

CONTACT INFORMATION	Lawrence Berkeley National Laboratory. Mailstop 50F1650 1 Cyclotron Road Berkeley, CA 94720 USA	Research team: https://passion.lbl.gov e-mail: abuluc@lbl.gov https://people.eecs.berkeley.edu/~aydin/
RESEARCH INTERESTS	Parallel Computing, Combinatorial Scientific Computing, Sparse Matrices, High-Performance Computing for Computational Biology, Graph Algorithms, and Machine Learning.	
EDUCATION AND TRAINING	UC Center for Executive Education , Berkeley, CA, USA LBNL Leadership Development Program for Emerging Leaders, 2015-16 University of California , Santa Barbara, CA, USA Ph.D. in Computer Science, 2010 <i>Thesis: Linear Algebraic Primitives for Parallel Computing on Large Graphs</i> <i>Advisor: John R. Gilbert</i> M.S. in Computer Science, 2009 Massachusetts Institute of Technology , Cambridge, MA, USA Visiting Student at the Department of Mathematics, 2008 Sabancı University , Istanbul, Turkey B.S. in Computer Science, 2005 (with Minors in Mathematics)	
RESEARCH AND PROFESSIONAL EXPERIENCE	CRD, Lawrence Berkeley National Laboratory , Berkeley, CA - Computational Senior Scientist, <i>August 2022 - Present</i> - Computational Staff Scientist (Career), <i>March 2016 - August 2022</i> - Computational Research Scientist (Career), <i>April 2014 - March 2016</i> - Computational Research Scientist (Career-track), <i>April 2012 - March 2014</i> - Luis W. Alvarez Postdoctoral Fellow , <i>April 2010 - April 2012</i> EECS Department, University of California , Berkeley, CA - Adjunct Assistant Professor, <i>January 2017 - Present</i> - Affiliated Researcher, <i>April 2010 - Present</i> Bebop (BERkeley Benchmarking and OPTimization) group Simons Institute for the Theory of Computing , Berkeley, CA Visiting Scientist (Long Term Participant), <i>August 2013 - December 2013</i> Program on Theoretical Foundations of Big Data Analysis CSRI, Sandia National Laboratory , Albuquerque, NM Graduate Research Intern, <i>June 2008 - August 2008</i> Citrix Online , Santa Barbara, CA Software Engineer Intern, <i>June 2006 - September 2006</i> Garanti Technology , Istanbul, Turkey Software Engineer Intern, <i>July 2004 - September 2004</i> Computer Science Department, University of California , Santa Barbara, CA Research Assistant, Combinatorial Scientific Computing Laboratory , 2007-2010 Teaching Assistant, 2005-2007	

SCHOLARSHIPS &
AWARDS

- Major Awards and Nominations
 - *ACM Gordon Bell Finalist*, 2023
 - *IEEE TCSC Award for Excellence for Early Career Research* by the IEEE Technical Committee on Scalable Computing, 2015
 - *DOE Early Career Award* by Office of Science, Department of Energy, 2013
- Paper Awards
 - Best Paper Nominee, SC'21: International Conference for High Performance Computing, Networking, Storage and Analysis, 2021
 - Distinguished Paper, 24th International European Conference on Parallel and Distributed Computing (Euro-Par 2018)
 - Best Paper Nominee, SC'18: International Conference for High Performance Computing, Networking, Storage and Analysis, 2018
- Magazine Awards
 - *Best Use of HPC Application in Life Sciences* (as part of HipMer team) for boosting the assembly of the human genome on the Cray XC30 'Edison' supercomputer by the HPCWire Magazine (Readers' Choice), 2015
 - *Best Use of HPC Application in Life Sciences* (as part of MetaHipMer team) for unprecedented assembly of environmental microbiomes – evolutionary processes of microbial communities on decadal timescales (25 terabyte dataset) by the HPCWire Magazine (Editors' Choice), 2021
- Scholarships
 - *Luis W. Alvarez Postdoctoral Fellowship in Computational Science* by Lawrence Berkeley National Laboratory, 2010
 - Merit scholarship support (tuition, fees, accommodation) for undergraduate studies, Sabancı University, 2001-2005.
- Minor Awards
 - *Lawrence Berkeley National Laboratory Spot Award*
For keeping NERSC at the top 2 of Graph500 list for three consecutive terms, 2012.

JOURNAL
PUBLICATIONS
(REFEREED)

1. Georgios A. Pavlopoulos, Fotis A. Baltoumas, Sirui Liu, Oguz Selvitopi, Antonio Pedro Camargo, Stephen Nayfach, Ariful Azad, Simon Roux, Lee Call, Natalia N. Ivanova, I. Min Chen, David Paez-Espino, Evangelos Karatzas, Novel Metagenome Protein Families Consortium, Ioannis Iliopoulos, Konstantinos Konstantinidis, James M. Tiedje, Jennifer Pett-Ridge, David Baker, Axel Visel, Christos A. Ouzounis, Sergey Ovchinnikov, **Aydın Buluç**, and Nikos C. Kyrpides. Unraveling the functional dark matter through global metagenomics. *Nature*, 622(7983):594–602, 2023.
2. Naw Safrin Sattar, **Aydın Buluç**, Khaled Z Ibrahim, and Shaikh Arifuzzaman. Exploring temporal community evolution: algorithmic approaches and parallel optimization for dynamic community detection. *Applied Network Science*, 8(1):64, 2023.
3. Can Kızılkale, Farid Rashidi Mehrabadi, Erfan Sadeqi Azer, Eva Pérez-Guijarro, Kerrie L. Marie, Maxwell P. Lee, Chi-Ping Day, Glenn Merlino, Funda Ergün, **Aydın Buluç**, S. Cenk Sahinalp, and Salem Malikić. Fast intratumor heterogeneity inference from single-cell sequencing data. *Nature Computational Science*, 2(9):577–583, 2022.
4. Fernando Meyer, Adrian Fritz, Zhi-Luo Deng, David Koslicki, Till Robin Lesker, Alexey Gurevich, Gary Robertson, Mohammed Alser, Dmitry Antipov, Francesco Beghini, Denis Bertrand, Jaqueline J. Brito, C. Titus Brown, Jan Buchmann, **Aydın Buluç**, Bo Chen, Rayan Chikhi, Philip T. L. C. Clausen, Alexandru Cristian, Piotr Wojciech Dabrowski, Aaron E. Darling, Rob Egan, Eleazar Eskin, Evangelos Geor-

