

Kamesh Madduri

The Pennsylvania State University
W312 Westgate Building, University Park, PA 16802
madduri@cse.psu.edu madduri.org +1 814-865-0883

Education

Georgia Institute of Technology, Atlanta, GA
Ph.D., Computer Science, 2008.
Specialization: Computational Science and Engineering.
Advisor: Prof. David A. Bader.

Indian Institute of Technology Madras, Chennai, India
B.Tech., Electrical Engineering, 2004.

Employment

The Pennsylvania State University, University Park, PA
College of Engineering, School of Electrical Engineering and Computer Science
Associate Professor with tenure, July 2017 – Present.
Assistant Professor (tenure-track), August 2011 – June 2017.
Graduate faculty of Computer Science and Engineering, Bioinformatics and Genomics,
Social Data Analytics, and Electrical Engineering. Member of the Center for Computa-
tional Biology and Bioinformatics and the Center for Social Data Analytics. Affiliated
with Huck Institutes of the Life Sciences and the Institute for CyberScience.

Lawrence Berkeley National Laboratory, Berkeley, CA
Guest Scientist, July 2011 – Present.
Research Scientist, August 2010 – July 2011.
Luis W. Alvarez Postdoctoral Fellow, August 2008 – August 2010.

Research Interests

Data Science, High-performance computing, Computational Biology, Graph analytics,
Computational Science applications, Transportation data analytics. Please visit the
Scalable Computing Lab website (scalablecomputinglab.org) for more details.

Honors and Awards

1. Paper titled “Design and Implementation of the HPCS Graph Analysis Benchmark on Symmetric Multiprocessors,” co-authored with David Bader and published at the 2005 IEEE International Conference on High-Performance Computing (HiPC), selected as one of seven impactful papers in 25 years of HiPC, 2018.

2. Best Artifact Award, 7th Workshop on Irregular Applications: Algorithms and Architectures (IA³), 2017.
3. Student Innovation Award, IEEE HPEC Graph Challenge, 2017.
4. Best Paper Award, 42nd International Conference on Parallel Processing (ICPP), 2013.
5. National Science Foundation (NSF) CAREER Award, 2013.
6. SIAM Activity Group on Supercomputing Junior Scientist Prize, 2010.
7. Luis W. Alvarez Postdoctoral Fellowship in Computational Science, Lawrence Berkeley National Laboratory, 2008 – 2010.
8. Outstanding Graduate Research Assistant award, College of Computing, Georgia Institute of Technology, 2008.
9. Best Poster Award, IEEE Technical Committee for Parallel Processing (TCPP) PhD Forum held at the 22nd IEEE International Parallel and Distributed Processing Symposium (IPDPS), 2008.
10. ACM/IEEE-CS High Performance Computing Ph.D. Fellowship Honorable Mention, 2007.
11. NASA Graduate Student Researchers Program (GSRP) Fellowship, 2006 – 2008.
12. NSF Graduate Research Fellowship Program (GRFP) Honorable Mention, 2005.

Current and Recent Research Grants

1. NSF, Office of Advanced Cyberinfrastructure, CAREER: Algorithmic and Software Foundations for Large-scale Graph Analysis, \$500,000, May 2013 – June 2019, PI.
2. NSF, Division of Computing and Communication Foundations, XPS: FULL: DSD: End-to-end Acceleration of Genomic Workflows on Emerging Heterogeneous Supercomputers, \$849,984, Sep. 2014 – Feb. 2020, PI, co-PIs: Kandemir, Medvedev.
3. Penn State, College of Engineering Multidisciplinary Research Seed Grant, High Performance Computing for Network Traffic Safety, \$25,000, July 2016 – July 2017, PI, co-PI: Shankar.
4. Microsoft, Azure Data Science Research, Nov. 2017 – Nov. 2018, Cloud-enabled Hydrology Data Portals and Workflows, \$20,000 (Azure credit), PI, co-PI: Leonard.

Publications

According to Elsevier Scopus, my peer-reviewed research papers have received over 1,200 citations with an *h*-index of 18. According to Google Scholar, my research artifacts have been cited over 2,950 times with an *h*-index of 27.

Refereed Journal Publications

1. B. Wang, S. Ethier, W. Tang, K. Ibrahim, K. Madduri, S. Williams, and L. Oliker, “Modern Gyrokinetic Particle-in-Cell Simulation of Fusion Plasmas on Top Supercomputers,” *SAGE Int’l. Journal of High Performance Computing Applications (IJHPCA)* 33(1):169–188, 2019.
2. H. Zhan, G. Gomes, X. S. Li, K. Madduri, A. Sim, and K. Wu, “Consensus Ensemble System for Traffic Flow Prediction,” *IEEE Trans. Intelligent Trans. Syst.* 19(12):3903–3914, 2018.
3. L. Leonard, A. M. MacEachren, and K. Madduri, “Graph-based Visual Analysis for Large-scale Hydrological Modeling,” *SAGE Information Visualization (InfoVis)* 16(3):205–216, 2017.
4. G. M. Slota, K. Madduri, and S. Rajamanickam, “Complex Network Partitioning using Label Propagation,” *SIAM Journal on Scientific Computing (SISC)* 38(5):S620–S645, 2016.
5. L. Leonard, K. Madduri, and C. Duffy, “Tuning Heterogeneous Computing Platforms for Large-scale Hydrology Data Management,” *IEEE Trans. Parallel Distrib. Syst.* 27(9):2753–65, 2016.
6. G. M. Slota and K. Madduri, “Parallel color-coding,” *Elsevier Parallel Computing* 47:51–69, 2015.
7. K. Z. Ibrahim, K. Madduri, S. Williams, B. Wang, S. Ethier, and L. Oliker, “Analysis and optimization of gyrokinetic toroidal simulations on homogenous and heterogenous platforms,” *SAGE Int’l. Journal of High Performance Computing Applications (IJHPCA)* 27(4):454–473, 2013.
8. K. Madduri, J. Su, S. Williams, L. Oliker, S. Ethier, and K. Yelick, “Optimization of Parallel Particle-to-Grid Interpolation on Leading Multicore Platforms,” *IEEE Trans. Parallel Distrib. Syst.* 23(10):1915–1922, 2012.
9. K. Madduri, E.-J. Im, K. Z. Ibrahim, S. Williams, S. Ethier, and L. Oliker, “Gyrokinetic particle-in-cell optimization on emerging multi- and manycore platforms,” *Parallel Computing* 37(9):501–520, 2011.
10. K. Subramani and K. Madduri, “Two-level heaps: a new priority queue structure with applications to the single source shortest path problem,” *Springer Computing* 90(3-4):113–130, 2010.
11. J. Orlin, K. Madduri, K. Subramani, and M. Williamson, “A faster algorithm for the single source shortest path problem with few distinct positive lengths,” *Elsevier Journal of Discrete Algorithms* 8(2):189–198, 2010.

