

Sidharth Kumar

Department of Computer Science, University of Illinois at Chicago

Phone: +1-202- 701-5268

Email: sidharth.science@gmail.com

Website: <https://sidharthkumar.io/>

Google Scholar: <https://scholar.google.com/citations?user=TDhtzHcAAAAJ&hl=en>

EDUCATION

Doctor of Philosophy in Computing 2016

Scientific Computing and Imaging Institute, University of Utah, Salt Lake City, UT, USA

Dissertation - **A Scalable and Tunable Adaptive Resolution Parallel I/O Framework**

Advisor - Valerio Pascucci

Committee Members: Martin Berzins, Mary Hall, Feifi Li, Venkatram Vishwanathan

Bachelor of Technology in Information and Communication Technology 2009

DA-IICT, Gandhinagar, Gujarat, India

Thesis - **Air for Human Simulator**

Advisor - Naresh Jotwani and Gautam Dutta

RESEARCH INTERESTS

High-performance computing, big and scalable data systems, large data processing, parallel I/O, storage and file systems, insitu analytics and visualization, scalable algorithms, large data file formats, progressive processing, scientific data, interactive techniques, and scientific visualization

CURRENT APPOINTMENT

Assistant Professor

Department of Computer Science

University of Illinois at Chicago

PRIOR PROFESSIONAL EXPERIENCE

Assistant professor 2018 - 2023

Department of Computer Science, University of Alabama at Birmingham, Birmingham, AL

Associate Scientist 2021 - 2023

Comprehensive Neuroscience Center, Heersink School of Medicine, University of Alabama at Birmingham

Postdoctoral Researcher 2016 - 2018

Scientific Computing and Imaging Institute, University of Utah, Salt Lake City, UT

Advisor - **Dr. Valerio Pascucci**

Research Assistant 2009 - 2015

Scientific Computing and Imaging Institute, University of Utah, Salt Lake City, UT

Advisor - **Dr. Valerio Pascucci**

Visiting Scholar Spring 2015

King Abdullah University of Science and Technology

Host - **Dr. David Keyes**

Research Intern Summer 2013

Los Alamos National Laboratory

Advisor - **Dr. James Ahrens**

Research Intern Summer 2012, 2011, 2010

Argonne National Laboratory

Advisor - **Dr. Venkatram Vishwanathan** and **Dr. Phil Carns**

RESEARCH

AWARDS

*PhD Student awards

1. *Best research paper award at IEEE Symposium on Large Data Analysis and Visualization LDAV 2023.
2. *Best research poster, Finalist at Supercomputing Conference 2023.
3. Best Research Poster Award (SRS), at 28th HiPC Student Research Symposium (SRS), 2021.
4. NSF EPSCoR Research Fellow Track 4 2022.
5. Hans Meuer best research paper award at The International Supercomputing Conference, 2020. ISC 2020, acceptance rate - 2.4% (2/81).
6. Best paper award at 26th IEEE International conference on high-performance computing, data, and analytics. HiPC 2019, acceptance rate - 1.1% (2/171).
7. Featured in ALCF's yearly Science magazine for Parallel relational algebra for logical inference at scale.
8. ALCF Director's discretionary award for over five million compute hours on supercomputers at ANL (2018-current).
9. Travel award for NSF Aspiring CSR PIs Workshop, 2019, Washington DC.
10. Best Poster Award- Graduate student poster competition 2014, University of Utah.
11. Travel award for 1st Argonne Training Program on Extreme-Scale Computing (ATPESC 2013).
12. Travel award for IEEE International Conference on Cluster Computing (Cluster 2011).

RESEARCH GRANTS

Active

1. A Full-stack Approach to Declarative Analytics at Scale
NSF PPOSS Large
Role: Principal Investigator, UIC
Total awarded amount: \$2,632,812
My Share at UIC: \$994,717, 2023-2027
URL: https://www.nsf.gov/awardsearch/showAward?AWD_ID=2316157
2. Scalable and Extensible I/O Runtime and Tools for Next Generation Adaptive Data Layouts
NSF Software and Hardware Foundation
Role: Principal Investigator, UIC
My Share at UIC: \$300,165, 2022-2025
URL: https://www.nsf.gov/awardsearch/showAward?AWD_ID=2221811
3. A Full-stack Approach to Declarative Analytics at Scale.
NSF PPOSS Planning.
Role: Principal Investigator, UIC
Total awarded amount: \$166,232
My Share at UIC: \$50,116, 2022-2024.
My Share at UAB: \$33,000, 2021-2022.
URL: https://www.nsf.gov/awardsearch/showAward?AWD_ID=2217036
4. Distributed relational algebra, Directors discretionary award.
Upwards of five million compute hours on Theta, Polaris, and ThetaGPU Supercomputers at Argonne Leadership Computing Facility (ALCF), ANL (2021).
Role: Principal Investigator, 2018-current.
URL: https://www.alcf.anl.gov/sites/default/files/2021-04/ALCF_2020ScienceReport.pdf
5. Relational Algebra on Heterogeneous Extreme-scale Systems.
NSF EPSCoR Research Infrastructure Improvement (RII) Track 4
Role: Principal Investigator, UAB
My share at UAB: \$264,755.00, 2022-2024.
URL: https://www.nsf.gov/awardsearch/showAward?AWD_ID=2132013