- ganas, Eugene Goltsman, Melissa A. Gray, Lars Hestbjerg Hansen, Steven Hofmeyr, Pingqin Huang, Luiz Irber, Huijue Jia, Tue Sparholt Jørgensen, Silas D. Kieser, Terje Klemetsen, Axel Kola, Mikhail Kolmogorov, Anton Korobeynikov, Jason Kwan, Nathan LaPierre, Claire Lemaitre, Chenhao Li, Antoine Limasset, Fabio Malcher-Miranda, Serghei Mangul, Vanessa R. Marcelino, Camille Marchet, Pierre Marijon, Dmitry Meleshko, Daniel R. Mende, Alessio Milanese, Niranjana Nagarajan, Jakob Nissen, Sergey Nurk, Leonid Oliner, Lucas Paoli, Pierre Peterlongo, Vitor C. Piro, Jacob S. Porter, Simon Rasmussen, Evan R. Rees, Knut Reinert, Bernhard Renard, Espen Mikal Robertsen, Gail L. Rosen, Hans-Joachim Ruscheweyh, Varuni Sarwal, Nicola Segata, Enrico Seiler, Lizhen Shi, Fengzhu Sun, Shinichi Sunagawa, Søren Johannes Sørensen, Ashleigh Thomas, Chengxuan Tong, Mirko Trajkovski, Julien Tremblay, Gherman Urtskiy, Riccardo Vicedomini, Zhengyang Wang, Ziyi Wang, Zhong Wang, Andrew Warren, Nils Peder Willassen, Katherine Yelick, Ronghui You, Georg Zeller, Zhengqiao Zhao, Shanfeng Zhu, Jie Zhu, Ruben Garrido-Oter, Petra Gastmeier, Stéphane Hacquard, Susanne Häubler, Ariane Khaledi, Friederike Maechler, Fantin Mesny, Simona Radutoiu, Paul Schulze-Lefert, Nathiana Smit, Till Strowig, Andreas Bremges, Alexander Sczyrba, and Alice Carolyn McHardy. Critical assessment of metagenome interpretation: the second round of challenges. *Nature Methods*, 19(4):429–440, 2022.
5. Ariful Azad, Oguz Selvitopi, Md Taufique Hussain, John Gilbert, and **Aydın Buluç**. Combinatorial BLAS 2.0: Scaling combinatorial algorithms on distributed-memory systems. *IEEE Transactions on Parallel and Distributed Systems*, 2021.
 6. Carl Yang, **Aydın Buluç**, and John D Owens. GraphBLAST: A high-performance linear algebra-based graph framework on the GPU. *ACM Transactions on Mathematical Software*, 48(1):1–51, 2022.
 7. Seher Acer, Ariful Azad, Erik Boman, **Aydın Buluç**, Karen Devine, Nitin Gawande, Sayan Ghosh, Mahantesh Halappanavar, Arif Khan, Ananth Kalyanaraman, Marco Minutoli, Alex Pothén, Sivasankaran Rajamanickam, Oguz Selvitopi, Nathan Tallent, and Antonino Tumeo. Exagraph: Graph and combinatorial methods for enabling exascale applications. *International Journal of High Performance Computing Applications (IJHPCA)*, 35(6):553–571, 2021.
 8. Ariful Azad, **Aydın Buluç**, Xiaoye S Li, Xinliang Wang, and Johannes Langguth. A distributed-memory algorithm for computing a heavy-weight perfect matching on bipartite graphs. *SIAM Journal on Scientific Computing*, 42(4):C143–C168, 2020.
 9. Steven Hofmeyr, Rob Egan, Evangelos Georganas, Alex Copeland, Robert Riley, Alicia Clum, Emiley Eloé-Fadrosh, Simon Roux, Eugene Goltsman, **Aydın Buluç**, Daniel Rokhsar, Leonid Oliner, and Katherine Yelick. Terabase-scale metagenome coassembly with MetaHipMer. *Scientific reports*, 2020.
 10. Muaaz G Awan, Jack Deslippe, **Aydın Buluç**, Oguz Selvitopi, Steven Hofmeyr, Leonid Oliner, and Katherine Yelick. ADEPT: a domain independent sequence alignment strategy for gpu architectures. *BMC bioinformatics*, 21(1):1–29, 2020.
 11. Yongzhe Zhang, Ariful Azad, and **Aydın Buluç**. Parallel algorithms for finding connected components using linear algebra. *Journal of Parallel and Distributed Computing*, 2020.
 12. Katherine Yelick, **Aydın Buluç**, Muaaz Awan, Ariful Azad, Benjamin Brock, Rob Egan, Saliya Ekanayake, Marquita Ellis, Evangelos Georganas, Giulia Guidi, et al. The parallelism motifs of genomic data analysis. *Philosophical Transactions of the Royal Society A*, 378(2166):20190394, 2020.
 13. Yusuke Nagasaka, Satoshi Matsuoka, Ariful Azad, and **Aydın Buluç**. Performance optimization, modeling and analysis of sparse matrix-matrix products on multi-core and many-core processors. *Parallel Computing*, page 102545, 2019.

14. Ariful Azad, Georgios A. Pavlopoulos, Christos A. Ouzounis, Nikos C. Kyrpides, and **Aydm Buluç**. HipMCL: A high-performance parallel implementation of the Markov clustering algorithm for large-scale networks. *Nucleic Acids Research (NAR)*, 2018.
15. Ariful Azad, **Aydm Buluç**, and Alex Pothen. Computing maximum cardinality matchings in parallel on bipartite graphs via tree-grafting. *IEEE Transactions on Parallel and Distributed Systems (TPDS)*, 28(1):44–59, 2017.
16. Hasan Metin Aktulga, Md. Afbuzzaman, Samuel Williams, **Aydm Buluç**, Meiyue Shao, Chao Yang, Esmund G. Ng, Pieter Maris, and James P. Vary. A high performance block eigensolver for nuclear configuration interaction calculations. *IEEE Transactions on Parallel and Distributed Systems (TPDS)*, 28(6):1550–1563, 2017.
17. Ariful Azad, Grey Ballard, **Aydm Buluç**, James Demmel, Laura Grigori, Oded Schwartz, Sivan Toledo, and Samuel Williams. Exploiting multiple levels of parallelism in sparse matrix-matrix multiplication. *SIAM Journal on Scientific Computing (SISC)*, 38(6):C624–C651, 2016.
18. Ariful Azad and **Aydm Buluç**. A matrix-algebraic formulation of distributed-memory maximal cardinality matching algorithms in bipartite graphs. *Parallel Computing*, 2016.
19. Jarrod Chapman, Martin Mascher, **Aydm Buluç**, Kerrie Barry, Evangelos Georganas, Adam Session, Veronika Strnadova, Jerry Jenkins, Sunish Sehgal, Leonid Oliker, Jeremy Schmutz, Katherine Yelick, Uwe Scholz, Robbie Waugh, Jesse Poland, Gary Muehlbauer, Nils Stein, and Daniel Rokhsar. A whole-genome shotgun approach for assembling and anchoring the hexaploid bread wheat genome. *Genome Biology*, 16(26), 2015.
20. Adam Lugowski, Shoaib Kamil, **Aydm Buluç**, Samuel Williams, Erika Duriakova, Leonid Oliker, Armando Fox, and John Gilbert. Parallel processing of filtered queries in attributed semantic graphs. *Journal of Parallel and Distributed Computing (JPDC)*, 79-80:115–131, 2015.
21. **Aydm Buluç** and John R. Gilbert. Parallel sparse matrix-matrix multiplication and indexing: Implementation and experiments. *SIAM Journal of Scientific Computing*, 34(4):170 – 191, 2012.
22. **Aydm Buluç** and John R. Gilbert. The Combinatorial BLAS: Design, implementation, and applications. *International Journal of High Performance Computing Applications (IJHPCA)*, 25(4):496–509, 2011.
23. **Aydm Buluç**, John R. Gilbert, and Ceren Budak. Solving path problems on the GPU. *Parallel Computing*, 36(5-6):241 – 253, 2010.

CONFERENCE
PUBLICATIONS
(REFEREED)

1. Tianyu Liang, Riley Murray, **Aydm Buluç**, and James Demmel. Fast multiplication of random dense matrices with fixed sparse matrices. In *International Parallel and Distributed Processing Symposium (IPDPS)*. IEEE, 2024.
2. Brian Wheatman, Randal Burns, **Aydm Buluç**, and Helen Xu. CPMA: An efficient batch-parallel compressed set without pointers. In *ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming (PPOPP)*, 2024.
3. Vivek Bharadwaj, Osman Asif Malik, Riley Murray, Laura Grigori, **Aydm Buluç**, and James Demmel. Fast exact leverage score sampling from Khatri-Rao products with applications to tensor decomposition. In *Thirty-seventh Conference on Neural Information Processing Systems (NeurIPS)*, 2023.

4. Luk Burchard, Max Xiaohang Zhao, Johannes Langguth, **Aydin Buluç**, and Giulia Guidi. Space efficient sequence alignment for SRAM-based computing: X-Drop on the Graphcore IPU. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, SC'23, New York, NY, USA, 2023. Association for Computing Machinery.
5. Brian Wheatman, Randal Burns, **Aydin Buluç**, and Helen Xu. Optimizing search layouts in packed memory arrays. In *Proceedings of the Symposium on Algorithm Engineering and Experiments (ALENEX)*, pages 148–161. SIAM, 2023.
6. Oguz Selvitopi, Saliya Ekanayakey, Giulia Guidiz, Muaaz G Awanx, Georgios A Pavlopoulos, Ariful Azadk, Nikos Kyrpides, Leonid Oliker, Katherine Yelick, and **Aydin Buluç**. Extreme-scale many-against-many protein similarity search. In *International Conference for High Performance Computing, Networking, Storage and Analysis (SC)*, pages 1–12. IEEE Computer Society, 2022.
7. Yuxin Chen, Benjamin Brock, Serban Porumbescu, **Aydin Buluç**, Katherine Yelick, and John D Owens. Scalable irregular parallelism with GPUs: getting CPUs out of the way. In *International Conference for High Performance Computing, Networking, Storage and Analysis (SC)*, pages 708–723. IEEE Computer Society, 2022.
8. Srdjan Milaković, Oguz Selvitopi, Israt Nisa, Zoran Budimlić, and **Aydin Buluç**. Parallel algorithms for masked sparse matrix-matrix products. In *Proceedings of the 51st International Conference on Parallel Processing*, pages 1–11, 2022.
9. Giulia Guidi, Gabriel Raulet, Daniel Rokhsar, Leonid Oliker, Katherine Yelick, and **Aydin Buluç**. Distributed-memory parallel contig generation for de novo long-read genome assembly. In *Proceedings of the 51st International Conference on Parallel Processing*, pages 1–11, 2022.
10. Yuxin Chen, Benjamin Brock, Serban Porumbescu, **Aydin Buluç**, Katherine Yelick, and John Owens. Atos: A task-parallel GPU scheduler for graph analytics. In *Proceedings of the 51st International Conference on Parallel Processing*, pages 1–11, 2022.
11. Vivek Bharadwaj, **Aydin Buluç**, and James Demmel. Distributed memory sparse kernels for machine learning. In *International Parallel and Distributed Processing Symposium (IPDPS)*, 2022.
12. Ivan Kawaminami, Arminda Estrada, Youssef Elsakkary, Hayden Jananthan, **Aydin Buluç**, Tim Davis, Daniel Grant, Michael Jones, Chad Meiners, Andrew Morris, et al. Large scale enrichment and statistical cyber characterization of network traffic. In *High Performance Extreme Computing Conference (HPEC)*, pages 1–7. IEEE, 2022.
13. 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.
14. Giulia Guidi, Marquita Ellis, Daniel Rokhsar, Katherine Yelick, and **Aydin Buluç**. BELLA: Berkeley efficient long-read to long-read aligner and overlapper. In *SIAM Conference on Applied and Computational Discrete Algorithms (ACDA21)*, pages 123–134. SIAM, 2021.
15. Oguz Selvitopi, Benjamin Brock, Israt Nisa, Alok Tripathy, Katherine Yelick, and **Aydin Buluç**. Distributed-memory parallel algorithms for sparse times tall-skinny-dense matrix multiplication. In *Proceedings of the ACM International Conference on Supercomputing (ICS)*, pages 431–442, 2021.