12. K. Subramani, C. Tauras, and K. Madduri, "Space-time tradeoffs in negative cycle detection - An empirical analysis of the Stressing algorithm," *Elsevier Applied Mathematics and Computation* 215(10):3563–3575, 2010.
13. D. A. Bader and K. Madduri, "A graph-theoretic analysis of the human protein-interaction network using multicore parallel algorithms," *Elsevier Parallel Computing* 34(11):627–639, 2008.
14. K. Subramani and K. Madduri, "A Randomized Queueless Algorithm for Breadth-First Search," *Int'l. Journal of Computers and their Applications* 15(3):177–186, 2008.
15. D. A. Bader, V. Agarwal, K. Madduri, and S. Kang, "High performance combinatorial algorithm design on the Cell Broadband Engine processor," *Elsevier Parallel Computing* 33(10-11):720–740, 2007.
16. D. A. Bader, K. Madduri, J. R. Gilbert, V. Shah, J. Kepner, T. Meuse, and A. Krishnamurthy, "Designing Scalable Synthetic Compact Applications for Benchmarking High Productivity Computing Systems," *CTWatch Quarterly* 2(4B):41–51, 2006.

Refereed Conference and Workshop Publications

17. V. Rengasamy, M. Kandemir, P. Medvedev, and K. Madduri, "Parallel Read Partitioning for Concurrent Assembly of Metagenomic Data," in *Proc. IEEE Int'l. Conf. on High Performance Computing, Data, and Analytics (HiPC)*, Dec. 2018.
18. H. Zhan, G. Gomes, X. S. Li, K. Madduri, and K. Wu, "Efficient Online Hyperparameter Learning for Traffic Flow Prediction," in *Proc. 21st IEEE Int'l. Conf. on Intelligent Transportation Systems (ITSC)*, pp. 164–169, Nov. 2018.
19. S. Kirmani and K. Madduri, "Spectral Graph Drawing: Building Blocks and Performance Analysis," in *Proc. Workshop on Graph Algorithms Building Blocks (GABB)*, May 2018.
20. H. Kabir and K. Madduri, "Shared-memory Graph Truss Decomposition," in *Proc. IEEE Int'l. Conf. on High Performance Computing, Data, and Analytics (HiPC)*, pp. 13–22, Dec. 2017.
21. V. Rengasamy, T.-Y. Fu, W.-C. Lee, and K. Madduri, "Optimizing Word2Vec Performance on Multicore Systems," in *Proc. Workshop on Irregular Applications: Architectures and Algorithms (IA³)*, Nov. 2017.
22. J. Kotra, S. Kim, K. Madduri, and M. T. Kandemir, "Congestion-Aware Memory Management on NUMA Platforms: A VMware ESXi case study," in *Proc. IEEE Int'l. Symp. on Workload Characterization (IISWC)*, Oct. 2017.
23. H. Kabir and K. Madduri, "Parallel k-truss Decomposition on Multicore Systems," in *Proc. IEEE High Performance Extreme Computing (HPEC)*, Sep. 2017.

24. S. Parimalarangan, G. M. Slota, and K. Madduri, "Fast Parallel Graph Triad Census and Triangle Counting on Shared-memory Platforms," in *Proc. 2nd IEEE Workshop on Parallel and Distributed Processing for Computational Social Systems (ParSocial)*, Jun. 2017.
25. H. Kabir and K. Madduri, "Parallel k-core Decomposition on Multicore Platforms," in *Proc. 2nd IEEE Workshop on Parallel and Distributed Processing for Computational Social Systems (ParSocial)*, Jun. 2017.
26. R. Narayanan and K. Madduri, "Parallel Particle-in-Cell Performance Optimization: A Case Study of Electrospray Simulation," in *Proc. 18th IEEE Int'l. Workshop on Parallel and Distributed Scientific and Engineering Computing (PDSEC)*, Jun. 2017.
27. H. Zhan and K. Madduri, "Analyzing Community Structure in Networks," in *Proc. 1st Workshop on the Intersection of Graph Algorithms and Machine Learning (GraML)*, Jun. 2017.
28. G. M. Slota, S. Rajamanickam, K. Madduri, and K. Devine, "Partitioning Trillion-edge Graphs in Minutes," in *Proc. 31st IEEE Int'l. Parallel and Distributed Processing Symposium (IPDPS)*, May 2017.
29. V. Rengasamy, P. Medvedev, and K. Madduri, "Parallel and Memory-efficient Preprocessing for Metagenome Assembly," in *Proc. 16th IEEE Int'l. Workshop on High Performance Computational Biology (HiCOMB)*, May 2017.
30. G. M. Slota, S. Rajamanickam, and K. Madduri, "Order or Shuffle: Empirically Evaluating Vertex Order Impact on Parallel Graph Computations," in *Proc. Workshop on Graph Algorithms Building Blocks (GABB)*, May 2017.
31. W. Tang, B. Wang, S. Ethier, G. Kwasniewski, T. Hoefler, K. Ibrahim, K. Madduri, S. Williams, L. Oliker, C. Rosales-Fernandez, and T. Williams, "Extreme Scale Plasma Turbulence Simulations on Top Supercomputers Worldwide," in *Proc. ACM/IEEE Conf. on Supercomputing (SC)*, Nov. 2016.
32. V. Rengasamy and K. Madduri, "SPRITE: A Fast Parallel SNP Detection Pipeline," in *Proc. 31st Int'l. Conf. ISC High Performance*, Jun. 2016.
33. G. M. Slota, S. Rajamanickam, and K. Madduri, "A Case Study of Complex Graph Analysis in Distributed Memory: Implementation and Optimization," in *Proc. 30th IEEE Int'l. Parallel and Distributed Processing Symposium (IPDPS)*, May 2016.
34. H. Zhan and K. Madduri, "GSK: Graph sparsification as a knapsack problem formulation," in *Proc. 3rd SDM Workshop on Mining Networks and Graphs (MNG)*, May 2016.
35. L. Leonard, K. Madduri, and C. J. Duffy, "Graph-based analysis for large-scale hydrological modeling," in *Proc. IEEE VIS Exploring Graphs at Scale (EGAS) Workshop*, Oct. 2015.