Under/Planned submission

1. Enabling Large Scale Scientific Analysis and Visualization on the Web with Progressive Data Streaming, GPU Acceleration, and AI. **NSF CISE Software and Hardware foundation**. *Role: Principal Investigator* Share: \$300,000. Years: 2025-2028. (Under submission)
2. Algorithm-centric Characterization, Modelling and Optimization of Sparse and Irregular Data Movements at Exascale. **NSF Faculty Early Career Development Program (CAREER)**. *Role: Principal Investigator* Share: \$497,996. Years: 2024-2029. (Planned for submission)

Completed

1. Virus Taxonomy: A Community Knowledgebase Supporting Virus Research. NIH U24. *Role: Investigator* My Share: \$60,000. Years: 2021-2023
2. Ease of use and deployment for a fast, scalable data movement infrastructure, **DOE SBIR Phase 1**, award No. DE-SC0017152. PI from The University of Utah. (\$224,927), *as PI from University of Utah*.

Submitted, but not-accepted

1. Algorithm-centric Characterization, Modelling and Optimization of Sparse and Irregular Data Movements at Exascale. **NSF Faculty Early Career Development Program (CAREER)**, *Role: Principal Investigator* Share: \$497,996, Years: 2024-2029. (Submitted at UIC)
2. Quality Controlled Progressive Data Encoding and Streaming for Trustworthy Sharing and Analytics of Petascale Scientific Data. **DOE Data reduction for science**. *Role: Principal Investigator*, Share: \$450,000, Years: 2024-2027. (Submitted at UIC)
3. Collaborative Research: SHF: Medium: A New Generation of GPU-powered High-Performance Analysis and Visualization for the Web. **NSF Software and Hardware Foundation**, *Role: Principal Investigator*, Share: \$600,000, Years: 2024-2028. (Submitted at UIC)
4. Optimization, Characterization, and Modeling of Global and Irregular Data Movements at Exascale. **DOE Early Career preproposal**. *Role: Principal Investigator*, Years: 2024-2019. (Submitted at UIC)
5. CAREER: General-purpose High-Performance Parallel Databases for Exascale Systems. **NSF Career**. *Role: Principal Investigator*, Share: \$447,428. Years: 2023-2028. (Submitted at UAB)
6. Optimizing Computation, Communication, and IO for Iterated Parallel Relational Algebra. **NSF OAC-Advanced Cyberinfrastructure Core 2021**, *Role: Principal Investigator*, Share \$310,000. (Submitted at UAB)
7. Topology-driven cyberinfrastructure for Brain imaging to facilitate Alabama-wide study of mental disorders and associated health disparities. **Alabama Research and Development Enhancement Fund**, *Role: Principal Investigator*. (Submitted at UAB)
8. Accessible tools for leveraging large publicly available datasets to improve reliability of small patient samples. **NIH U24**, *Role: Co-Principal Investigator*. (Submitted at UAB)
9. Engaging in STEM Education through Vertical Integration of Big Data Analytics and Emerging Technologies. **NSF Improving Undergraduate STEM Education (IUSE)**, *Role: Co-Principal Investigator*. (Submitted at UAB)
10. An AI Guided Multimodal Decision Making Framework for Environmental Exposure and Respiratory Response Assessment Connecting People, Data, and Systems. **NSF Smart Health and Biomedical Research in the Era of Artificial Intelligence and Advanced Data Science (SCH)**, *Role: Co-Principal Investigator*, (Submitted at UAB)
11. Acquisition of An Immersive Virtual Reality and Visualization System (VRVS) To Integrate Data and Simulation Intensive Research And Education: Bringing VR to Deep South. **NSF Major Research Instrumentation Program**, *Role: Co-Principal Investigator*. (Submitted at UAB)
12. Parallel relational algebra for logical inferencing at scale. **Incite award, 2020** for compute hours on Argonne National Laboratory's Theta supercomputer, *Role: Principal Investigator* (9 million computer hours). (Submitted at UAB)
13. Parallel Relational Algebra for Logical Inference at Scale **NSF Computer and Information Science and Engineering (CISE) 2019**, *Role: Principal Investigator* (\$499,000). (Submitted at UAB)
14. Parallel relational algebra for logical inferencing at scale. **Incite award, 2019** for compute hours on Argonne National Laboratory's Theta supercomputer, *Role: Principal Investigator* (9 million computer hours). (Submitted at UAB)