16. Prashant Pandey, Brian Wheatman, Helen Xu, and **Aydın Buluç**. Terrace: A hierarchical graph container for skewed dynamic graphs. In *Proceedings of the 2021 International Conference on Management of Data*, pages 1372–1385, 2021.
17. Muaaz Awan, Steven Hofmeyr, Rob Egan, Nan Ding, **Aydın Buluç**, Jack Deslippe, Leonid Oliker, and Katherine Yelick. Accelerating large scale de novo metagenome assembly using GPUs. In *International Conference for High Performance Computing, Networking, Storage and Analysis, SC'21*, 2021.
18. Marquita Ellis, **Aydın Buluç**, and Katherine Yelick. Scaling generalized n-body problems, a case study from genomics. In *50th International Conference on Parallel Processing (ICPP)*, 2021.
19. Giulia Guidi, Oguz Selvitopi, Marquita Ellis, Leonid Oliker, Katherine Yelick, and **Aydın Buluç**. Parallel string graph construction and transitive reduction for de novo genome assembly. In *International Parallel and Distributed Processing Symposium (IPDPS)*, pages 517–526. IEEE, 2021.
20. Md Taufique Hussain, Oguz Selvitopi, **Aydın Buluç**, and Ariful Azad. Communication-avoiding and memory-constrained sparse matrix-matrix multiplication at extreme scale. In *International Parallel and Distributed Processing Symposium (IPDPS)*, pages 90–100. IEEE, 2021.
21. Israt Nisa, Prashant Pandey, Marquita Ellis, Leonid Oliker, **Aydın Buluç**, and Katherine Yelick. Distributed-memory k-mer counting on GPUs. In *International Parallel and Distributed Processing Symposium (IPDPS)*, pages 527–536. IEEE, 2021.
22. Jeremy Kepner, Michael Jones, Daniel Andersen, **Aydın Buluç**, Chansup Byun, K Claffy, Timothy Davis, William Arcand, Jonathan Bernays, David Bestor, et al. Spatial temporal analysis of 40,000,000,000,000 internet darkspace packets. In *2021 IEEE High Performance Extreme Computing Conference (HPEC)*, pages 1–8. IEEE, 2021.
23. Alok Tripathy, Katherine Yelick, and **Aydın Buluç**. Reducing communication in graph neural network training. In *International Conference for High Performance Computing, Networking, Storage and Analysis, SC'20*, 2020.
24. Oguz Selvitopi, Saliya Ekanayake, Giulia Guidi, Georgios Pavlopoulos, Ariful Azad, and **Aydın Buluç**. Distributed many-to-many protein sequence alignment using sparse matrices. In *International Conference for High Performance Computing, Networking, Storage and Analysis, SC'20*, 2020.
25. Yu-Hang Tang, Oguz Selvitopi, Doru Popovici, and **Aydın Buluç**. A high-throughput solver for marginalized graph kernels on GPU. In *International Parallel and Distributed Processing Symposium (IPDPS)*, 2020.
26. Oguz Selvitopi, Md Taufique Hussain, Ariful Azad, and **Aydın Buluç**. Optimizing high performance Markov clustering for pre-exascale architectures. In *International Parallel and Distributed Processing Symposium (IPDPS)*, 2020.
27. Alberto Zeni, Giulia Guidi, Marquita Ellis, Nan Ding, Marco D. Santambrogio, Steven Hofmeyr, **Aydın Buluç**, Leonid Oliker, and Katherine Yelick. LOGAN: High-performance GPU-based x-drop long-read alignment. In *International Parallel and Distributed Processing Symposium (IPDPS)*, 2020.
28. Santosh Pandey, Xiaoye Sherry Li, **Aydın Buluç**, Jiejun Xu, and Hang Liu. H-index: Hash-indexing for parallel triangle counting on GPUs. In *2019 IEEE High Performance Extreme Computing Conference (HPEC)*, pages 1–7. IEEE, 2019. Graph Challenge Champion.
29. Ben Brock, **Aydın Buluç**, and Katherine Yelick. BCL: A cross-platform distributed data structures library. In *48th International Conference on Parallel Processing (ICPP)*, Kyoto, Japan, 2019.

30. Marquita Ellis, Giulia Guidi, **Aydın Buluç**, Leonid Oliker, and Katherine Yelick. diBELLA: Distributed long read to long read alignment. In *48th International Conference on Parallel Processing (ICPP)*, Kyoto, Japan, 2019.
31. Ariful Azad and **Aydın Buluç**. LACC: a linear-algebraic algorithm for finding connected components in distributed memory. In *International Parallel and Distributed Processing Symposium (IPDPS)*, 2019.
32. Evangelos Georganas, Rob Egan, Steven Hofmeyr, Eugene Goltsman, Bill Arndt, Andrew Tritt, **Aydın Buluç**, Leonid Oliker, and Katherine Yelick. Extreme scale de novo metagenome assembly. In *International Conference for High Performance Computing, Networking, Storage and Analysis (SC'18)*, 2018. Best Paper Nominee.
33. Amir Gholami, Ariful Azad, Peter Jin, Kurt Keutzer, and **Aydın Buluç**. Integrated model, batch, and domain parallelism in training neural networks. In *The 30th ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*, 2018.
34. Carl Yang, **Aydın Buluç**, and John D Owens. Implementing push-pull efficiently in GraphBLAS. In *47th International Conference on Parallel Processing (ICPP)*, 2018.
35. Carl Yang, **Aydın Buluç**, and John D Owens. Design principles for sparse matrix multiplication on the GPU. In *24th International European Conference on Parallel and Distributed Computing (Euro-Par)*, 2018. Distinguished Paper.
36. Penporn Koanantakool, Alnur Ali, Ariful Azad, **Aydın Buluç**, Dmitriy Morozov, Leonid Oliker, Katherine Yelick, and Sang-Yun Oh. Communication-avoiding optimization methods for distributed massive-scale sparse inverse covariance estimation. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 1376–1386, 2018.
37. Yang You, **Aydın Buluç**, and James Demmel. Scaling deep learning on GPU and Knights Landing clusters. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (SC'17)*, 2017.
38. Timothy G Mattson, Carl Yang, Scott McMillan, **Aydın Buluç**, and José E Moreira. GraphBLAS C API: Ideas for future versions of the specification. In *High Performance Extreme Computing Conference (HPEC), 2017 IEEE*, pages 1–6. IEEE, 2017.
39. Ariful Azad and **Aydın Buluç**. A work-efficient parallel sparse matrix-sparse vector multiplication algorithm. In *International Parallel and Distributed Processing Symposium (IPDPS)*, 2017.
40. Ariful Azad, Mathias Jacquelin, **Aydın Buluç**, and Esmond G. Ng. The reverse Cuthill-McKee algorithm in distributed-memory. In *International Parallel and Distributed Processing Symposium (IPDPS)*, 2017.
41. Marquita Ellis, Evangelos Georganas, Rob Egan, Steven Hofmeyr, **Aydın Buluç**, Brandon Cook, Leonid Oliker, and Katherine Yelick. Performance characterization of de novo genome assembly on leading parallel systems. In *23rd International European Conference on Parallel and Distributed Computing (Euro-Par)*, 2017.
42. Jeremy Kepner, Peter Aaltonen, David Bader, **Aydın Buluç**, Franz Franchetti, John Gilbert, Dylan Hutchison, Manoj Kumar, Andrew Lumsdaine, Henning Meyerhenke, Scott McMillan, José Moreira, John Owens, Carl Yang, Marcin Zalewski, and Timothy Mattson. Mathematical foundations of the GraphBLAS. In *IEEE High Performance Extreme Computing (HPEC)*, 2016.
43. Ariful Azad and **Aydın Buluç**. Distributed-memory algorithms for maximum cardinality matching in bipartite graphs. In *International Parallel and Distributed Processing Symposium (IPDPS)*, 2016.

