, Hayden Jananthan, Lauren Milechin, and Sid Samsi. Sparse deep neural network exact solutions. In *2018 IEEE High Performance extreme Computing Conference (HPEC)*, pages 1–8. IEEE, 2018.
- 2018 Vijay Gadepally, Jeremy Kepner, Lauren Milechin, William Arcand, David Bestor, Bill Bergeron, Chansup Byun, Matthew Hubbell, Micheal Houle, Micheal Jones, et al. Hyperscaling internet graph analysis with d4m on the mit supercloud. In *2018 IEEE High Performance extreme Computing Conference (HPEC)*, pages 1–6. IEEE, 2018.
- 2017 Katherine Yu, Vijay Gadepally, and Michael Stonebraker. Database engine integration and performance analysis of the bigdawg polystore system. In *2017 IEEE High Performance Extreme Computing Conference (HPEC)*, pages 1–7. IEEE, 2017.
- 2017 Ran Tan, Rada Chirkova, Vijay Gadepally, and Timothy G Mattson. Enabling query processing across heterogeneous data models: A survey. In *2017 IEEE International Conference on Big Data (Big Data)*, pages 3211–3220. IEEE, 2017.
- 2017 Siddharth Samsi, Vijay Gadepally, Michael Hurley, Michael Jones, Edward Kao, Sanjeev Mohindra, Paul Monticciolo, Albert Reuther, Steven Smith, William Song, et al. Static graph challenge: Subgraph isomorphism. In *2017 IEEE High Performance Extreme Computing Conference (HPEC)*, pages 1–6. IEEE, 2017.
- 2017 Andrew Prout, William Arcand, David Bestor, Bill Bergeron, Chansup Byun, Vijay Gadepally, Matthew Hubbell, Michael Houle, Michael Jones, Peter Michaleas, et al. Mit supercloud portal workspace: Enabling hpc web application deployment. In *2017 IEEE High Performance Extreme Computing Conference (HPEC)*, pages 1–6. IEEE, 2017.
- 2017 Oscar Moll, Aaron Zalewski, Sudeep Pillai, Sam Madden, Michael Stonebraker, and Vijay Gadepally. Exploring big volume sensor data with vroom. *Proceedings of the VLDB Endowment*, volume 10, pages 1973–1976. VLDB Endowment, 2017.
- 2017 Tim Mattson, Vijay Gadepally, Zuohao She, Adam Dziedzic, and Jeff Parkhurst. Demonstrating the bigdawg polystore system for ocean metagenomics analysis. In *CIDR*, 2017.