15. GeneTerrain: Tools for organizing, searching, and analyzing dynamic knowledge models of cancer. **NIH U24 (2019)**, *Role: Co-Investigator*. (Submitted at UAB)
16. An AI Guided Multimodal Decision Making Framework for Environmental Exposure and Respiratory Response Assessment Connecting People, Data, and Systems **NSF Div Of Information and Intelligent Systems**, *Role: Co-Investigator*. (Submitted at UAB)

PUBLICATIONS

PhD student papers, Papers as an assistant professor*

Refereed Journal papers (J)

- J55*** Landon Richard Dyken, Will Usher, Sidharth Kumar. Interactive Isosurface Visualization in Memory Constrained Environments Using Deep Learning and Speculative Raycasting. *IEEE Transactions on Visualization and Computer Graphics*, 2024.
Impact factor: 4.579
- J54*** Sidharth Kumar, Ahmedur Rahman Shovon, Gopikrishna Deshpande, Sidharth Kumar. The robustness of persistent homology of brain networks to data acquisition-related non-neural variability in resting state fMRI. *Human Brain Mapping*, 2023.
Impact factor: 5.038
- J53*** Kyungmi Lee, Tim Cheongho Lee, Maria Yefimova, Sidharth Kumar, Frank Puga, Andres Azuero, Arif Kamal, Marie A Bakitas, Alexi A Wright, George Demiris, Christine S Ritchie, Carolyn EZ Pickering, J Nicholas Dionne-Odom. Using Digital phenotyping to understand health-related outcomes: A scoping review. *International Journal of Medical Informatics, Elsevier* 2022.
Impact factor: 4.063
- J52*** Darshan Shimoga Chandrashekar, Santhosh Kumar Karthikeyan, Praveen Kumar Korla, Henalben Patel, Ahmedur Rahman Shovon, Mohammad Athar, George J Netto, Zhaohui S Qin, Sidharth Kumar, Upender Manne, Chad J Crieghton, Sooryanarayana Varambally. UALCAN: An update to the integrated cancer data analysis platform. *Neoplasia* 25 (2022): 18-27.
Impact factor: 6.318

Refereed book chapters (B)

- B51** John Edwards, Sidharth Kumar, Valerio Pascucci. Big Data From Scientific Simulations. *Cloud Computing and Big Data, Advances in Parallel Computing, Volume 24*, C. Catlett, W. Gentsch, L. Grandinetti, G. Joubert, J. L. Vazquez-Poletti, Eds. IOS Press, 2013, pages 212-230.
- B50** Valerio Pascucci, Giorgio Scorzelli, Brian Summa, Peer-T. Bremer, Attila Gyulassy, Cameron Christensen, Sidharth Kumar. Scalable visualization and interactive analysis using massive data streams. *Cloud Computing and Big Data, Advances in Parallel Computing, Volume 23*, C. Catlett, W. Gentsch, L. Grandinetti, G. Joubert, J. L. Vazquez-Poletti, Eds. IOS Press, 2013, pages 212-230.
- B49** Valerio Pascucci, Giorgio Scorzelli, Brian Summa, Peer-T. Bremer, Attila Gyulassy, Cameron Christensen, Sujin Philip, Sidharth Kumar. The ViSUS visualization framework. In *High-Performance Visualization: Enabling Extreme-Scale Scientific Insight*, E. W. Bethel, H. Childs, C. Hansen, Eds. Chapman & Hall/CRC Computational Science, 2012.