44. Penporn Koanantakool, Ariful Azad, **Aydın Buluç**, Dmitriy Morozov, Sang-Yun Oh, Leonid Oliker, and Katherine Yelick. Communication-avoiding parallel sparse-dense matrix-matrix multiplication. In *International Parallel and Distributed Processing Symposium (IPDPS)*, 2016.
45. Evangelos Georganas, **Aydın Buluç**, Jarrod Chapman, Steven Hofmeyr, Chaitanya Aluru, Rob Egan, Leonid Oliker, Daniel Rokhsar, and Katherine Yelick. HiPMer: An extreme-scale de novo genome assembler. In *International Conference for High Performance Computing, Networking, Storage and Analysis (SC'15)*, 2015.
46. Ariful Azad and **Aydın Buluç**. Distributed-memory algorithms for maximal cardinality matching using matrix algebra. In *IEEE International Conference on Cluster Computing (CLUSTER)*, 2015.
47. Veronika Strnadova, **Aydın Buluç**, Joseph Gonzalez, Jarrod Chapman, John Gilbert, and Leonid Oliker. Efficient data reduction for large-scale genetic mapping. In *ACM Conference on Bioinformatics, Computational Biology, and Health Informatics (ACM BCB)*, 2015.
48. Ariful Azad, **Aydın Buluç**, and Alex Pothen. A parallel tree grafting algorithm for maximum cardinality matching in bipartite graphs. In *International Parallel and Distributed Processing Symposium (IPDPS)*, 2015.
49. Evangelos Georganas, **Aydın Buluç**, Jarrod Chapman, Leonid Oliker, Daniel Rokhsar, and Katherine Yelick. meraligner: A fully parallel sequence aligner. In *International Parallel and Distributed Processing Symposium (IPDPS)*, 2015.
50. V. Strnadova, **A. Buluç**, J. Gonzalez, S. Jegelka, J. Chapman, J. Gilbert, D. Rokhsar, and L. Oliker. Efficient and accurate clustering for large-scale genetic mapping. In *The IEEE International Conference on Bioinformatics and Biomedicine (BIBM'14)*, 2014.
51. Evangelos Georganas, **Aydın Buluç**, Jarrod Chapman, Leonid Oliker, Daniel Rokhsar, and Katherine Yelick. Parallel de bruijn graph construction and traversal for de novo genome assembly. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (SC'14)*, 2014.
52. H. Metin Aktulga, **Aydın Buluç**, Samuel Williams, and Chao Yang. Optimizing sparse matrix-multiple vectors multiplication for nuclear configuration interaction calculations. In *International Parallel and Distributed Processing Symposium (IPDPS)*. IEEE Computer Society, 2014.
53. Grey Ballard, **Aydın Buluç**, James Demmel, Laura Grigori, Benjamin Lipshitz, Oded Schwartz, and Sivan Toledo. Communication optimal parallel multiplication of sparse random matrices. In *SPAA 2013: The 25th ACM Symposium on Parallelism in Algorithms and Architectures*, Montreal, Canada, 2013.
54. **Aydın Buluç**, Erika Duriakova, Armando Fox, John Gilbert, Shoaib Kamil, Adam Lugowski, Leonid Oliker, and Samuel Williams. High-productivity and high-performance analysis of filtered semantic graphs. In *International Parallel and Distributed Processing Symposium (IPDPS)*. IEEE Computer Society, 2013.
55. Edgar Solomonik, **Aydın Buluç**, and James Demmel. Minimizing communication in all-pairs shortest paths. In *International Parallel and Distributed Processing Symposium (IPDPS)*. IEEE Computer Society, 2013.
56. **Aydın Buluç** and Kamesh Madduri. Graph partitioning for scalable distributed graph computations. In *Graph Partitioning and Graph Clustering (Proc. 10th DIMACS Implementation Challenge)*, volume 588 of *Contemporary Mathematics*, pages 83–101. AMS, 2013.

57. Adam Lugowski, David Alber, **Aydın Buluç**, John Gilbert, Steve Reinhardt, Yun Teng, and Andrew Waranis. A flexible open-source toolbox for scalable complex graph analysis. In *SIAM Conference on Data Mining (SDM)*, 2012.
58. **Aydın Buluç** and Kamesh Madduri. Parallel breadth-first search on distributed memory systems. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (SC'11)*, pages 65:1–65:12, 2011.
59. **Aydın Buluç**, Samuel Williams, Leonid Oliker, and James Demmel. Reduced-bandwidth multithreaded algorithms for sparse matrix-vector multiplication. In *IPDPS'11: Proceedings of the 25th IEEE International Symposium on Parallel&Distributed Processing*, pages 721–733. IEEE Computer Society, 2011.
60. **Aydın Buluç**, Jeremy T. Fineman, Matteo Frigo, John R. Gilbert, and Charles E. Leiserson. Parallel sparse matrix-vector and matrix-transpose-vector multiplication using compressed sparse blocks. In *SPAA '09: Proceedings of the 21st Annual ACM Symposium on Parallel Algorithms and Architectures*, pages 233–244, 2009.
61. **Aydın Buluç** and John R. Gilbert. Challenges and advances in parallel sparse matrix-matrix multiplication. In *ICPP'08: Proc. of the 37th Intl. Conf. on Parallel Processing*, pages 503–510, Portland, Oregon, USA, 2008. IEEE Computer Society.
62. **Aydın Buluç** and John R. Gilbert. On the representation and multiplication of hypersparse matrices. In *IPDPS'08: Proceedings of the 22nd IEEE International Symposium on Parallel&Distributed Processing*, pages 1–11. IEEE Computer Society, 2008.

WORKSHOP AND
SHORT PAPERS WITH
PROCEEDINGS
(REFEREED)

1. Md Taufique Hussain, Guttu Sai Abhishek, **Aydın Buluç**, and Ariful Azad. Parallel algorithms for adding a collection of sparse matrices. In *International Parallel and Distributed Processing Symposium Workshops (IPDPSW)*, pages 285–294. IEEE, 2022.
2. Benjamin Brock, Scott McMillan, **Aydın Buluç**, Timothy G Mattson, and José E Moreira. GraphBLAS: C++ iterators for sparse matrices. In *International Parallel and Distributed Processing Symposium Workshops (IPDPSW)*, pages 238–246. IEEE, 2022.
3. Jeremy Kepner, Michael Jones, Daniel Andersen, **Aydın Buluç**, Chansup Byun, K Claffy, Timothy Davis, William Arcand, Jonathan Bernays, David Bestor, et al. Temporal correlation of internet observatories and outposts. In *International Parallel and Distributed Processing Symposium Workshops (IPDPSW)*, pages 247–254. IEEE, 2022.
4. Giulia Guidi, Marquita Ellis, **Aydın Buluç**, Katherine Yelick, and David Culler. 10 years later: Cloud computing is closing the performance gap. In *Companion of the ACM/SPEC International Conference on Performance Engineering*, pages 41–48, 2021.
5. Benjamin Brock, **Aydın Buluç**, Timothy G Mattson, Scott McMillan, and José E Moreira. Introduction to GraphBLAS 2.0. In *International Parallel and Distributed Processing Symposium Workshops (IPDPSW)*, pages 253–262. IEEE, 2021.
6. Francesco Peverelli, Lorenzo Di Tucci, Marco D. Santambrogio, Nan Ding, Steven Hofmeyr, **Aydın Buluç**, Leonid Oliker, and Katherine Yelick. GPU accelerated partial order multiple sequence alignment for long reads self-correction. In *Intl. Parallel & Distributed Processing Symposium Workshop (IPDPSW)*, 2020.
7. Benjamin Brock, **Aydın Buluç**, Timothy G Mattson, Scott McMillan, José E Moreira, Roger Pearce, Oguz Selvitopi, and Trevor Steil. Considerations for a distributed GraphBLAS API. In *Intl. Parallel & Distributed Processing Symposium Workshop (IPDPSW)*, 2020.