- 2017 Edward Kao, Vijay Gadepally, Michael Hurley, Michael Jones, Jeremy Kepner, Sanjeev Mohindra, Paul Monticciolo, Albert Reuther, Siddharth Samsi, William Song, et al. Streaming graph challenge: Stochastic block partition. In *2017 IEEE High performance extreme computing conference (HPEC)*, pages 1–12. IEEE, 2017.
- 2017 Hayden Jananathan, Ziqi Zhou, Vijay Gadepally, Dylan Hutchison, Suna Kim, and Jeremy Kepner. Polystore mathematics of relational algebra. In *2017 IEEE International Conference on Big Data (Big Data)*, pages 3180–3189. IEEE, 2017.
- 2017 Vijay Gadepally, Kyle O'Brien, Adam Dziedzic, Aaron Elmore, Jeremy Kepner, Samuel Madden, Tim Mattson, Jennie Rogers, Zuohao She, and Michael Stonebraker. Bigdawg version 0.1. In *2017 IEEE High Performance Extreme Computing Conference (HPEC)*, pages 1–7. IEEE, 2017.
- 2017 Laura Brattain, Mihnea Bulugoiu, Adam Brewster, Mark Hernandez, Heejin Choi, Taeyun Ku, Kwanghun Chung, and Vijay Gadepally. A cloud-based brain connectivity analysis tool. In *2017 IEEE High Performance Extreme Computing Conference (HPEC)*, pages 1–6. IEEE, 2017. **Outstanding Student Paper Award Finalist.**
- 2017 Edmon Begoli, J Blair Christian, Vijay Gadepally, and Stavros Papadopoulos. An emerging role for polystores in precision medicine. In *VLDB Workshop on Data Management and Analytics for Medicine and Healthcare*, pages 41–52. Springer, Cham, 2017.
- 2016 Siddharth Samsi, Laura Brattain, William Arcand, David Bestor, Bill Bergeron, Chansup Byun, Vijay Gadepally, Matthew Hubbell, Michael Jones, Anna Klein, et al. Benchmarking scidb data import on hpc systems. In *2016 IEEE High Performance Extreme Computing Conference (HPEC)*, pages 1–5. IEEE, 2016. **Best Paper Award Finalist.**
- 2016 Thomas Moyer and Vijay Gadepally. High-throughput ingest of data provenance records into accumulo. In *2016 IEEE High Performance Extreme Computing Conference (HPEC)*, pages 1–6. IEEE, 2016.
- 2016 Jeremy Kepner, Vijay Gadepally, Dylan Hutchison, Hayden Jananathan, Timothy Mattson, Siddharth Samsi, and Albert Reuther. Associative array model of sql, nosql, and newsql databases. In *2016 IEEE High Performance Extreme Computing Conference (HPEC)*, pages 1–9. IEEE, 2016.
- 2016 Mike Jones, Bill Arcand, Bill Bergeron, David Bestor, Chansup Byun, Lauren Milechin, Vijay Gadepally, Matt Hubbell, Jeremy Kepner, Pete Michaleas, et al. Scalability of vm provisioning systems. In *2016 IEEE High Performance Extreme Computing Conference (HPEC)*, pages 1–5. IEEE, 2016.
- 2016 Dylan Hutchison, Jeremy Kepner, Vijay Gadepally, and Bill Howe. From nosql accumulo to newsql graphulo: Design and utility of graph algorithms inside a bigtable database. In *2016 IEEE High Performance Extreme Computing Conference (HPEC)*, pages 1–9. IEEE, 2016. **Best Student Paper Award.**
- 2016 Ankush M Gupta, Vijay Gadepally, and Michael Stonebraker. Cross-engine query execution in federated database systems. In *2016 IEEE High Performance Extreme Computing Conference (HPEC)*, pages 1–6. IEEE, 2016. **Outstanding Student Paper Award Finalist.**
- 2016 Brendan Gavin, Vijay Gadepally, and Jeremy Kepner. Enforced sparse non-negative matrix factorization. In *2016 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW)*, pages 902–911. IEEE, 2016.
- 2016 Vijay Gadepally, Peinan Chen, Jennie Duggan, Aaron Elmore, Brandon Haynes, Jeremy Kepner, Samuel Madden, Tim Mattson, and Michael Stonebraker. The bigdawg polystore system and architecture. In *2016 IEEE High Performance Extreme Computing Conference (HPEC)*, pages 1–6. IEEE, 2016.