Refereed conference papers (C) and Refereed workshop papers (W)

- C48*** Ke Fan, Steve Petruzza, Thomas Gilray, Sidharth Kumar. Configurable Algorithms for All-to-all Collectives. *International Supercomputing Conference ISC 2024*, acceptance rate - 28%
- C47*** Ke Fan, Suraj Kesavan, Steve Petruzza, Sidharth Kumar. TinyProf: Towards Continuous Performance Introspection through Scalable Parallel I/O. *International Supercomputing Conference ISC 2024*, acceptance rate - 28%
- C46*** John-Paul Robinson, Ke Fan, Steve Petruzza, Thomas Gilray, Sidharth Kumar. Analysis of MPI Communication Time for Distribution of Repartitioned Data. *Human powered computing PEARC 2024*

- C45*** Andres Sewell, Ke Fan, Ahmedur Rahman Shovon, Landon Dyken, Sidharth Kumar, Steve Petruzza. Bruck Algorithm Performance Analysis for Multi-GPU All-to-All Communication. *2024 International Conference on High-Performance Computing in Asia-Pacific Region HPC Asia 2024*
- C44*** Will Usher, Landon Dyken, Sidharth Kumar. Speculative Progressive Raycasting for Memory Constrained Isosurface Visualization of Massive Volumes. *The 13th IEEE Symposium on Large Data Analysis and Visualization*.
LDAV 2023, acceptance rate - 38%
Best research paper award
- C43*** Ahmedur Rahman Shovon, Thomas Gilray, Kristopher Micinski, Sidharth Kumar. Towards iterated relational algebra on the GPU. *2023 USENIX Annual Technical Conference*
USENIX 2023, acceptance rate - 19%
- W42*** Kashyap Balakavi, Rushitha Janga, Ahmedur Rahman Shovon, Don Dempsey, Elliot Lefkowitz, Sidharth Kumar. Scalable, interactive and hierarchical visualization of virus taxonomic data. *Workshop on Visual Analytics in Healthcare, co-located with IEEE Vis*.
VAHC 2023
- C41*** Yihao Sun, Sidharth Kumar, Thomas Gilray, Kristopher Micinski. Communication-Avoiding Recursive Aggregation. *IEEE International Conference on Cluster Computing*.
Cluster 2023, acceptance rate - 24%
- C40*** Landon Dyken, Pravin Poudel, Steve Petruzza, Will Usher, Jake Chen, Sidharth Kumar. GraphWaGu: GPU Powered Large Scale Graph Layout Computation and Rendering for the Web. *Eurographics Symposium on Parallel Graphics and Visualization*.
EGPGV 2022, acceptance rate - 58%
- C39*** Ke Fan, Thomas Gilray, Kristopher Micinski, Valerio Pascucci, Sidharth Kumar. Optimizing the Bruck algorithm for non-uniform all-to-all communication. *ACM International Symposium on High-Performance Parallel and Distributed Computing*.
HPDC 2022, acceptance rate - 19%
- W38*** Nick Nettekville, Ke Fan, Sidharth Kumar, Thomas Gilray. A Visual Guide to MPI All-to-all. *4th Workshop on Education for High-Performance Computing (held with HiPC)*
EduHiPC 2022
- W37*** Ahmedur Rahman Shovon, Landon Dyken, Oded Green, Thomas Gilray, Sidharth Kumar. Accelerating Datalog applications with cuDF. *12th Workshop on Irregular Applications: Architectures and Algorithms (Held with SC 2022)*
IA3 2022, acceptance rate - 42%
- C36*** Ke Fan, Duong Hoang, Steve Petruzza, Thomas Gilray, Valerio Pascucci, Sidharth Kumar. Load-balancing Parallel I/O of Compressed Hierarchical Layouts. *IEEE Conference On High-Performance Computing, Data, and Analytics*.
HiPC 2021, acceptance rate - 23%
- C35*** Thomas Gilray, Sidharth Kumar, Kristopher Micinski. Compiling Data-parallel Datalog, *International Conference on Compiler Construction, 2021*.
CC 2021
- C34*** Will Usher, Xuan Huang, Steve Petruzza, Sidharth Kumar, Stuart R. Slattery, Sam T. Reeve, Feng Wang, Chris R. Johnson and, Valerio Pascucci. Adaptive Spatially Aware I/O for Multiresolution Particle Data Layouts. *International Parallel and distributed processing symposium*
IPDPS 2021, acceptance rate - 19%
- W33*** Arkaprabha Banerjee, Pratvi Shah, Shivani Nandani, Shantanu Tyagi, Sidharth Kumar, and Bhaskar Chaudhury. An empirical investigation of OpenMP-based implementation of Simplex Algorithm *International Workshop on OpenMP*.
IWOMP 2021
- W32*** Sarthak Patel, Bhrugu Dave, Smit Kumbhani, Mihir Desai, Sidharth Kumar, Bhaskar Chaudhury. Scalable parallel algorithm for fast computation of Transitive Closure of Graphs on Shared Memory Architectures. *2021 IEEE/ACM 6th International Workshop on Extreme Scale Programming Models and Middleware (Held with SC 2022)*.
ESPM2 2021