8. Benjamin Brock, **Aydın Buluç**, Timothy G Mattson, Scott McMillan, and José E Moreira. A roadmap for the GraphBLAS C++ API. In *Intl. Parallel & Distributed Processing Symposium Workshop (IPDPSW)*, 2020.
9. Benjamin A Brock, Yuxin Chen, Jiakun Yan, John Owens, **Aydın Buluç**, and Katherine Yelick. RDMA vs. RPC for implementing distributed data structures. In *9th IEEE/ACM Workshop on Irregular Applications: Architectures and Algorithms (IA3)*, pages 17–22. IEEE, 2019.
10. Timothy Mattson, Timothy A. Davis, Manoj Kumar, **Aydın Buluç**, Scott McMillan, José Moreira, and Carl Yang. LAGraph: A community effort to collect graph algorithms built on top of the GraphBLAS. In *GrAPL: Workshop on Graphs, Architectures, Programming, and Learning (IPDPSW)*, 2019.
11. Muhammad Osama, Minh Truong, Carl Yang, **Aydın Buluç**, and John D Owens. Graph coloring on the GPU. In *GrAPL: Workshop on Graphs, Architectures, Programming, and Learning (IPDPSW)*, 2019.
12. Yusuke Nagasaka, Satoshi Matsuoka, Ariful Azad, and **Aydın Buluç**. High-performance sparse matrix-matrix products on Intel KNL and multicore architectures. In *47th International Conference on Parallel Processing Workshops (ICPPW)*, 2018.
13. **Aydın Buluç**, Timothy Mattson, Scott McMillan, Jose Moreira, and Carl Yang. Design of the GraphBLAS API for C. In *IEEE Workshop on Graph Algorithm Building Blocks, IPDPSW*, 2017.
14. Ariful Azad and **Aydın Buluç**. Towards a GraphBLAS library in Chapel. In *Intl. Parallel & Distributed Processing Symposium Workshop (IPDPSW)*, 2017.
15. Ariful Azad, **Aydın Buluç**, and John R Gilbert. Parallel triangle counting and enumeration using matrix algebra. In *Proceedings of the IPDPSW, Workshop on Graph Algorithm Building Blocks (GABB)*, 2015.
16. Scott Beamer, **Aydın Buluç**, Krste Asanović, and David Patterson. Distributed memory breadth-first search revisited: Enabling bottom-up search. In *Workshop on Multithreaded Architectures and Applications (MTAAP), in conjunction with IPDPS*. IEEE Computer Society, 2013.
17. Adam Lugowski, **Aydın Buluç**, John R. Gilbert, and Steve Reinhardt. Scalable complex graph analysis with the knowledge discovery toolbox. In *ICASSP, Special session on “Big Data”*, 2012.
18. K. Kandalla, **A. Buluç**, H. Subramoni, K. Tomko, J. Vienne, L. Oliker, and D. K. Panda. Can network-offload based non-blocking neighborhood mpi collectives improve communication overheads of irregular graph algorithms? In *International Workshop on Parallel Algorithms and Parallel Software (IWPAPS 2012), held in Conjunction with IEEE Cluster*, 2012.

BOOK CHAPTERS

1. Evangelos Georganas, Steven Hofmeyr, Rob Egan, **Aydın Buluç**, Leonid Oliker, Daniel Rokhsar, and Katherine Yelick. Extreme-scale de novo genome assembly. In T. Straatsma, K. Antypas, and T. Williams, editors, *Exascale Scientific Applications: Programming Approaches for Scalability, Performance, and Portability*. CRC Press, 2017.
2. **A. Buluç**, H. Meyerhenke, I. Safro, P. Sanders, and C. Schulz. Recent advances in graph partitioning. In *Algorithm Engineering - Selected Results and Surveys*, volume 9220. Lecture Notes in Computer Science, 2016.

EXTENDED
ABSTRACTS
WITHOUT
PROCEEDINGS
(REFEREED)

3. **Aydın Buluç**, John R. Gilbert, and Viral B. Shah. Implementing sparse matrices for graph algorithms. In J. Kepner and J. Gilbert, editors, *Graph Algorithms in the Language of Linear Algebra*. SIAM, Philadelphia, 2011.
4. **Aydın Buluç** and John R. Gilbert. New ideas in sparse matrix-matrix multiplication. In J. Kepner and J. Gilbert, editors, *Graph Algorithms in the Language of Linear Algebra*. SIAM, Philadelphia, 2011.
1. Srdjan Milaković, Oguz Selvitopi, Israt Nisa, Zoran Budimlić, and **Aydın Buluç**. Poster: Parallel algorithms for masked sparse matrix-matrix products. In *Proceedings of the 26th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming (PPOPP)*, 2022.
2. Marquita Ellis, **Aydın Buluç**, and Katherine Yelick. Asynchrony versus bulk-synchrony for a generalized n-body problem from genomics. In *Proceedings of the 26th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming (PPOPP)*, pages 465–466, 2021.
3. Nicolas Swenson, Aditi S Krishnapriyan, **Aydın Buluç**, Dmitriy Morozov, and Katherine Yelick. PersGNN: Applying topological data analysis and geometric deep learning to structure-based protein function prediction. In *Learning Meaningful Representations of Life (LMRL) workshop at NeurIPS*, 2020.
4. Amir Gholami, Ariful Azad, Kurt Keutzer, and **Aydın Buluç**. Communication analysis of hybrid model and data parallelism in training neural networks. In *NIPS Workshop: Deep Learning At Supercomputer Scale*, 2017.
5. Veronika Strnadova-Neeley, **Aydın Buluç**, John R Gilbert, Leonid Oliker, and Weimin Ouyang. LiRa: A new likelihood-based similarity score for collaborative filtering. In *The Workshop on Large Scale Recommendation Systems (LSRS'16), in conjunction with RecSys*, 2016.
6. D. Bader, **A. Buluç**, J. Gilbert, J. Gonzalez, J. Kepner, and T. Mattson. The Graph BLAS effort and its implications for Exascale. In *SIAM Workshop on Exascale Applied Mathematics Challenges and Opportunities (EX14)*, 2014.
7. **Aydın Buluç**, John Gilbert, and Sivan Toledo. Computing shortest paths using sparse gaussian elimination. In *SIAM Workshop on Network Science (NS14)*, 2014. (Poster presentation).
8. **Aydın Buluç**, Armando Fox, John Gilbert, Shoaib Kamil, Adam Lugowski, Leonid Oliker, and Samuel Williams. High-performance analysis of filtered semantic graphs. In *International Conference on Parallel Architectures and Compilation Techniques, PACT 2012*, Minneapolis, MN, 2012.
9. **Aydın Buluç** and John R. Gilbert. Parallel sparse matrix indexing and assignment. In *CSC'11: The 5th SIAM Workshop on Combinatorial Scientific Computing*, Darmstadt, Germany.
10. Steve Reinhardt, Adam Lugowski, John R. Gilbert, and **Aydın Buluç**. Enabling non-graph-expert use of very-large-scale graph analysis. In *CSC'11: The 5th SIAM Workshop on Combinatorial Scientific Computing*, Darmstadt, Germany.
11. **Aydın Buluç**. Parallel primitives for computation with large graphs. In *TCPP PhD Forum at the IEEE International Parallel and Distributed Processing Symposium*, Rome, Italy, 2009. (Doctoral Colloquium Poster Presentation).

TECH REPORTS AND
NON-REFEREED
PUBLICATIONS

1. Alok Tripathy, Katherine Yelick, and **Aydm Buluç**. Distributed matrix-based sampling for graph neural network training. *arXiv preprint arXiv:2311.02909*, 2023.
2. Vivek Bharadwaj, Osman Asif Malik, Riley Murray, **Aydm Buluç**, and James Demmel. Distributed-memory randomized algorithms for sparse tensor CP decomposition. *arXiv preprint arXiv:2210.05105*, 2022.
3. **Aydm Buluç**. Sparse matrices powering three pillars of science: Simulation, data, and learning. In *Proceedings of the International Symposium on Symbolic and Algebraic Computation*, pages 31–31, 2022.
4. **Aydm Buluç**, Tamara G Kolda, Stefan M Wild, Mihai Anitescu, Anthony DeGennaro, John Jakeman, Chandrika Kamath, Miles E Lopes, Per-Gunnar Martinsson, Kary Myers, et al. Randomized algorithms for scientific computing (RASC). *arXiv preprint arXiv:2104.11079*, 2021. DOE Report.
5. Farid Rashidi Mehrabadi, Kerrie L Marie, Eva Pérez-Guijarro, Salem Malikić, Erfan Sadeqi Azer, Howard H Yang, Can Kızılkale, Charli Gruen, Welles Robinson, Huaitian Liu, et al. Profiles of expressed mutations in single cells reveal subclonal expansion patterns and therapeutic impact of intratumor heterogeneity. *bioRxiv*, pages 2021–03, 2021.
6. **Aydm Buluç**, Scott Beamer, Kamesh Madduri, Krste Asanović, and David Patterson. Distributed-memory breadth-first search on massive graphs. *arXiv preprint arXiv:1705.04590*, 2017.
7. Evangelos Georganas, Marquita Ellis, Rob Egan, Steven Hofmeyr, **Aydm Buluç**, Brandon Cook, Leonid Oliker, and Katherine Yelick. Merbench: PGAS benchmarks for high performance genome assembly. In *Proceedings of the Second Annual PGAS Applications Workshop (PAW17), in conjunction with SC'17*, 2017.
8. **Aydm Buluç**, Timothy Mattson, Scott McMillan, José Moreira, and Carl Yang. The GraphBLAS C API Specification, version 1.3.0. Technical report, The GraphBLAS Signatures Subgroup, Sept 2019. http://people.eecs.berkeley.edu/aydin/GraphBLAS_API_C_v13.pdf.
9. J. Kepner, D. Bader, **A. Buluç**, J. Gilbert, T. Mattson, and H. Meyerhenke. Graphs, matrices, and the GraphBLAS: Seven good reasons. In *The International Conference on Computational Science (ICCS)*, 2015.
10. T. Mattson, D. Bader, J. Berry, **A. Buluç**, J. Dongarra, C. Faloutsos, J. Feo, J. Gilbert, J. Gonzalez, B. Hendrickson, J. Kepner, C. Leiserson, A. Lumsdaine, D. Padua, S. Poole, S. Reinhardt, M. Stonebraker, S. Wallach, and A. Yoo. Standards for graph algorithm primitives. In *High Performance Extreme Computing Conference (HPEC '13)*. IEEE, September 2013. (position paper).
11. **Aydm Buluç** and John R. Gilbert. Highly parallel sparse matrix-matrix multiplication. Technical Report UCSB-CS-2010-10, Computer Science Department, University of California, Santa Barbara, 2010.
12. **Aydm Buluç** and Erik G. Boman. Towards scalable parallel hypergraph partitioning. In *CSRI Summer Proceedings*, Albuquerque, NM and Livermore, CA, 2008.