36. G. M. Slota, S. Rajamanickam, and K. Madduri, "High-Performance Graph Analytics on Manycore Processors," in *Proc. 29th IEEE Int'l. Parallel and Distributed Processing Symposium (IPDPS)*, pp. 17–27, May 2015.
37. G. M. Slota and K. Madduri, "Simple Parallel Biconnectivity Algorithms for Multicore Platforms," in *Proc. 20th IEEE Int'l. Conf. on High Performance Computing (HiPC)*, pp. 1–10, Dec. 2014.
38. G. M. Slota, K. Madduri, and S. Rajamanickam, "PuLP: Scalable Multi-Objective Multi-Constraint Partitioning for Small-World Networks," in *Proc. 2nd IEEE Int'l. Conf. on Big Data (BigData)*, pp. 481–490, Oct. 2014.
39. G. M. Slota, S. Rajamanickam, and K. Madduri, "BFS and Coloring-based Parallel Algorithms for Strongly Connected Components and Related Problems," in *Proc. 28th IEEE Int'l. Parallel and Distributed Processing Symposium (IPDPS)*, pp. 550–559, May 2014.
40. G. M. Slota and K. Madduri, "Complex network analysis using parallel approximate motif counting," in *Proc. 28th IEEE Int'l. Parallel and Distributed Processing Symposium (IPDPS)*, pp. 405–414, May 2014.
41. J. Choi, A. Chandramowlishwaran, K. Madduri, and R. Vuduc, "A CPU-GPU Hybrid Implementation and Model-Driven Scheduling of the Fast Multipole Method," in *Proc. 7th Workshop on General Purpose Processing using GPUs (GPGPU)*, pp. 64:1–64:8, Mar. 2014.
42. B. Wang, S. Ethier, W. Tang, T. Williams, K. Ibrahim, K. Madduri, S. Williams, and L. Oliker, "Kinetic Turbulence Simulations at Extreme Scale on Leadership-Class Systems," in *Proc. ACM/IEEE Conf. on Supercomputing (SC)*, pp. 82:1–82:12, Nov. 2013.
43. G. Slota and K. Madduri, "Fast Approximate Subgraph Counting and Enumeration," in *Proc. 42nd Int'l. Conf. on Parallel Processing (ICPP)*, pp. 210–219, Oct. 2013.
44. D. Hadka, P. Reed, and K. Madduri, "Scalability Analysis of the Asynchronous, Master-slave Multiobjective Evolutionary Algorithm," in *Proc. 16th Int'l. Workshop on Nature Inspired Distributed Computing (NIDISC)*, pp. 425–434, May 2013.
45. M. Frasca, K. Madduri, and P. Raghavan, "NUMA-aware graph mining techniques for performance and energy efficiency," in *Proc. ACM/IEEE Conf. on Supercomputing (SC)*, pp. 95:1–95:11, Nov. 2012.
46. A. Chandramowlishwaran, J. Choi, K. Madduri, and R. W. Vuduc, "Brief announcement: Towards a Communication optimal Fast Multipole Method and its implications at Exascale," in *Proc. 24th ACM Symp. on Parallelism in Algorithms and Architectures (SPAA)*, pp. 182–184, Jun. 2012.
47. A. Buluç and K. Madduri, "Graph partitioning for scalable distributed graph computations," in *Proc. 10th DIMACS Implementation Challenge Workshop – Graph Partitioning and Graph Clustering*, Feb. 2012.

48. K. Madduri, K. Z. Ibrahim, S. Williams, E.-J. Im, S. Ethier, J. Shalf, and L. Oliker, "Gyrokinetic Toroidal Simulations on leading multi- and manycore HPC systems," in *Proc. Conf. on High Performance Computing, Networking, Storage and Analysis (SC)*, p. 23, Nov. 2011.
49. A. Buluç and K. Madduri, "Parallel breadth-first search on distributed memory systems," in *Proc. Conf. on High Performance Computing, Networking, Storage and Analysis (SC)*, p. 65, Nov. 2011.
50. K. Madduri and K. Wu, "Massive-Scale RDF Processing using Compressed Bitmap Indexes," in J. B. Cushing, J. C. French, and S. Bowers (editors), *Proc. 23rd Int'l. Conf. on Scientific and Statistical Database Management (SSDBM)*, vol. 6809 of LNCS, pp. 470–479, Springer, Jul. 2011.
51. R. Sudarsan, J. Borrill, C. Cantalupo, T. Kisner, K. Madduri, L. Oliker, Y. Zheng, and H. Simon, "Cosmic microwave background map-making at the petascale and beyond," in *Proc. 25th ACM Int'l. Conf. on Supercomputing (ICS)*, pp. 305–316, May-June 2011.
52. A. Chandramowlishwaran, K. Madduri, and R. Vuduc, "Diagnosis, tuning, and redesign for multicore performance: A case study of the Fast Multipole Method," in *Proc. Conf. on High Performance Computing, Networking, Storage and Analysis (SC)*, pp. 1–12, Nov. 2010.
53. E. Strohmaier, S. Williams, A. Kaiser, K. Madduri, K. Ibrahim, D. Bailey, and J. W. Demmel, "A Kernel Testbed for Parallel Architecture, Language, and Performance Research," in *Proc. 8th Int'l. Conf. of Numerical Analysis and Applied Mathematics (ICNAAM)*, vol. 1281 of AIP Conference Proceedings, pp. 1297–1300, Sep. 2010.
54. K. Wu, K. Madduri, and S. Canon, "Multi-level bitmap indexes for flash memory storage," in *Proc. 14th Int'l. Database Engineering & Applications Symposium (IDEAS)*, pp. 114–116, Aug. 2010.
55. A. Kaiser, S. Williams, K. Madduri, K. Ibrahim, D. H. Bailey, J. Demmel, and E. Strohmaier, "A case for a testbed of kernels for software/hardware co-design research," in *Proc. 2nd USENIX Workshop on Hot Topics in Parallelism (HotPar)*, Jun. 2010.
56. K. Madduri, S. Williams, S. Ethier, L. Oliker, J. Shalf, E. Strohmaier, and K. Yelick, "Memory-efficient Optimization of Gyrokinetic Particle-to-Grid Interpolation for Multi-core Processors," in *Proc. ACM/IEEE Conf. on High Performance Computing (SC)*, Nov. 2009.
57. K. Madduri and K. Wu, "Efficient Joins with Compressed Bitmap Indices," in *Proc. 18th ACM Conf. on Information and Knowledge Management (CIKM)*, pp. 1017–1026, Nov. 2009.
58. X. Gu, K. Madduri, K. Subramani, and H.-J. Lai, "Improved Algorithms for Detecting Negative Cost Cycles in Undirected Graphs," in X. Deng, J. Hopcroft, and J. Xue (edi-