- 2016 Patrick Dreher, Chansup Byun, Chris Hill, Vijay Gadepally, Bradley Kuszmaul, and Jeremy Kepner. Pagerank pipeline benchmark: Proposal for a holistic system benchmark for big-data platforms. In *2016 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW)*, pages 929–937. IEEE, 2016.
- 2016 Peinan Chen, Vijay Gadepally, and Michael Stonebraker. The bigdawg monitoring framework. In *2016 IEEE High Performance Extreme Computing Conference (HPEC)*, pages 1–6. IEEE, 2016.
- 2016 Alexander Chen, Alan Edelman, Jeremy Kepner, Vijay Gadepally, and Dylan Hutchison. Julia implementation of the dynamic distributed dimensional data model. In *2016 IEEE High Performance Extreme Computing Conference (HPEC)*, pages 1–7. IEEE, 2016. **Best Paper Award.**
- 2015 Andrew Prout, Jeremy Kepner, Peter Michaleas, William Arcand, David Bestor, Bill Bergeron, Chansup Byun, Lauren Edwards, Vijay Gadepally, Matthew Hubbell, et al. Enabling on-demand database computing with mit supercloud database management system. In *2015 IEEE High Performance Extreme Computing Conference (HPEC)*, pages 1–6. IEEE, 2015.
- 2015 Andrew Moran, Vijay Gadepally, Matthew Hubbell, and Jeremy Kepner. Improving big data visual analytics with interactive virtual reality. In *2015 IEEE high performance extreme computing conference (HPEC)*, pages 1–6. IEEE, 2015.
- 2015 Jeremy Kepner, Vijay Gadepally, and Pete Michaleas. Percolation model of insider threats to assess the optimum number of rules. *Environment Systems and Decisions*, volume 35, pages 504–510. Springer US, 2015.
- 2015 Dylan Hutchison, Jeremy Kepner, Vijay Gadepally, and Adam Fuchs. Graphulo implementation of server-side sparse matrix multiply in the accumulo database. In *2015 IEEE High Performance Extreme Computing Conference (HPEC)*, pages 1–7. IEEE, 2015.
- 2015 Matthew Hubbell, Andrew Moran, William Arcand, David Bestor, Bill Bergeron, Chansup Byun, Vijay Gadepally, Peter Michaleas, Julie Mullen, Andrew Prout, et al. Big data strategies for data center infrastructure management using a 3d gaming platform. In *2015 IEEE High Performance Extreme Computing Conference (HPEC)*, pages 1–6. IEEE, 2015.
- 2015 Vijay Gadepally, Jeremy Kepner, William Arcand, David Bestor, Bill Bergeron, Chansup Byun, Lauren Edwards, Matthew Hubbell, Peter Michaleas, Julie Mullen, et al. D4m: Bringing associative arrays to database engines. In *2015 IEEE High Performance Extreme Computing Conference (HPEC)*, pages 1–6. IEEE, 2015.
- 2015 Vijay Gadepally and Jeremy Kepner. Using a power law distribution to describe big data. In *2015 IEEE High Performance Extreme Computing Conference (HPEC)*, pages 1–5. IEEE, 2015.
- 2015 Vijay Gadepally, Taylor Herr, Luke Johnson, Lauren Milechin, Maja Milosavljevic, and Benjamin A Miller. Sampling operations on big data. In *2015 49th Asilomar Conference on Signals, Systems and Computers*, pages 1515–1519. IEEE, 2015.
- 2015 Vijay Gadepally, Braden Hancock, Benjamin Kaiser, Jeremy Kepner, Pete Michaleas, Mayank Varia, and Arkady Yerukhimovich. Computing on masked data to improve the security of big data. In *2015 IEEE International Symposium on Technologies for Homeland Security (HST)*, pages 1–6. IEEE, 2015.
- 2015 Vijay Gadepally, Jake Bolewski, Dan Hook, Dylan Hutchison, Ben Miller, and Jeremy Kepner. Graphulo: Linear algebra graph kernels for nosql databases. In *2015 IEEE International Parallel and Distributed Processing Symposium Workshop*, pages 822–830. IEEE, 2015.
- 2015 Aaron J Elmore, Jennie Duggan, Michael Stonebraker, Magdalena Balazinska, Ugur Cetintemel, Vijay Gadepally, Jeffrey Heer, Bill Howe, Jeremy Kepner, Tim Kraska, et al. A demonstration of the bigdawg polystore system. *Proceedings of the VLDB Endowment*, volume 8, page 1908, 2015.

- 2015 Adam Dziedzic, Jennie Duggan, Aaron J Elmore, Vijay Gadepally, and Michael Stonebraker. Bigdawg: a polystore for diverse interactive applications. In *Data Syst. Interactive Anal. Workshop 2015*, 2015.
- 2014 Sophia Yakoubov, Vijay Gadepally, Nabil Schear, Emily Shen, and Arkady Yerukhimovich. A survey of cryptographic approaches to securing big-data analytics in the cloud. In *2014 IEEE High Performance Extreme Computing Conference (HPEC)*, pages 1–6. IEEE, 2014.
- 2014 Zachary Weber and Vijay Gadepally. Using 3d printing to visualize social media big data. *arXiv preprint arXiv:1409.7724*, 2014.
- 2014 Jeremy Kepner, Vijay Gadepally, Pete Michaleas, Nabil Schear, Mayank Varia, Arkady Yerukhimovich, and Robert K Cunningham. Computing on masked data: a high performance method for improving big data veracity. In *2014 IEEE High Performance Extreme Computing Conference (HPEC)*, pages 1–6. IEEE, 2014.
- 2014 Jeremy Kepner, William Arcand, David Bestor, Bill Bergeron, Chansup Byun, Vijay Gadepally, Matthew Hubbell, Peter Michaleas, Julie Mullen, Andrew Prout, et al. Achieving 100,000,000 database inserts per second using accumulo and d4m. In *2014 IEEE High Performance Extreme Computing Conference (HPEC)*, pages 1–6. IEEE, 2014.
- 2014 Vijay Gadepally and Jeremy Kepner. Big data dimensional analysis. In *2014 IEEE High Performance Extreme Computing Conference (HPEC)*, pages 1–6. IEEE, 2014.
- 2011 Vijay Gadepally, Arda Kurt, Ashok Krishnamurthy, and Ümit Özgüner. Driver/vehicle state estimation and detection. In *2011 14th International IEEE Conference on Intelligent Transportation Systems (ITSC)*, pages 582–587. IEEE, 2011.
- 2009 Juan C Chaves, Alan Chalker, David Hudak, Vijay Gadepally, Fernando Escobar, and Patrick Longhini. Enabling high-productivity sip application development: Modeling and simulation of superconducting quantum interference filters. In *2009 DoD High Performance Computing Modernization Program Users Group Conference*, pages 302–306. IEEE, 2009.
- 2008 Ashok Krishnamurthy, David Hudak, John Nehrbass, Siddharth Samsi, and Vijay Gadepally. Parallel matlab in production supercomputing with applications in signal and image processing. In *SIAM Parallel Processing for Scientific Computing*. SIAM, 2008.
- 2007 David E Hudak, Neil Ludban, Vijay Gadepally, and Ashok Krishnamurthy. Developing a computational science ide for hpc systems. In *Third International Workshop on Software Engineering for High Performance Computing Applications (SE-HPC'07)*, pages 5–5. IEEE, 2007.