CITATION INFO

According to Google Scholar
- Total: 7050
- H-Index: 41

FUNDING AND
OTHER GRANTS

- External Funding (Active)
 - *Lead PI (Center Director), Sparsitute: A Mathematical Institute for Sparse Computations in Science and Engineering*
DOE ASCR Applied Math MMICCs, Sep 2022-2027, \$13.5M total (\$4.5M for LBNL)
 - *Inst. Co-PI, Fundamental nuclear physics at the exascale and beyond*
DOE NP/ASCR SciDAC, Sep 2022-2027, \$1.743M for LBNL. (Inst. PI: Li, Lead PI: Edwards)
 - *PI, Scalable Graph Learning for Scientific Discovery*
DOE ASCR Applied Math Base Program, Sep 2021-2024, \$1.5M
 - *Inst. PI, ExaGraph: Combinatorial Methods for Enabling Exascale Applications*
DOE ECP Co-Design Center, FY17-23: \$2.74M for LBNL. (Lead PI: Halappanavar)
 - *Co-I, Surrogating High Dimensional Probability Distributions with Deep Learning for Scientific Inference and Data Analysis*
DOE ASCR Applied Math, Sep 2021-2024, \$2.4M. (PI: Seljak)
 - *Co-I, Exascale Solutions for Microbiome Analysis*
DOE ECP Application, FY17-23: \$11.5M. (PI: Yelick)
 - *Co-PI, Global Address Programming with Accelerators*
NSF SPX Program, October 2018-2023, \$465K for UC Berkeley. (Lead PI: Owens)
- External Funding (Past)
 - *PI, Parallel Primitives for Randomized Algorithms on Sparse Data*
DOE ASCR Applied Math Base Program, Sep 2018-2021, \$1.5M
 - *Co-PI, Dynamic Runtime Fusion of Parallel Operators for Graph Analytics and Machine Learning in GraphBLAS*
ARO, September 2019-2022, \$1.25M (PI: Iancu)
 - *PI: Energy-efficient Parallel Graph and Data Mining Algorithms*
DOE Early Career Research Award, July 2013-2018, \$2.5M
 - *Co-I: Scalable Statistics and Machine Learning for Data-Centric Science*
DOE Applied Math Grant, Oct 2013-2016, \$2M (PI: Prabhat)
 - *Co-PI: THOR: Throughput-Oriented Runtimes*
DOD, FY16-17: \$450K, FY19: \$250K (PI: Iancu)
- Internal Funding
 - *PI: Graph-based analysis and visualization of multi-modal multi-resolution large-scale neuroimaging data*
LBNL LDRD, October 2013-2015, \$543K (unburdened)
 - *Co-PI: High-Performance Parallel Analysis for Key Genomics Computations*
LBNL LDRD, October 2012-2015, \$665K (unburdened)

ADVISING

Current

- Alok Tripathy (UC Berkeley, EECS, Ph.D.), co-advisor, 2019- (w/ Kathy Yelick)
- Vivek Bharadwaj (UC Berkeley, EECS, Ph.D.), co-advisor, 2020- (w/ Jim Demmel)
- Tianyu Liang (UC Berkeley, EECS, Ph.D.), co-advisor, 2021- (w/ Jim Demmel)
- Yen-Hsiang Chang (UC Berkeley, EECS, Ph.D.), co-advisor (temp), 2023- (w/ Jim Demmel)
- Gabriel Raulet (UC Berkeley, EECS, Ph.D.), co-advisor (temp), 2023- (w/ Kathy Yelick)
- David Tench (LBNL), Hopper Postdoctoral Fellow, 2023-
- Yuxi Hong (LBNL), Postdoctoral Fellow, 2023-
- Richard Lettich (LBNL), Computer Systems Engineer, 2022-

Postdoctoral Fellows supervised at LBNL

- Helen Xu, Hopper Postdoctoral Fellow, 2022-23 (now at Georgia Tech)

- Can Kizilkale, Postdoctoral Fellow, 2019-23 (now at UC Berkeley)
- Prashant Pandey, Postdoctoral Fellow, 2019-21 (w/ Kathy Yelick) (now at University of Utah)
- Israt Nisa, Postdoctoral Fellow, 2020-21 (w/ Costin Iancu) (now at Amazon)
- Oguz Selvitopi, Postdoctoral Fellow, 2018-20 (now at LBNL)
- Saliya Ekanayake, Postdoctoral Fellow, 2018-2020 (now at d-Matrix)
- Ariful Azad, Postdoctoral Fellow, 2014-2016 (now at Indiana University)
- Sang-Yun Oh, Simons Postdoctoral Fellow, 2013-15 (w/ Lenny Oliker) (now at UC Santa Barbara)
- Harsha Vardhan Simhadri, Postdoctoral Fellow, 2013-14 (w/ Kathy Yelick)

PhD Dissertations supervised

- Giulia Guidi (UC Berkeley, EECS, Ph.D.), co-advisor, 2018-22 (w/ Kathy Yelick)
Dissertation: [Parallel Algorithms for De Novo Long Read Genome Assembly via Sparse Linear Algebra](#) (now at Cornell University).
- Benjamin Brock (UC Berkeley, EECS, Ph.D.), co-advisor, 2017-22 (w/ Kathy Yelick)
Dissertation: [RDMA-Based Distributed Data Structures for Large-Scale Parallel Systems](#) (now at Intel Labs).
- Carl Yang (UC Davis, ECE, Ph.D.), co-advisor, 2015-2019 (w/ John Owens)
Dissertation: [High-Performance Linear Algebra-based Graph Framework on the GPU](#) (now at Waymo, Inc).

Masters students

- Ujjaini Mukhopadhyay (Masters Thesis, Second Reader, EECS, UC Berkeley), 2021-23 (w/ Kathy Yelick)
- Nick Swenson, Jiali Chen, Jude Fernandes (MEng Capstone Advisor, EECS, UC Berkeley), 2019-2020 (w/ Kathy Yelick)
- Richard Barnes (Masters Thesis, Second Reader, EECS, UC Berkeley), 2018-2020 (advisor: Kathy Yelick)
- Cristina Teodoropol (Masters Thesis, Second Reader, EECS, UC Berkeley), 2018-2020 (advisor: Kathy Yelick)

Other students supervised

- Muhammed Emin Ozturk (University of Utah), Graduate Summer Intern, 2023
- Abdullah Al Raqibul Islam (University of North Carolina at Charlotte), Graduate Summer Intern, 2023
- Koby Hayashi (Georgia Tech), CSGF Summer Intern, 2021-22
- Elizabeth Koning, SUPERB REU undergraduate intern (2019) and Graduate Summer Intern (2022)
- Jiakun Yan, Undergraduate researcher, 2019
- Caitlin Whitter (Purdue), CSGF Summer Intern, 2019
- Sureyya Emre Kurt (Utah), Graduate Summer Intern, 2019
- Yusuke Nagasaka (Tokyo Institute of Technology), Graduate Summer Intern, 2017 (w/ John Shalf) (now at Fujitsu Laboratories)
- Veronika Strnadova-Neeley (UCSB), GSRA, 2013-2016 (w/ Lenny Oliker, campus advisor: John Gilbert) (now at Montana State University)
- Patrick Flick (Georgia Tech), Graduate Summer Intern, 2016
- Chaitanya Aluru (UC Berkeley), Undergraduate Research Assistant, 2014-15
- Adam Sealfon (MIT), CSGF Summer Intern, 2015
- Eric Lee (UC Berkeley), Undergraduate Summer Intern, 2014
- Adam Lugowski (UCSB), Graduate Summer Intern, 2012 (w/ L. Oliker & S. Williams)