- tors), *Proc. 3rd Int'l. Frontiers of Algorithmics Workshop (FAW)*, vol. 5598 of LNCS, pp. 40–50, Springer, Jun. 2009.
59. K. Subramani and K. Madduri, “Two-level heaps: a new priority queue structure with applications to the single source shortest path problem,” in D.-Z. Du, X. Hu, and P. M. Pardalos (editors), *Proc. 3rd Int'l. Conf. on Combinatorial Optimization and Applications (COCOA)*, vol. 5573 of LNCS, pp. 186–196, Springer, Jun. 2009.
 60. K. Madduri, D. Ediger, K. Jiang, D. A. Bader, and D. Chavarria-Miranda, “A Faster Parallel Algorithm and Efficient Multithreaded Implementations for Evaluating Betweenness Centrality on Massive Datasets,” in *Proc. 3rd Workshop on Multithreaded Architectures and Applications (MTAAP)*, May 2009.
 61. K. Madduri and D. A. Bader, “Compact Graph Representations and Parallel Connectivity Algorithms for Massive Dynamic Network Analysis,” in *Proc. 23rd IEEE Int'l. Parallel and Distributed Processing Symposium (IPDPS)*, May 2009.
 62. D. A. Bader and K. Madduri, “SNAP: Small-world Network Analysis and Partitioning: an open-source parallel graph framework for the exploration of large-scale networks,” in *Proc. 22nd IEEE Int'l. Parallel and Distributed Processing Symposium (IPDPS)*, Apr. 2008.
 63. K. Subramani and K. Madduri, “Accomplishing Approximate FCFS fairness without queues,” in S. Aluru, M. Parashar, R. Badrinath, and V. K. Prasanna (editors), *Proc. 14th Int'l. Conf. on High Performance Computing (HiPC)*, vol. 4873 of LNCS, pp. 540–551, Springer, Dec. 2007.
 64. D. A. Bader, S. Kintali, K. Madduri, and M. Mihail, “Approximating Betweenness Centrality,” in A. Bonato and F. R. K. Chung (editors), *Proc. 5th Int'l. Workshop on Algorithms and Models for the Web-Graph (WAW)*, vol. 4863 of LNCS, pp. 124–137, Springer, Dec. 2007.
 65. J. R. Crobak, J. Berry, K. Madduri, and D. A. Bader, “Advanced Shortest Paths Algorithms on a Massively-Multithreaded Architecture,” in *Proc. 1st Workshop on Multithreaded Architectures and Applications (MTAAP)*, Mar. 2007.
 66. D. A. Bader, V. Kanade, and K. Madduri, “SWARM: A Parallel Programming Framework for Multicore Processors,” in *Proc. 1st Workshop on Multithreaded Architectures and Applications (MTAAP)*, Mar. 2007.
 67. D. A. Bader, V. Agarwal, and K. Madduri, “On the Design and Analysis of Irregular Algorithms on the Cell Processor: A case study of list ranking,” in *Proc. 21st IEEE Int'l. Parallel and Distributed Processing Symposium (IPDPS)*, Mar. 2007.
 68. D. A. Bader and K. Madduri, “A Graph-Theoretic Analysis of the Human Protein-Interaction Network Using Multi-core Parallel Algorithms,” in *Proc. 6th IEEE Int'l. Workshop on High-Performance Computational Biology (HiCOMB)*, Mar. 2007.

69. K. Madduri, D. A. Bader, J. W. Berry, and J. R. Crobak, “An experimental study of a parallel shortest path algorithm for solving large-scale graph instances,” in *Proc. 9th Workshop on Algorithm Engineering and Experiments (ALENEX)*, SIAM, Jan. 2007.
70. D. A. Bader and K. Madduri, “Parallel Algorithms for Evaluating Centrality Indices in Real-world Networks,” in *Proc. 35th Int’l. Conf. on Parallel Processing (ICPP)*, pp. 539–550, Aug. 2006.
71. D. A. Bader and K. Madduri, “Designing Multithreaded Algorithms for Breadth-First Search and st-connectivity on the Cray MTA-2,” in *Proc. 35th Int’l. Conf. on Parallel Processing (ICPP)*, pp. 523–530, Aug. 2006.
72. D. A. Bader and K. Madduri, “Design and Implementation of the HPCS Graph Analysis Benchmark on Symmetric Multiprocessors,” in D. A. Bader, M. Parashar, S. Varadarajan, and V. K. Prasanna (editors), *Proc. 12th Int’l. Conf. on High Performance Computing (HiPC)*, vol. 3769 of *LNCS*, pp. 465–476, Springer, Dec. 2005.
73. D. A. Bader and K. Madduri, “A Parallel State Assignment Algorithm for Finite State Machines,” in L. Bougé and V. K. Prasanna (editors), *Proc. 11th Int’l. Conf. on High Performance Computing (HiPC)*, vol. 3296 of *LNCS*, pp. 297–308, Springer, Dec. 2004.
74. K. Madduri, K. H. Aparna, and V. S. Chakravarthy, “PATRAM: A Handwritten Word Processor for Indian Languages,” in *Proc. 9th Int’l. Workshop on Frontiers in Handwriting Recognition (IWFHR)*, pp. 557–562, Aug. 2004.

Book Chapters

75. A. Buluç and K. Madduri, “Graph partitioning for scalable distributed graph computations,” in D. Bader, H. Meyerhenke, P. Sanders, and D. Wagner (editors), *Graph Partitioning and Graph Clustering*, chap. 6, pp. 81–100, AMS, 2013.
76. D. A. Bader and K. Madduri, “Computational challenges in emerging combinatorial scientific computing applications,” in U. Naumann and O. Schenk (editors), *Combinatorial Scientific Computing*, chap. 17, pp. 471–494, Chapman and Hall/CRC, Boca Raton, FL, 2012.
77. K. Madduri, “SNAP (Small-World Network Analysis and Partitioning) Framework,” in D. A. Padua (editor), *Encyclopedia of Parallel Computing*, pp. 1832–1837, Springer, 2011.
78. D. A. Bader, C. E. Heitsch, and K. Madduri, “Large-scale network analysis,” in J. Kepner and J. Gilbert (editors), *Graph Algorithms in the Language of Linear Algebra*, chap. 12, pp. 253–285, SIAM, Philadelphia, PA, 2011.
79. D. A. Bader, V. Agarwal, K. Madduri, and F. Petrini, “Combinatorial algorithm design on the Cell/B.E. processor,” in J. Kurzak, D. A. Bader, and J. Dongarra (editors), *Scientific Computing with Multicore and Accelerators*, chap. 10, pp. 195–216, CRC Press, Boca Raton, FL, 2010.

80. K. Madduri, D. A. Bader, J. W. Berry, and J. R. Crobak, "Parallel shortest path algorithms for solving large-scale instances," in C. Demetrescu, A. V. Goldberg, and D. Johnson (editors), *The Shortest Path Problem: Ninth DIMACS Implementation Challenge*, vol. 74, pp. 249–290, AMS, Providence, RI, 2009.
81. K. Madduri, D. A. Bader, J. W. Berry, J. R. Crobak, and B. A. Hendrickson, "Multi-threaded algorithms for processing massive graphs," in D. Bader (editor), *Petascale Computing: Algorithms and Applications*, chap. 12, pp. 237–262, Chapman and Hall/CRC, Boca Raton, FL, 2007.
82. D. A. Bader, K. Madduri, G. Cong, and J. Feo, "Design of multithreaded algorithms for combinatorial problems," in S. Rajasekaran and J. Reif (editors), *Handbook of Parallel Computing: Models, Algorithms, and Applications*, chap. 31, pp. 1–29, Chapman and Hall/CRC, Boca Raton, FL, 2007.