- W31*** [Ke Fan](#), Kristopher Micinski, Thomas Gilray, Sidharth Kumar. Exploring MPI Collective I/O and File-per-process I/O for Checkpointing a Logical Inference task. *Workshop on High-Performance Storage. HPS 2021*
- C30*** Sidharth Kumar, Thomas Gilray. Load-balancing Parallel Relational Algebra, *The International Supercomputing Conference, 2020*.
ISC 2020, acceptance rate - 31% (27/81)
Hans Meuer Best paper award
- C29*** Sidharth Kumar, Thomas Gilray. Distributed Relational Algebra at Scale. *IEEE Conference On High-Performance Computing, Data, and Analytics*.
HiPC 2019, acceptance rate - 23% (39/171)
Best paper award
- C28*** Sidharth kumar, Will Usher, Steve Petruzza, Valerio Pascucci. Spatially-aware Parallel I/O for Particle Data. *48th International Conference on Parallel Processing*.
ICPP 2019, acceptance rate - 26% (106/405)
- C27** Qi Wu, Will Usher, Steve Petruzza, Sidharth Kumar, Feng Wang, Ingo Wald, Valerio Pascucci and Charles D. Hansen. VisIt-OSPRay: Toward an Exascale Volume Visualization System. *Eurographics Symposium on Parallel Graphics and Visualization (2018)*.
EGPGV 2018, acceptance rate - 43% (10/23)
- C26** Sidharth kumar, Alan Humphrey, Will Usher, Steve Petruzza, Brad Peterson, John A. Schmidt, Derek Harris, Ben Isaac, Jeremy Thornock, Todd Harman, Valerio Pascucci, Martin Berzins. A Case Study for Scaling Task-Based Runtime Systems for Next Generation Engineering Problems. *Supercomputing Asia. SCA 2018*
- C25** Sidharth Kumar, Duong Hoang, Steve Petruzza, John Edwards, Valerio Pascucci. Reducing network congestion and synchronization overhead during aggregation of hierarchical data. *IEEE Conference On High-Performance Computing, Data, and Analytics*.
HiPC 2017, acceptance rate - 23% (42/184)
- W24** Thomas Gilray, Sidharth Kumar. Toward Parallelizing Control-flow Analysis with Datalog. *Scheme and Functional Programming Workshop 2017* (Held in conjunction with 22nd ACM SIGPLAN International Conference on Functional Programming).
SCHEME 2017.
- C23** Aaditya Landge, Ivan Rodero, Sidharth Kumar, Manish Parasar, Valerio Pascucci, Peer-T. Bremer. Evaluation of In-Situ Analysis Strategies at Scale for Power Efficiency and Scalability. *IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing*.
CCGRID 2016, acceptance rate - 20% (40/200)
- C22** Sidharth Kumar, John Edwards, Peer-T. Bremer, Aaron Knoll, Cameron Christensen, Venkatram Vishwanath, Phil Carns, John Schmidt, Valerio Pascucci. Efficient I/O and storage of adaptive resolution data. *IEEE Conference On High-Performance Computing Networking, Storage And Analysis*.
SC 2014, acceptance rate - 20% (82/394)
- C21** Sidharth Kumar, Cameron Christensen, John Schmidt, Peer-T. Bremer, Eric Brugger, Venkatram Vishwanath, Philip Carns, Giorgio Scorzelli, Hemanth Kolla, Ray Grout, Jacqueline Chen, and Valerio Pascucci. Fast multi-resolution reads of massive simulation datasets. *The International Supercomputing Conference. ISC 2014*
- C20** Sidharth Kumar, Avishek Saha, Venkatram Vishwanath, Philip Carns, John Schmidt, Giorgio Scorzelli, Hemanth Kolla, Ray Grout, Robert Latham, Robert Ross, Michael E. Papka, Jacqueline Chen, and Valerio Pascucci. Characterization and modeling of PIDX parallel I/O for performance optimization. *IEEE Conference On High-Performance Computing Networking, Storage And Analysis*.
SC 2013, acceptance rate - 20% (92/457)
- C19** Sidharth Kumar, Venkatram Vishwanath, Phil Carns, Joshua A. Levine, Robert Latham, Giorgio Scorzelli, Robert Ross, Hemanth Kolla, Ray Grout, Jackie Chen, Michael E. Papka, Valerio Pascucci. Efficient data restructuring and aggregation for I/O acceleration in PIDX. *IEEE Conference On High-Performance Computing Networking, Storage And Analysis*.
SC 2012, acceptance rate - 21% (100/472)