Ph.D. Committee Membership

- Koby Hayashi (CSE, Georgia Tech), 2023-2025 (advisors: Haesun Park and Rich

Vuduc)

- Elizabeth Koning (CS, UIUC), 2023-2025 (advisor: Bill Gropp)
- Toluwanimi Odemuyiwa (ECE, UC Davis), 2023-2025 (advisor: John D. Owens)
- Yuxin Chen (ECE, UC Davis), 2018-2023 (advisor: John D. Owens)
- Marquita Ellis (EECS, UC Berkeley), 2017-2020 (advisor: Kathy Yelick)
- Fazle Sadi (ECE, CMU), 2017-2019 (advisors: Larry Pileggi and Franz Franchetti)

TEACHING

- Applications of Parallel Computers (CS267), Spring 2024 (w/ J. Demmel)
- Applications of Parallel Computers (CS267), Spring 2023 (w/ J. Demmel & L. Grigori)
- Applications of Parallel Computers (CS267), Spring 2022 (w/ J. Demmel)
- Applications of Parallel Computers (CS267), Spring 2021 (w/ J. Demmel & K. Yelick)
- Applications of Parallel Computers (CS267), Spring 2020 (w/ J. Demmel & K. Yelick)
- Applications of Parallel Computers (CS267), Spring 2019 (w/ J. Demmel & K. Yelick)
- Applications of Parallel Computers (CS267), Spring 2018 (w/ J. Demmel & K. Yelick)
- Applications of Parallel Computers (CS267), Spring 2017 (w/ K. Yelick)

INVITED TALKS AND SEMINARS

1. Sparse BLAS is not just for numerical linear algebra. In *Workshop on the design and standardization of Basic and Advanced Sparse Linear Algebra Routines*, Knoxville, Tennessee, 2023. (Keynote talk).
2. Computational journeys in a sparse universe. In *Oak Ridge National Laboratory*, Oak Ridge, Tennessee, 2023. (Invited Talk).
3. Computational journeys in a sparse universe. In *University of Utah, Data Science Seminar*, Virtual, 2023. (Invited Talk).
4. Distributed sparse matrices in graph algorithms and graph learning. In *GTA3 Workshop at IEEE BigData*, Virtual, 2022. (Keynote talk).
5. Sparse matrices beyond solvers: Graphs, biology, and machine learning. In *International Symposium of Symbolic and Algebraic Computation (ISSAC)*, Lille, France, 2022. (Invited Tutorial Lecturer).
6. Parallel sparse matrix algorithms for data analysis and machine learning. In *ETH Zurich (SPCL)*, Virtual, 2022. (Invited Seminar).
7. Large-scale graph representation learning and computational biology through sparse matrices. In *New Jersey Institute of Technology*, Virtual, 2021. (Invited Seminar).
8. Parallelism in deep neural network training, with an emphasis on graph neural networks. In *Artificial Intelligence Summer School, Turkish Academy of Sciences*, Virtual, 2021. (Invited Seminar, in Turkish).
9. Large-scale graph representation learning and computational biology through sparse matrices. In *Queen's University Belfast*, Virtual, 2021. (Invited Seminar).
10. Sparse matrices beyond solvers: Graphs, biology, and machine learning. In *XXI Householder Symposium on Numerical Linear Algebra*, Selva di Fasano (Br), Italy, 2020. (Postponed due to COVID-19).
11. Sparse Matrices Beyond Solvers: Graphs, Biology and Machine Learning. In *EECS, MIT*, Cambridge, MA, June 2020. (Invited Seminar, Online due to COVID-19, Host: Julian Shun).
12. Sparse Matrices Beyond Solvers: Graphs, Biology and Machine Learning. In *SIAM Conference on Mathematics of Data Science*, Cincinnati, OH, May 2020. (Invited Minisymposium Talk, Online due to COVID-19).

13. Sparse Matrices Beyond Solvers: Graphs, Biology and Machine Learning. In *EECS, University of California*, Irvine, CA, March 2020. (Invited Departmental Seminar, Cancelled due to COVID-19).
14. GraphBLAST: A high-performance linear algebra-based graph framework on the GPU. In *ACS HPC and Data Analytics Workshop*, Baltimore, MD, 2019. (Invited talk).
15. GraphBLAS: Concepts, algorithms, and applications. In *The 14th Scheduling for Large Scale Systems Workshop*, Bordeaux, France, June 2019. (Invited Workshop Talk).
16. Communication-avoiding sparse matrix algorithms for large graph and machine learning problems. In *SIAM Conference on Computational Science and Engineering*, Spokane, WA, 2019. (Invited Minisymposium Talk).
17. Large-scale parallel computing for computational genomics. In *Hellenic Bioinformatics 11*, Thessaloniki, Greece, November 2018. (Invited Conference Talk).
18. Communication-avoiding sparse matrix algorithms for large graph and machine learning problems. In *Workshop on New Architectures and Algorithms, IPAM (UCLA)*, Los Angeles, CA, November 2018. (Invited Workshop Talk).
19. The GraphBLAS effort and its software ecosystem. In *OpenSuCo Workshop at SC'18*, Dallas, TX, November 2018. (Invited Workshop Talk).
20. Scaling parallel graph analysis & machine learning using sparse matrix operations. In *NERSC Data Seminar*, Berkeley, CA, May 2018. Invited Seminar.
21. Scaling parallel graph analysis & machine learning using sparse matrix operations. In *Michigan State University*, East Lansing, MI, April 2018. Invited Departmental Seminar (Host: H. Metin Aktulga).
22. Genomics, graphs and the GraphBLAS. In *Graphs Across Domains (GraphXD) Workshop*, Berkeley Institute for Data Science (BIDS), Berkeley, CA, March 2018. (Invited Workshop Talk).
23. Parallel algorithms across the GraphBLAS stack. In *ACS HPC and Data Analytics Workshop*, Baltimore, MD, 2017. (Invited talk).
24. Faster parallel GraphBLAS kernels and new graph algorithms in matrix algebra. In *Google Inc.*, Mountain View, CA, November 2016. (Invited Talk).
25. Faster parallel GraphBLAS kernels and new graph algorithms in matrix algebra. In *EECS, University of California*, Berkeley, CA, October 2016. (Invited Departmental Seminar).
26. A tour of contemporary genome assembly algorithms and software. In *Workshop on Parallel Software Libraries for Sequence Analysis (pSALSA)*, at *ACM BCB*, Seattle, WA, October 2016. (Invited Keynote Talk).
27. Parallel de novo assembly of complex genomes via HipMer. In *IEEE International Workshop on High Performance Computational Biology (HiCOMB)*, at *IPDPS*, Chicago, IL, May 2016. (Invited Talk).
28. The GraphBLAS effort: Kernels, API, and parallel implementations. In *SIAM Conference on Parallel Processing for Scientific Computing*, Paris, France, 2016. (Invited Minisymposium Talk).
29. Faster parallel GraphBLAS kernels and new graph algorithms in matrix algebra. In *High-Performance Extreme Computing (HPEC)*, Waltham, MA, September 2015. (Invited Conference Talk).
30. Faster parallel GraphBLAS kernels and new graph algorithms in matrix algebra. In *HP Labs*, Palo Alto, CA, August 2015. (Invited Seminar).