Conference and Workshop Publications without Proceedings

83. K. Madduri, "Distributed-memory graph algorithms for bioinformatics," SIAM Conf. on Computational Science and Engineering (CSE), Feb 2019.
84. S. Kirmani and K. Madduri, "A fast and simple multilevel spectral graph drawing algorithm," 8th SIAM Workshop on Combinatorial Scientific Computing (CSC), Jun. 2018.
85. G. M. Slota, S. Rajamanickam, K. Madduri, and K. D. Devine, "Partitioning irregular graphs at the trillion-edge scale," SIAM Conf. on Computational Science and Engineering (CSE), Feb 2017.
86. V. Rengasamy and K. Madduri, "High-performance graph traversal for De Bruijn graph-based metagenome assembly," SIAM Conf. on Computational Science and Engineering (CSE), Feb 2017.
87. G. Slota, S. Rajamanickam, and K. Madduri, "HPCGraph: Benchmarking massive graph analytics on supercomputers," 7th SIAM Workshop on Combinatorial Scientific Computing (CSC), Oct. 2016.
88. H. Zhan and K. Madduri, "A combinatorially-interpretable matrix factorization for network community structure evaluation," SIAM Annual Meeting, Jul. 2016.
89. G. M. Slota, S. Rajamanickam, and K. Madduri, "PuLP: Complex objective partitioning of small-world networks using label propagation," SIAM Conf. on Computational Science and Engineering, Mar. 2015.
90. T. Panitanarak and K. Madduri, "Performance analysis of single-source shortest path algorithms on distributed-memory systems," 6th SIAM Workshop on Combinatorial Scientific Computing (CSC), Jul. 2014.

91. G. Slota and K. Madduri, "Characterizing biological networks using subgraph counting and enumeration," SIAM Conf. on Parallel Processing for Scientific Computing, Feb. 2014.
92. G. Slota, S. Rajamanickam, and K. Madduri, "Parallel strongly connected components in shared memory architectures," SIAM Conf. on Parallel Processing for Scientific Computing, Feb. 2014.
93. K. Madduri, "Parallel analysis of graph-structured data in genomics and proteomics," First Int'l. Workshop on Big Data in Life Sciences (BigLS), Jun. 2013.
94. K. Madduri, "High-performance metagenomic data clustering and assembly," SIAM Annual Meeting, Jul. 2012.
95. K. Madduri, "Scalable SPARQL querying with compressed bitmap indexes," SIAM Conf. on Parallel Processing for Scientific Computing, Mar. 2012.
96. K. Madduri, "Optimizing short-read genome assembly algorithms for emerging multi-core platforms," SIAM Conf. on Computational Science and Engineering, February-March 2011.
97. K. Madduri, "Hybrid parallel programming for massive graph analysis," SIAM Annual Meeting, Jul. 2010.
98. K. Madduri, "Scaling up graph algorithms on emerging multicore systems," SIAM Annual Meeting, Jul. 2009.
99. K. Madduri, "High performance combinatorial techniques for processing dynamic interaction networks," SIAM Conf. on Parallel Processing for Scientific Computing, Mar. 2008.
100. D. A. Bader and K. Madduri, "High-performance combinatorial techniques for analyzing massive dynamic interaction networks," DIMACS/DyDAn Workshop on Computational Methods for Dynamic Interaction Networks, Sep. 2007.
101. K. Madduri, D. A. Bader, J. W. Berry, and J. R. Crobak, "Parallel shortest path algorithms for solving large-scale instances," 9th DIMACS Implementation Challenge workshop (The Shortest Path Problem), Nov. 2006.
102. D. A. Bader and K. Madduri, "Efficient shared-memory algorithms and implementations for solving large-scale graph problems," SIAM Annual Meeting, Jul. 2006.

Thesis

- K. Madduri, "A High-Performance Framework for Analyzing Massive Complex Networks," PhD dissertation, Georgia Institute of Technology, Atlanta, GA, July 2008.

Poster Presentations

- V. Rengasamy, P. Medvedev, and K. Madduri, “Harnessing Multinode Parallelism for Terascale Genomic Data Analysis,” *Biological Data Science*, Cold Spring Harbor Laboratory, Laurel Hollow, NY, Nov. 2018.
- K. Madduri, V. Rengasamy, and P. Medvedev, “SPRITE: A fast parallel SNP detection pipeline,” *American Society of Human Genetics (ASHG) Annual Meeting*, Baltimore, MD, Oct. 2015.
- G. M. Slota, K. Madduri, S. Rajamanickam, “Parallel Complex Network Partitioning,” *Supercomputing 2014*, New Orleans, LA, Nov. 2014.
- G. Slota, K. Madduri, “FASCIA: Fast Approximate Subgraph Counting and Enumeration,” *SIAM Workshop on Network Science*, San Diego, CA, July 2013.
- B. Wang, S. Ethier, W. Tang, K. Ibrahim, K. Madduri, S. Williams, “Advances in Gyrokinetic Particle-in-Cell Simulation for Fusion Plasmas to Extreme Scale,” *Supercomputing 2012*, Salt Lake City, UT, Nov. 2012.
- K. Madduri, “High-Performance Computing for Massive Graph Analysis,” *ATIP First Workshop on High-Performance Computing in India*, Portland, OR, Nov. 2009.
- K. Madduri, “SNAP: A Parallel Graph Framework for Large-scale Network Analysis,” *SIAM Conf. on Parallel Processing for Scientific Computing*, Atlanta, GA, March 2008.
- V. Agarwal and K. Madduri, “Efficient Implementation of Irregular Algorithms on the Cell Multi-core Architecture,” *Supercomputing 2006 Workshop – General-Purpose GPU Computing: Practice And Experience*, Tampa, FL, Nov. 2006.
- D. A. Bader and K. Madduri, “Efficient Graph Algorithms and Implementations on the Cray MTA-2,” *Fall Creek Falls Conference: Computational Science at Scale*, Pikeville, TN, Oct. 2005.