- C18** Sidharth Kumar, Venkatram Vishwanath, Phil Carns, Brian Summa, Giorgio Scorzelli, Valerio Pascucci, Robert Ross, Jackie Chen, Hemanth Kolla. PIDX: efficient parallel I/O for multi-resolution multi-dimensional scientific datasets. *IEEE International Conference on Cluster Computing. Cluster 2011, acceptance rate - 27% (39/140)*
- W17** Sidharth Kumar, Valerio Pascucci, Venkatram Vishwanath, Phil Carns, Robert Latham, Tom Peterka, Michael Papka, Robert Ross. Towards parallel access of multi-dimensional, multi-resolution scientific data. *Petascale Data Storage Workshop (Held in conjunction with ACM/IEEE Supercomputing Conference). PDSW 2010.*

Extended abstracts (A)

- A16*** Ke Fan, Sidharth Kumar. Two-phase IO Enabling Large-scale Introspection. SC 2023 Research Poster: Proceedings of the International Conference for High-Performance Computing, Networking, Storage and Analysis, 2023.
Best poster finalist
- A15*** Santhosh Kumar Karthikeyan, Darshan S Chandrashekar, Upender Manne, Chad Creighton, Zhaohui S Qin, Sidharth Kumar, Sooryanaraya Varambally. Mammonc-DB: A web-based, user-friendly tool for comprehensive multi-omics data analysis in breast cancer. Research poster: The American Association for Cancer Research, 2023.
- A14*** Darshan Shimoga Chandrashekar, Sooryanarayana Varambally, Santhosh Kumar Karthikeyan, Praveen Kumar Korla, Henalben Patel, Mohammad Athar, Upender Manne, George J Netto, Ahmedur Rahman Shovon, Sidharth Kumar, Zhaohui S Qin, Chad J Crieghton. Updates to UALCAN, a comprehensive cancer proteogenomic data analysis platform for discovery research. Research poster: The American Association for Cancer Research.
- A13*** Ke Fan, Sidharth Kumar. Generalized Radix-r Bruck Algorithm for All-to-all Communication. SC 2022 Research Poster: Proceedings of the International Conference for High-Performance Computing, Networking, Storage and Analysis, 2022.
- A12*** Vishwa Shah, Shruti Agrawal, Riddhi Tanna, Dishita Thaker, Sidharth Kumar, Bhaskar Chaudhury. Parallel Implementations of Arithmetic Encoding on Shared Memory Systems. Students Research Symposium, HiPC 2021.
Best poster award
- A11*** Ke Fan, Thomas Gilray, Sidharth Kumar. Padding to Extend the Bruck Algorithm for Non-uniform All-to-all Communication. SC 2021 Research Poster: Proceedings of the International Conference for High-Performance Computing, Networking, Storage and Analysis, 2021.
- A10** Sidharth Kumar, Steve Petruzza, Duong Hoang, Valerio Pascucci. Accelerating In-situ Analysis and Visualization with Subsampling. In ACM Symposium on High-Performance Parallel and Distributed Computing, (HPDC 2017).
- A9** Xiao Xu, Bok Jik Lee, Ramanan Sankaran, Steve Petruzza, Sidharth Kumar, Giorgio Scorzelli, Francisco E. Hernaandez Perez, Valerio Pascucci, Hong G. Im. KARFS/Direct Numerical Simulation and PIDX I/O and Visualization of Dimethylether (DME)/Air Turbulent Mixture Auto-ignition with Thermal Stratification. High Performance Computing Saudi Arabia, 2017.
- A8** Sidharth Kumar, Valerio Pascucci. Multi-resolution I/O for Massive Simulations: Enabling Scalable Visualization and Processing. Graduate student poster competition 2014, University of Utah.
Best poster award

Manuscripts Under Review/Revision (M)

- M7** David Rubey, Darshan Shimoga Chandrashekar, Santhosh Kumar Karthikeyan, Ahmedur Rahman Shovon, Sooryanarayana Varambally, Sidharth Kumar. UALCAN MOBILE: An app for cancer gene expression data analysis.
- M6** Ke Fan, Jens Domke, Sidharth Kumar. An ensemble of novel parameterized algorithms for non-uniform all-to-all collectives.
- M5** Ahmedur Rahman Shovon, Gopikrishna Deshpande, Sidharth Kumar. Topology-assisted clustering of temporal fMRI brain networks with use-case in mitigating non-neural multi-site variability.