Technical Reports (not peer-reviewed)

- 2021 El Kindi Rezig, Michael Cafarella, and Vijay Gadepally. Technical report on data integration and preparation. *arXiv preprint arXiv:2103.01986*, 2021.
- 2020 Oscar Moll, Favien Bastani, Sam Madden, Mike Stonebraker, Vijay Gadepally, and Tim Kraska. Exsample: Efficient searches on video repositories through adaptive sampling. *arXiv preprint arXiv:2005.09141*, 2020.
- 2020 Jeremy Kepner, Vijay Gadepally, Hayden Jananthan, Lauren Milechin, and Siddharth Samsi. Ai data wrangling with associative arrays. *arXiv preprint arXiv:2001.06731*, 2020.
- 2020 Matthew Hutchinson and Vijay Gadepally. Video action understanding: A tutorial. *arXiv preprint arXiv:2010.06647*, 2020.
- 2020 Vijay Gadepally and Jeremy Kepner. Technical report: Developing a working data hub. *arXiv preprint arXiv:2004.00190*, 2020.
- 2020 Michael Cafarella, David DeWitt, Vijay Gadepally, Jeremy Kepner, Christos Kozyrakis, Tim Kraska, Michael Stonebraker, and Matei Zaharia. Dbos: A proposal for a data-centric operating system. *arXiv preprint arXiv:2007.11112*, 2020.

2019 Vijay Gadepally, Justin Goodwin, Jeremy Kepner, Albert Reuther, Hayley Reynolds, Siddharth Samsi, Jonathan Su, and David Martinez. Ai enabling technologies: A survey. *arXiv preprint arXiv:1905.03592*, 2019.

Selected Seminars, Talks, and Posters

- 2022–Present "Understanding (and reducing) the Energy Impact of AI", Talk Given at:
- HPCA NetZero Workshop (2023)
 - TTI/Vanguard (2023)
 - Open Compute Project Meeting (2023)
 - JP Morgan Chase (2022)
- 2022–Present Using AI to improve AI Development and Deployment
- CLSAC (2022)
 - Computational Research in Boston and Beyond Seminar Series (2021)
 - JP Morgan Chase (2022)
- 2020–Present "AI Accelerator Challenge Problems", Talk given at:
- USAF-MIT AIA Seminar (2023)
 - IEEE HPEC (2022)
 - DoD High Performance Computing Modernization Program UGM 2022
 - IEEE HPEC (2021)
 - USAF-MIT AIA Seminar (2021)
 - Recent Advances in Artificial Intelligence for National Security (RAAINS) Workshop (2020)
- 2017–Present "Fast AI: AI Enabling Technologies", Talk given at:
- MIT Lincoln Laboratory (Labwide Seminar) (2021)
 - MIRACLE Forum, Air Force Research Laboratory (2021)
 - Louisiana State University AI Colloquium (2021)
 - Open Data Science Conference (ODSC) East (2020)
 - Wright Brothers Institute (2019)
 - Air Force Research Laboratory (2018)
 - MGHPCC Day, Northeastern University (2018)
 - Enterprise Resilience Conference: Software for HPC session (2018)
 - The Ohio State University (2017)
- 2021 Panelist, Polstore Systems and their Future, Poly@VLDB'21 Workshop
- 2021 "An Open Data Center Dataset for AI-enabled Optimization", Talk given with Siddharth Samsi at:
- NVIDIA GTC (2021)
 - Computational Research in Boston and Beyond Seminar Series (2021)
- 2021 Panelist, Legal Implications of Artificial Intelligence, MIT IAP course