Open-source research code

1. SpectralGraphDrawing. S Kirmani and K. Madduri, 2019. <https://github.com/kmadduri/SpectralGraphDrawing>. This package implements a collection of graph drawing algorithms based on computing eigenvectors of the normalized adjacency matrix. Related publication: GABB18.
2. MetaPart. V. Rengasamy (primary developer, supervised graduate student), M.T. Kandemir, P. Medvedev, K. Madduri, 2018. <https://github.com/vasupsu/MetaPart>. MetaPart is a collection of preprocessing routines that can be performed on short-read DNA sequence data prior to *de novo* assembly. Related publication: HiPC18.
3. PKT. H. Kabir (primary developer, supervised graduate student), K. Madduri, 2017. <https://github.com/humayunk1/PKT>. PKT implements two shared-memory paral-

- lel algorithms for k -truss decomposition of large sparse graphs. Related publication: HiPC17.
4. MetaPrep. V. Rengasamy (primary developer, supervised graduate student), P. Medvedev, K. Madduri, 2017. <https://github.com/vasupsu/MetaPrep>. MetaPrep is a collection of preprocessing routines that can be performed on short-read DNA sequence data prior to *de novo* assembly. Related publication: HiCOMB17.
 5. PKC. H. Kabir (primary developer, supervised graduate student), K. Madduri, 2017. <https://github.com/humayunk1/PKC>. The PKC package implements several algorithms for k -core decomposition of large sparse graphs. Related publication: ParSocial17.
 6. HPCGraph. G. Slota (primary developer, supervised graduate student), S. Rajamanickam, K. Madduri, 2016. <https://github.com/HPCGraphAnalysis/HPCGraph>. HPCGraph is a collection of optimized graph analysis routines for distributed-memory systems. Related publications: IPDPS16, CSC16.
 7. ES-PICBench. R. Narayanan (primary developer, supervised graduate student), K. Madduri, 2016. <https://psu.app.box.com/s/9821pzzg31f1yptn4g7b51wpms8itngi>. ES-PICBench (PIC benchmark for Electrospray Simulations) is a parallel implementation of the particle-in-cell method for electrospray simulations. Related publication: PDSEC17.
 8. XtraPuLP and PuLP. G. Slota (primary developer, supervised graduate student), S. Rajamanickam, K. Madduri, 2016. <https://github.com/HPCGraphAnalysis/PuLP>. XtraPuLP and PuLP are parallel graph partitioning tools. Related publications: IPDPS17, SISC16, BigData14.
 9. KokkosConnectivity. G. Slota (primary developer, supervised graduate student), S. Rajamanickam, K. Madduri, 2016. <https://github.com/HPCGraphAnalysis/Connectivity>. KokkosConnectivity is an implementation of the Multistep algorithm and related parallel graph connectivity algorithms in the Kokkos shared-memory programming model. Related publications: IPDPS15a, IPDPS16.
 10. SPRITE. V. Rengasamy (primary developer, supervised graduate student) and K. Madduri, 2015. <http://sprite-psu.sourceforge.net>. SPRITE is an open-source bioinformatics software tool for accelerating the genetic variant detection pipeline. SPRITE consists of three new parallel tools: PRUNE, SAMPA, and PARSNIP. Related publications: ASHG15, ISC-HPC16.
 11. BiCC-BFS and BiCC-Coloring. G. Slota (primary developer, supervised graduate student) and K. Madduri, 2014. <http://www.graphanalysis.info>. BiCC-BFS and BiCC-Coloring are two new parallel algorithms and their corresponding software implementations for biconnected components decomposition of large-scale social networks. Related publication: HiPC 2014.

12. FastPath. G. Slota (primary developer, supervised graduate student), K. Madduri, 2014. <https://sourceforge.net/projects/fastpath-psu>. FastPath is an open-source graph analysis package for enumerating weighted paths using the color-coding method. Related publication: ParCo15.
13. DistSSSP. T. Panitanarak (primary developer, supervised graduate student) and K. Madduri, 2014. <http://www.graphanalysis.info>. DistSSSP has implementations of several parallel single-source shortest path algorithms designed for distributed-memory systems. Related publication: CSC14.
14. RDF3x-MPI. S. Chirravuri (primary developer, supervised graduate student) and K. Madduri, 2014. <https://bitbucket.org/saikrishnan/rdf3x-mpi/>. RDF3x-MPI is a parallel implementation of the RDF3x RDF data management and SPARQL querying platform. Related publication: Sai Chirravuri's MS thesis.
15. MultiStep. G. Slota (primary developer, supervised graduate student), S. Rajamanickam, K. Madduri, 2014. <https://github.com/HPCGraphAnalysis/Connectivity>. Multistep is an algorithmic strategy for various graph connectivity problems. Related publication: IPDPS14.
16. FASCIA. G. Slota (primary developer, supervised graduate student) and K. Madduri, 2014. <http://fascia-psu.sourceforge.net>. FASCIA is a software tool for determining approximate counts of tree-structured subgraphs in large networks. It is a parallel implementation based on the color-coding method. Related publications: ICPP13, IPDPS14, ParCo15.
17. Graph500-BFS-PSU, 2010-14. <http://graphanalysis.info>. Graph500 is a benchmark program and specification for ranking supercomputers. Madduri has been developing various parallel implementations of this specification since May 2010. The current version of Madduri's code is called Graph500-BFS-PSU (July 2014) and is available at graphanalysis.info as part of the DistSSSP package. Madduri has shared his past implementations with several collaborators in the academia and national laboratories, as well as industrial research groups.
18. GTC-P, 2009 – Present. Gyrokinetic Toroidal Code-Princeton (GTC-P) is a high-performance software program for ITER-sized global fusion simulations. Several versions of GTC (Gyrokinetic Toroidal Code) have been in development since 1998. GTC-P is maintained by Princeton Plasma Physics Laboratory and is available under the Theory Code licensing agreement. A closely-related “benchmark version” of this code is available at <http://www.nersc.gov/research-and-development/apex/apex-benchmarks/gtc-p/> as part of the NERSC benchmark suite. Related publications: SC16, SC13, IJH-PCA13, TPDS12, SC11, SC09.
19. SNAP, 2008-13. <http://snap-graph.sourceforge.net>. SNAP (Small-world Network Analysis and Partitioning) is an open-source graph analysis framework for multicore

platforms. This project is no longer actively maintained. Madduri was the lead developer of this project. Various versions of SNAP have been cumulatively downloaded 2100+ times on SourceForge. Related publication: IPDPS08.

20. HumanPINAnalysis, K. Madduri and D. A. Bader, 2007. <http://www.cse.psu.edu/~kxm85/software/HumanPPI>. HumanPINAnalysis is a collection of analysis routines, data sets, and supplemental results collected for the topological analysis of a large-scale human protein interaction network. Related publications: HiCOMB07, ParCo08.
21. ParallelSSSP, K. Madduri, 2006. ParallelSSSP is an implementation of the Delta stepping single-source shortest paths algorithm for the Cray MTA-2 platform. The code is no longer actively maintained, but is still available at <http://www.cse.psu.edu/~kxm85/software/ParallelSSSP/>. Related publication: DIMACS07.
22. GTgraph, K. Madduri and D. A. Bader, 2006. GTgraph is a toolkit of synthetic random graph generators. The code is no longer maintained, and the version from Feb 2006 is available at <http://www.cse.psu.edu/~kxm85/software/GTgraph>. Related publication: documentation at website.
23. SSCA#2, K. Madduri and D. A. Bader, 2005-12. SSCA#2 is a graph analysis benchmark designed as part of the DARPA High Productivity Computing Systems (HPCS) project to characterize performance of novel architectures and programming languages on graph-theoretic kernels. Madduri and Bader developed the first reference C/OpenMP parallel implementation of this benchmark in 2005, and Madduri has continued developing and maintaining new versions. Old versions of the benchmark (v1.0, v2.1, v2.2) are available at <http://www.graphanalysis.org>. The most recent version from 2012 is available upon request. Related publications: HiPC05, ICPP06, MTAAP09, SC12.