- M4** Thomas Gilray, Arash Sahebolamri, Sidharth Kumar, Kristopher Micinski. Higher-order, Data-Parallel Structured Deduction.
- M3** Landon Dyken, Pravin Poudel, Saugat Adhikari, Da Yan, Steve Petruzza, Will Usher, Sidharth Kumar. Elevation-Aware Region Annotation in the Web.
- M2** John-Paul Robinson, Steve Petruzza, Suren Byna, Ke Fan, Thomas Gilray, Sidharth Kumar. I/O Performance Analysis of Non-uniform Data in Parallel Applications.
- M1** Yihao Sun, Ahmedur Rahman Shovon, Thomas Gilray, Kristopher Micinski, Sidharth Kumar. Datalog on the GPU.

TALKS

Invited talks

1. Parallel Deductive Databases June 2022
Data science group, Argonne National Laboratory
2. Optimizing Bruck algorithm for all-to-all communication July 2022
MPICH Group, Argonne National Laboratory
3. Revisiting a Classic: Bruck algorithm for non-uniform all-to-all communication November 2021
Center for Understandable, Performant Exascale Communication Systems (CUP-ECS)
4. Parallel I/O and Data-parallel Deductive databases in the age of HPC January 2021
DAHCT Alumni Series, Gandhinagar, Gujarat, India
5. Parallel Relational Algebra for Logical Inference at Scale February 2020
NSF Smart cyberinfrastructure workshop, Washington DC, USA
6. Data movement for massive scientific applications January 2020
Indian Institute of Technology, Ropar, India
7. Data movement for massive scientific applications August 2019
University of Alabama Tuscaloosa, Alabama
8. Scientific Data processing at Scale August 2019
Annual Alabama Advanced Imaging Consortium Imaging Retreat
9. Foray into the world of Supercomputers September 2019
Indian Springs School, Birmingham, Alabama
10. PIDX: Scalable parallel Multiresolution I/O at extreme scale August 2017
Department of Computer Science, University of Maryland, College Park
11. Exploring the landscape of Supercomputing with focus on High-Performance I/O Systems January 2017
Indian Institute of Technology (IIT) Patna, India
12. PIDX: Scalable parallel Multiresolution I/O at extreme scale November 2016
Prace booth, Supercomputing conference
13. A Scalable and Tunable Adaptive Resolution Parallel I/O Framework August 2016
Symposium at Idaho Academy of science and Engineering, Idaho State University
14. High-Performance Multi-resolution I/O: Achieving Fast Simulation Checkpointing while Enabling Interactive Data Access for Analytics and Visualization December 2014
Extreme Computing Research Center, King Abdullah University of Science and Technology, Saudi Arabia

Contributed conference talks

1. Configurable Algorithms for All-to-all Collectives. May 2024
International Supercomputing Conference (ISC) Technical Paper Presentation.
2. TinyProf: Towards Continuous Performance Introspection through Scalable Parallel I/O. May 2024
International Supercomputing Conference (ISC) Technical Paper Presentation.
3. A Visual Guide to MPI all to all December 2022
Workshop on Education for High-Performance Computing (eduHiPC) Technical Paper Presentation.
4. Load balancing Parallel Relational Algebra June 2020
International Supercomputing Conference (ISC) Technical Paper Presentation.

5. Distributed Relational Algebra at Scale November 2019
IEEE Conference On High-Performance Computing, Data, and Analytics Technical Paper Presentation
6. Reducing network congestion and synchronization during aggregation of hierarchical data November 2017
IEEE Conference On High-Performance Computing, Data, and Analytics Technical Paper Presentation
7. Case Study for scaling task-based runtime systems for next-generation engineering problems. May 2018
Supercomputing Asia (SCA 2018)
8. Efficient I/O and storage of adaptive-resolution data November 2014
IEEE/ACM Supercomputing Conference (SC) Technical Paper Presentation
9. Fast multi-resolution reads of massive simulation datasets June 2014
International Supercomputing Conference (ISC) Technical Paper Presentation.
10. Characterization and modeling of PIDX parallel I/O for performance optimization November 2013
IEEE/ACM Supercomputing Conference (SC) Technical Paper Presentation
11. Efficient data restructuring and aggregation for I/O acceleration in PIDX November 2012
IEEE/ACM Supercomputing Conference (SC) Technical Paper Presentation
12. PIDX: efficient parallel I/O for multi-resolution multi-dimensional scientific datasets November 2011
IEEE International Conference on Cluster Computing Technical Paper Presentation
13. Towards parallel access of multi-dimensional, multi-resolution scientific data November 2010
Petascale Data Storage Workshop (PDSW 2010) Technical Paper Presentation