- 2014–2017 “Addressing Big Data Challenges through Innovative Architecture, Databases and Software”, Talk given at:
- Raytheon Technology Networks Symposium (2017)
 - IoT Summit, *Keynote* (2016)
 - University of Michigan (2016)
 - The Ohio State University (2016)
 - Smith College (2015)
 - Tata Consultancy Services (2015)
 - Big Data Innovation Boston (October 2014)
 - Accumulo Summit (June 2014)
 - The Data Warehousing Institute (February 2014)
 - MITRE Corporation (December 2014)
- 2017 Panelist, “Data Analytics”, NSF Smart Grids Workshop (2018)
- 2017 “Introduction to Polystore Databases and the BigDAWG Polystore System”, Tutorial at IEEE High Performance Extreme Computing (HPEC) Conference, 2017
- 2016 – 2017 “BigDAWG Polystore and Architecture”, Talk given at:
- University of California, San Diego (2017)
 - Oak Ridge National Laboratory (2016)
 - North Carolina State University (2016)
 - Chesapeake Bay Large Scale Analytics Conference (2016)
 - Intel (2015, 2016)
 - New England Database Day (2016)
- 2016 “Genomics Data, Analytics and the Future of Climate Change”, Blog Post for Intel Science and Technology Center for Big Data, August 12, 2016
- 2017 Panelist, “Open edX Depolymnet Panel”, OpenEdX Conference (2016)
- 2015 – 2016 “Graphulo: Graph Analytics in NoSQL databases”, Talk given at:
- Accumulo Summit (2016)
 - Graph Exploitation Symposium (2015)
- 2015 “Using D4M for rapid prototyping of analytics for Apache Accumulo”, Talk given at Accumulo Summit (2015)
- 2015 “Sifter: A Data Corpus Generator”, Vijay Gadepally, Sherwin Wu, Jeremy Kepner, Sam Madden, Poster at New England Database Day (2015)
- 2014–2015 “Computing on Masked Data”, Talk given at:
- MIT Lincoln Laboratory (2015)
 - Advanced Research and Technology Symposium (2015)
 - Cyber Netcentric Workshop at MIT Lincoln Laboratory (2014)
 - Intel Big Data Retreat (2014)
- 2014 “SciDB - Manage and Analyze Terabytes of Array Data”, Tutorial at SuperComputing - SC'14 (2014)
- 2014–2015 “Novel Methods to Visualize and Interact with Big Data”, Andrew Moran, Manuela Caicedo-Santiago, Zachary J. Weber, Matt Hubbell, Vijay Gadepally, Poster at:
- IEEE High Performance Extreme Computing (2014)
 - Big Data at MIT Media Laboratory Showcase (2015)

- 2014 “MIMICViz: Enabling Visualization of Medical Big Data”, Sherwin Wu, Vijay Gadepally, Andrew Whitaker, Jeremy Kepner, Bill Howe, Magdalena Balazinska, Sam Madden, Poster at Intel Big Data Retreat, (2014).
- 2011 “The future of personal transportation”, Invited Talk at STIR Symposium (2011)
- 2010 “Policy questions in autonomous vehicle design – an analysis of what has been done and what needs to be done”, Term project for Public Policy 880.05
- 2008 “GRAPE Parallelization”, Vijay Gadepally, Siddharth Samsi, John Nehrbass, Ashok Krishnamurthy, Technical Report for Army Research Laboratory (2008)
- 2007 “Survey of Parallel Extensions to Octave and Python”, Vijay Gadepally, Technical Report for Department of Defense High Performance Computing Modernization Program (2007)

Citizenship

United States of America

References

Available upon request