Selected Invited Talks

1. Computing and Information Sciences Invited Lecture, Florida International University, Oct. 26, 2018.
2. Whitacre College of Engineering seminar, Texas Tech, Oct. 23, 2017.
3. Computer Science seminar, Virginia Tech, Oct. 20, 2017.
4. Computer Science and Engineering seminar, The University at Buffalo, Nov. 5, 2015.
5. Dagstuhl seminar on High-performance Graph Algorithms and Applications in Computational Science, Nov. 13, 2014.
6. CScADS Workshop on Libraries and Autotuning for Extreme-scale Applications, Aug. 13, 2012.
7. Workshop on Algorithms for Modern Massive Data Sets (MMDS 2012), July 13, 2012.
8. The Pennsylvania State University, Nov. 8, 2010.

9. SIAM Conf. on Parallel Processing for Scientific Computing (SIAM PP10), Feb. 25, 2010.
10. ParLab Winter Retreat 2010, UC Berkeley, Jan. 14, 2010.
11. National Security Agency, Oct. 4, 2009.
12. Bay Area Scientific Computing Day 2009, Lawrence Berkeley National Laboratory, May 9, 2009.
13. Department of Computer Science and Electrical Engineering, West Virginia University, April 24, 2009.
14. Computer Science Research Institute, Sandia National Laboratories, April 9, 2008.
15. Computer Science Department, University of California Davis, April 1, 2008.
16. Computer Science and Mathematics Division, Oak Ridge National Laboratory, March 17, 2008.
17. Computational Research Division, Lawrence Berkeley National Laboratory, March 10, 2008.

Teaching and Advising

Classes at Penn State

- CMPSC 450, Concurrent Scientific Programming, Spring 2017 (64 students), Spring 2016 (41 students), Spring 2015 (29 students), Spring 2014 (30 students), and Spring 2013 (25 students).
- CMPSC/MATH 451, Numerical Computations, Spring 2018 (160 students, 2 sections), Spring 2016 (83 students, 2 sections), Fall 2013 (60 students, 2 sections), and Fall 2011 (62 students, 2 sections).
- CMPSC/MATH 455, Numerical Analysis I, Spring 2019 (63 students), Spring 2013 (32 students), and Fall 2012 (30 students).
- CMPSC 465, Data Structures and Algorithms, Fall 2018 (260 students), Fall 2016 (209 students), and Fall 2014 (154 students).
- CSE 557, Concurrent Matrix Computations, Spring 2014 (19 students) and Spring 2013 (12 students).
- CSE 597C, Graph Mining, Spring 2015 (18 students).

Current Advisees

- Ashirbad Mishra, CSE PhD student, Spring 2019 – Present.
- Hongyuan Zhan. CSE PhD candidate, Fall 2014 – Spring 2019 (expected).

Alumni

- David Hadka. PhD in Computer Science and Engineering, May 2013. Thesis: The Foundations, Parallelization, and Application of the BORG MOEA. Co-advised with Prof. Patrick Reed, Department of Civil and Environmental Engineering. First employment: Penn State Applied Research Laboratory.
- Sai Chirravuri. MS in Computer Science and Engineering, July 2014. Thesis: RDF3X-MPI: A Partitioned RDF engine for Data-Parallel SPARQL Querying. First employment: Cloudera.
- Ramachandran K. Narayanan. MS in Computer Science and Engineering, May 2016. Thesis: Parallel Particle-in-cell Performance Optimization: A Case Study of Electrospray Simulation. First employment: RNET Technologies Inc.
- Sindhuja Parimalarangan. MS in Computer Science and Engineering, May 2016. Thesis: Fast Parallel Triad Census and Triangle Listing on Shared-Memory Platforms. First employment: MathWorks.
- George Slota. PhD in Computer Science and Engineering, May 2016. Thesis: Irregular Graph Algorithms on Modern Multicore, Manycore, and Distributed Processing Systems. First employment: Rensselaer Polytechnic Institute.
- Thap Panitanarak. PhD in Computer Science and Engineering, August 2017. Thesis: Scalable graph and mesh algorithms on distributed-memory systems. Co-advised with Prof. Suzanne Shontz, Univ. of Kansas. First employment: Chulalongkorn University.
- Humayun Kabir. PhD in Computer Science and Engineering, May 2018. Thesis: Hierarchical Sparse Graph Computations on Multicore Platforms. Advised 2011-16 by Prof. Padma Raghavan, Vanderbilt. First employment: AQR Capital Management.
- Daniel Witman. B.S. in Computer Science, Spring 2018. Schreyer Honors College Thesis: Large Scale Interactive Graph Visualization and Automated Social Network Centrality Analysis.
- Vasudevan Rengasamy. PhD in Computer Science and Engineering, December 2018. Thesis: Engineering high performance workflows for end-to-end acceleration of genomic applications. First employment: Intel.

Other

- Emmanuel Oppong, Summer 2014 SROP intern.
- Ucheoma Ukah, Summer 2014 SROP intern.
- Aaron Altman, CSE PhD program, 2012-13.
- Fei Wu, CSE PhD program, 2012-13.
- Lab rotation adviser/co-adviser: Vijay Kumar (Spring 2017), Ayaan Hossain (Fall 2016), and Nick Stoler (Fall 2012).

Student Thesis Committee Membership

In progress: 2 PhD students. Completed: 18 PhD students, 7 MS students, 7 MEng students, 1 CMPSC Honors student.

Service

Conference Organization

1. Program Chair, 18th Int'l. Workshop on High Performance Computational Biology (HiCOMB), 2019.
2. Doctoral Research Showcase Vice Chair, Supercomputing, 2018.
3. Algorithms track Chair, 47th Int'l. Conf. on Parallel Processing (ICPP), 2018.
4. Co-Organizer, Int'l. Workshop on High Performance Big Graph Data Management, Analysis, and Mining (BigGraphs), 2014-19.
5. Publicity Co-Chair, 46th Int'l. Conf. on Parallel Processing (ICPP), 2017.
6. Applications track Co-Chair, The 10th Int'l. Conf. on Contemporary Computing, 2017.
7. Minisymposium Co-Organizer, SIAM Conf. on Computational Science and Engineering, 2017 and 2015.
8. Co-Chair, Student Research Symposium at the 21st Annual Int'l. Conf. on High Performance Computing (HiPC), 2014.
9. Algorithms track Co-Chair, The 7th Int'l. Conf. on Contemporary Computing, 2014.
10. Organizing Committee, SIAM Conf. on Parallel Processing for Scientific Computing, 2012.
11. Cyber Co-Chair, The 15th Int'l. Conf. on High Performance Computing (HiPC), 2008.