TEACHING

TEACHING EXPERIENCE

Courses taught at University of Illinois at Chicago

1. Database Systems, Fall 2023, **Student course evaluation score of 4.4/5** (both under-graduate and graduate students)

Courses taught at University of Alabama at Birmingham

1. Data structures and algorithms, Spring 2020, 2021. **Idea evaluation score of 4.4/5 and 4.1/5** (under-graduate class)
2. Advanced Algorithms and apps; Fall 2018, 2019. **Idea evaluation score of 4.6/5 and 4.2/5** (graduate class)
3. Data visualization; Spring 2018, 2019. **Idea evaluation score of 4.1, 4.4** (Self-developed graduate course)

Course taught at University of Utah

1. Parallel Computing (CS 6230), (co-instructed with Prof. Martin Berzins) Fall 2017

TUTORIALS ORGANIZED

1. **Managing Extreme Datasets for HPC** Spring 2016
Lawrence Livermore National Laboratory, 2016
Valerio Pascucci, **Sidharth Kumar**, Cameron Christensen, Steve Petruzza, Duong Hoang.
2. **Big Scientific Data Made Simple** Summer 2015
King Abdullah University of Science and technology, 2015
Valerio Pascucci, Steve Petruzza, Bilel Hadri, **Sidharth Kumar**, Duong Hoang and Peer-Time Bremer.

MENTORING

PhD students

1. *Ke Fan* (started Fall 2019)
2. *Ahmedur Rahman Shovon* (started Fall 2021)
3. *Landon Dyken* (started Fall 2022)
4. *John-Paul Robinson* (started Spring 2022)
5. *Christopher Hulu Charis* (starting Fall 2024)
6. *Jiaxin Lu* (starting Fall 2024)
7. *Kunting Qi* (starting Fall 2024)

Masters students

1. *Kashyap Balakavi* (performance modeling of MPI collectives)
2. *Sri Rushitha Janga* (hierarchical visualisation for virus taxonomy)
3. *Nick Netterville* (developing education tools)
4. *Gopichand Pulli* (developing mobile app for cancer data)
5. *David Rubey* (developing mobile app for cancer data)
6. *Feng Zhang* (Placed at Twitch)
7. *Jiaewei Wang* (Placed at Amazon)

Member of PhD Committee

1. *Mirza Tanzin Sami* (Department of computer Science)
2. *Aminul Hoque* (Department of computer Science)
3. *Raiful Hasan* (Department of computer Science)
4. *Jalal Khalil* (Department of computer Science)
5. *Guimu Gu* (Department of computer Science)
6. *Anuradha Mandal* (Department of computer Science)
7. *Kyungmee Lee* (Department of nursing)
8. *Md Yasser Karim* (Department of computer Science)
9. *Walker Haddock* (Department of computer Science)
10. *Jubur Mohammed* (Department of computer Science)

SERVICE

EXTERNAL SERVICE ROLES

Program Vice-chair

1. Applications track at HiPC 2024.

Panel Discussion

1. Challenges of Early Adoption of PDC/HPC Education in Computing and Engineering Curriculum in India and Developing Countries.

Program Committee

1. International Parallel and Distributed Processing Symposium (IPDPS 2023), Technical program committee member
2. ACM/IEEE Supercomputing Conference (SC 2023), technical program committee member
3. International Supercomputing Conference (ISC 2023), workshops committee member
4. International Supercomputing Conference (ISC 2023), technical program committee member
5. IEEE International Parallel and Distributed Processing Symposium (IPDPS 2022), PC Chairs Team
6. 29th IEEE International Conference on high-performance Computing, data, and analytics (HiPC), technical program committee member
7. ACM/IEEE Supercomputing Conference (SC 2022), technical program committee member
8. ACM/IEEE Supercomputing Conference (SC 2022), tutorials committee member
9. EuroMPI/USA 2022, technical program committee member
10. 28th IEEE International Conference on high-performance computing, data, and analytics (HiPC), technical program committee member
11. ACM/IEEE Supercomputing Conference (SC 2021), tutorials committee member
12. 27th IEEE International Conference on high-performance computing, data, and analytics (HiPC), technical program committee member
13. 19th International Conference on Algorithms and Architectures for Parallel Processing (ICA3PP), technical program committee member
14. 4th International Parallel Data Systems Workshop (PDSW), technical program committee member

Proposal Panelist

1. Two NSF review panels since 2018

Journal Reviewer

1. ACM Transactions on Storage
2. IEEE/ACM Transactions on Computational Biology and Bioinformatics