Conference Program Committee Service

- SIAM Conf. on Parallel Processing for Scientific Computing, 2020.
- Int'l. Conf. for High Performance Computing, Networking, Storage and Analysis (SC), 2019 (Applications track), 2015 (Algorithms), 2014 (Algorithms), 2012 (Applications).
- Int'l. Conf. on Parallel Processing (ICPP), 2019 (Applications track), 2017 (Algorithms), 2013 (Applications), 2010 (Algorithms and Applications).
- IEEE Int'l. Conf. on High Performance Computing, Data, and Analytics (HiPC), 2019 (Applications track), 2018 (Applications), 2017, 2016 (Applications), 2013, 2008.
- IEEE Int'l. Parallel and Distributed Processing Symp. (IPDPS), 2018 (Algorithms track), 2017 (Applications), 2016 (Algorithms), 2015 (Applications), 2014 (Algorithms), 2013 (Algorithms), 2010.

- IEEE Int'l. Conf. on Cluster Computing (Cluster), 2017 (Algorithms, Applications, and Libraries track).
- ACM Conf. on Bioinformatics, Computational Biology, and Health Informatics (BCB), 2017.
- IEEE Int'l. Conf. on Contemporary Computing (IC3), 2019, 2016, 2015.
- IEEE Int'l. Symp. on Signal Processing and Information Technology (ISSPIT), 2014.
- Int'l. Conf. on Algorithms and Architectures for Parallel Processing (ICA3PP), 2014, 2012.
- IEEE Int'l. Symp. on Parallel and Distributed Processing with Applications (ISPA), 2014, 2012.
- ACM Int'l. Conf. on Supercomputing (ICS), External Committee, 2014.
- IEEE/ACM Int'l. Symp. on Cluster, Cloud and Grid Computing (CCGrid), 2014 (Architecture and Accelerators track), 2010.
- Biotechnology and Bioinformatics Symp. (BIOT), 2013, 2012.
- IEEE Int'l. Conf. on High Performance Computing and Communications (HPCC), 2011-13.
- ACM SIGPLAN Symp. on Principles and Practice of Parallel Programming (PPoPP), External Committee, 2013, 2012.
- Symp. on Application Accelerators in High Performance Computing (SAAHPC), 2010-12.
- Int'l. Symp. on Computer Architecture and High Performance Computing (SBAC-PAD), 2009.

Workshop Program Committee Service

- Workshop on Graphs, Architectures, Programming, and Learning (GrAPL), 2019.
- Graph Algorithms Building Blocks (GABB) Workshop, 2015-18.
- IEEE Int'l. Workshop on High Performance Computational Biology (HiCOMB), 2018, 2016, 2012, 2011.
- Workshop on Irregular Applications: Architectures & Algorithms (IA³), 2013-18.
- Workshop on Converged Computing Infrastructure (CCIW), 2019.
- Workshop on Domain-Specific Languages and High-Level Frameworks for High Performance Computing (WOLFHPC), 2012-17.
- High Performance Graph Data Management and Processing (HPGDMP) Workshop, 2016.
- SIAM Workshop on Combinatorial Scientific Computing, 2016.

- The Int'l. Workshop on Accelerators and Hybrid Exascale Systems (AsHES), 2016.
- Workshop on High Performance Computing for Big Data (HPC4BD), 2015.
- Programming Models, Languages and Compilers for Manycore and Heterogeneous Architectures (PLC) Workshop, 2015.
- ACM Int'l. Workshop on Big Data in Life Sciences (BigLS), 2014.
- Workshop on High Performance Computing for Big Data (HPC4BD), 2014.
- Workshop on Multithreaded Architectures and Applications (MTAAP), 2014.
- Workshop on Mining Networks and Graphs: A Big Data Analytics Challenge, 2014.
- Workshop on Parallel Algorithms and Software for Analysis of Massive Graphs (Par-Graph), 2011-13.
- Network-aware Data Management Workshop (NDM), 2011.
- Workshop on Parallel Computational Biology (PBC), 2011.
- Workshop on Hybrid Multi-core Computing (WHMC), 2010.

Other Program Committee Service

- Doctoral Research Showcase, Supercomputing, 2017, 2016, 2014.
- Student Research Symp., IEEE Int'l. Conf. on High Performance Computing (HiPC), 2013, 2010.
- PhD Forum, IEEE Int'l. Parallel and Distributed Processing Symp. (IPDPS), 2012, 2010, 2009.

Journal Editorial Service

- Associate Editor, IEEE Transactions on Parallel and Distributed Systems, Oct. 2018 – Present.

Journal Manuscript Reviewing

ACM Journal of Computational Science, ACM Journal of Experimental Algorithmics, ACM Transactions on Parallel Computing, Elsevier Artificial Intelligence in Medicine, Elsevier Discrete Applied Mathematics, Elsevier Information Sciences, Elsevier Knowledge-based Systems, Elsevier Journal of Parallel and Distributed Computing, Elsevier Parallel Computing, Elsevier Theoretical Computer Science, IEEE Design and Test of Computers, IEEE Transactions on Knowledge and Data Engineering, IEEE Transactions on Parallel and Distributed Systems, IEEE/ACM Transactions on Computational Biology and Bioinformatics, IEEE Transactions on Network Science and Engineering, Oxford Journals Bioinformatics, PLOS ONE, SAGE International Journal of High Performance Computing Applications, SIAM Journal on Scientific Computing, Wiley Concurrency and Computation: Practice and Experience, Wiley Interscience Networks. 2006 – Present.

Research Proposal Review

- National Science Foundation, review panels, 2012 – Present.
- Department of Energy, 2018 (ad-hoc), 2016 (ad-hoc), 2015 (ad-hoc), 2013 (panel), 2011 (panel).

Other

- GLCPC Blue Waters Proposals Committee, 2013, 2014, 2015.
- SIAG/Supercomputing Prize Selection Committee, 2014.
- Judge, OpenMP Programming Contest, Supercomputing 2005.
- Student volunteer, Supercomputing (SC) conference, 2005, 2006, 2007.

Penn State, University, College, and Department Service

- Penn State Research Computing and Cyberinfrastructure, High-Performance Computing Working Group, 2018-19.
- College of Engineering Research Computing committee, 2016-17, 2018-19.
- CSE graduate committee, 2011-19.
- CSE curriculum committee, 2016-19.
- Numerical Analysis area candidacy exam chair, 2012-13, 2014-15, 2016-17, 2018-19.
- CSE colloquium chair, 2013-16.
- CSE curriculum/ABET committee, 2014-16.
- CMPEN major faculty adviser, College of Engineering Advising Center, 2012-13.