31. Scalable parallel algorithms for de novo assembly of complex (meta)genomes. In *DOE Joint Genome Institute*, Walnut Creek, CA, August 2015.
32. Parallel graph algorithms & primitives. In *Swift Navigation Inc.*, San Francisco, CA, May 2015. Invited Talk (Host: Timothy Harris).
33. Scalable algorithms for complex genome assembly, alignment, and genetic mapping. In *School of CSE, Georgia Tech*, Atlanta, GA, January 2015. (Invited Departmental Seminar, Host: Srinivas Aluru).
34. Distributed-memory parallel algorithms for graph traversal & genome assembly. In *Dept. CS&E, University at Buffalo (SUNY)*, Buffalo, NY, December 2014. (Departmental Seminar, Host: Murat Demirbas).
35. Distributed-memory parallel algorithms for graph traversal & genome assembly. In *Dept. CS, University at Albany (SUNY)*, Buffalo, NY, December 2014. (Departmental Seminar, Host: Petko Bogdanov).
36. Distributed-memory parallel algorithms for graph traversal & genome assembly. In *Dept. CS, Stony Brook University (SUNY)*, Buffalo, NY, December 2014. (Departmental Seminar, Host: Leman Akoglu).
37. Reducing communication in parallel graph computations. In *Workshop on Algorithms for Modern Massive Data Sets (MMDS)*, Berkeley, CA, 2014. (Invited talk).
38. Three goals in parallel graph computations: High performance, high productivity, and reduced communication. In *College of Engineering, Montana State University*, Bozeman, MT, September 2013. (Invited Departmental Seminar, Host: Mike Wittie).
39. Three goals in parallel graph computations: High performance, high productivity, and reduced communication. In *Workshop on Parallel and Distributed Algorithms for Inference and Optimization*, Simons Institute, Berkeley, CA, 2013. (Invited talk).
40. Parallel software for high-performance and high-productivity graph analysis. In *CS Dept., Stanford University*, Stanford, CA, February 2013. (Invited Seminar at Prof. Widom's group meeting, Host: Semih Salihoglu).
41. High-performance analysis of filtered semantic graphs (in a high-productivity language). In *ICIS Workshop on Graph and Hypergraph Problems in Computational Science*, Park City, UT, 2012. (Invited talk).
42. A sustainable software stack for parallel graph analysis. In *School of CSE, Georgia Tech*, Atlanta, GA, March 2012. (Invited Departmental Seminar, Host: David Bader).
43. A sustainable software stack for parallel graph analysis. In *Discovery 2015: HPC and Cloud Computing Workshop*, Berkeley, CA, 2012. (Invited talk).
44. Parallel graph libraries: Where do we go from here? In *KDT Spring Mind Meld*, Santa Barbara, CA, 2012. (Invited Talk).
45. Scalable parallel primitives for massive graph computation. In *EECS, İstanbul Şehir University*, İstanbul, Turkey, March 2011. (Invited Departmental Seminar, Host: Ahmet Bulut).
46. Scalable parallel primitives for massive graph computation. In *Dept. CS, Purdue University*, West Lafayette, IN, April 2010. (Invited Departmental Seminar, Host: Alex Pothén).
47. Scalable parallel primitives for massive graph computation. In *MCS Division, Argonne National Laboratory*, Argonne, IL, January 2010. (Invited Seminar, Host: Mihai Anitescu).
48. Scalable parallel primitives for massive graph computation. In *CSRI, Sandia Labs*, Albuquerque, NM, January 2010. (Invited Seminar, Host: Danny Rintoul).

49. Scalable parallel primitives for massive graph computation. In *CRD, Lawrence Berkeley National Laboratory*, Berkeley, CA, January 2010. (Invited Seminar, Host: Juan Meza).
50. Scalable parallel primitives for massive graph computation. In *Sandia Labs*, Livermore, CA, November 2009. (Invited Seminar, Host: Tammy Kolda).

SOFTWARE

The Combinatorial BLAS

A parallel library for running large-scale graph algorithms on distributed-memory architectures. The library consists of a growing set of optimized primitives such as sparse matrix-sparse matrix multiplication and sparse matrix-sparse vector multiplication.

HipMCL

A high-performance parallel implementation of the Markov Cluster (MCL) algorithm for large-scale network clustering. It can easily cluster a network of 75 million nodes with 68 billion edges in 2.4 hours using 2000 nodes of Cori supercomputer at NERSC.

(Meta)HipMer

An extreme-scale de novo genome assembler for large complex genomes and metagenomes. Primarily written in UPC with some MPI/C++ pieces. Scales over O(10,000) processors.

CAGNET

A family of parallel algorithms for training GNNs that can asymptotically reduce communication compared to previous parallel GNN training methods. CAGNET algorithms are based on 1D, 1.5D, 2D, and 3D sparse-dense matrix multiplication, and are implemented with torch.distributed on GPU-equipped clusters.

PASTIS

A distributed-memory parallel many-to-many protein sequence aligner that uses sparse matrices.

mergeSpMM

Sparse matrix multi-vector (aka tall-skinny dense matrix) products on the GPU, also used in [PyTorch Geometric](#).

BELLA

A computationally-efficient and highly-accurate long-read to long-read aligner and over-lapper for DNA sequences.

GraphBLAST

A high-performance linear algebra-based graph framework on the GPU, which closely approximates the [GraphBLAS API](#).

MS-BFS-Graft

Multithreaded OpenMP code for computing maximum cardinality matching on bipartite graphs. Performs multi-source breadth-first search with tree-grafting for exploiting parallelism.

HP-CONCORD and SpDM³

SpDM³ does communication-avoiding Sparse-Dense Matrix-Matrix Multiplication on distributed-memory parallel computers and HP-CONCORD is high-performance inverse covariance matrix estimation using the CONCORD-ISTA algorithm.

Compressed Sparse Blocks

Shared memory Cilk++ implementation of sparse matrix-vector and sparse matrix-transpose-vector multiplication using compressed sparse blocks.

BubbleCluster and MarkerReduce

Tools for large-scale genetic mapping. BubbleCluster is an algorithm used to efficiently cluster genetic markers into linkage groups. MarkerReduce is an algorithm for efficient data reduction of genetic marker data. Both codes are written in C++.

Knowledge Discovery Toolbox (KDT)

High-performance parallel graph analysis and mining in a very-high-level language (Python), based on the Combinatorial BLAS.

GPU-APSP

A CUDA program that computes the distances for all-pairs shortest paths in a dense directed graph using the R-Kleene algorithm.

SERVICES AND LEADERSHIP SKILLS

- *Founding Associate Editor (2013-2020)*: ACM Transactions on Parallel Computing
- *Organizing Committee Leadership*:
 - *Co-Chair (2023)*: **JRG70: Workshop celebrating the scientific contributions of John R. Gilbert**
 - *Co-Chair (2022)*: **1st ACDA Workshop in Aussois**
 - *Co-Chair (2020-21)*: **RASC: Randomized Algorithms for Scientific Computing** (DOE Sponsored Workshop)
 - *Co-Chair (2019)*: ASCR Applied Mathematics Principal Investigators Meeting
 - **2019: 1st Applied Math Visioning Workshop: The Future of Machine Learning and Data Analytics Across the Department of Energy**
- *Guest Editor*:
 - ACM Transactions on Parallel Computing, 2021, Special issue on ‘Special Issue for SPAA 2018’
 - Parallel Computing (Elsevier), 2015, Special issue on ‘Graph Analysis for Scientific Discovery’
- *Program Committee Leadership*
 - *Chair, Algorithms Track (2018)*: ACM/IEEE Intl. Conf. for High Performance Computing, Networking, Storage and Analysis (SC)
 - *Vice-chair, Applications Track (2017)*: ACM/IEEE Intl. Conf. for High Performance Computing, Networking, Storage and Analysis (SC)
 - *Co-chair (2017)*: IEEE Graph Algorithms Building Blocks (GABB) workshop at IPDPS
 - *Publications Chair (2018)*: SIAM Workshop on Combinatorial Scientific Computing (CSC).
- *Organizing Committee Member*:
 - **Workshop on Sparse Tensor Computations**, 2023
 - SIAM Conference on Parallel Processing for Scientific Computing, 2020
- *Program Committee Member*:
 - ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming (PPOPP), 2023-24
 - ACM Symposium on Parallelism in Algorithms and Architectures (SPAA), 2018, 2022, 2024
 - ACM/IEEE Intl. Conf. for High Perf. Computing, Networking, Storage and Analysis (SC), 2013-14, 2016, 2020, 2024
 - International Conference on Parallel Architectures and Compilation Techniques (PACT), 2022
 - IEEE Intl. Parallel & Dist. Processing Symp. (IPDPS), 2013-2019, 2021-22
 - SIAM Conference on Applied and Computational Discrete Algorithms (ACDA),

2021

- AAAI Fall 2020 Symposium on Physics-Guided AI to Accelerate Scientific Discovery
- SIAM Conference on Parallel Processing for Scientific Computing, 2020
- Workshop on Irregular Applications, Architectures & Algorithms (IA³), 2017-21
- The 21st Meeting on Algorithm Engineering and Experiments (ALENEX), 2019
- ACM Symposium on High-Performance Parallel and Distributed Computing (HPDC), 2018
- International Conference on Parallel Processing (ICPP), 2013, 2020, 2021
- IEEE Intl. Conf. on High Performance Comp., Data, & Analytics (HiPC), 2019
- Workshop on Graphs, Architectures, Programming, and Learning (GrAPL), 2019-21
- IEEE High Performance Extreme Computing Conference (HPEC), 2016, 2019
- ACM Conf. on Bioinformatics, Comp. Biology, & Health Informatics (BCB), 2017
- IEEE Cluster, 2017
- PLDI ARRAY Workshop on Libraries, Languages & Compilers for Prog., 2017
- ACM International Conference on Supercomputing (ICS), 2015
- IEEE Graph Algorithms Building Blocks (GABB) workshop at IPDPS, 2015-16.
- High Performance Graph Processing (HPGP) workshop at HPDC, 2016
- SIAM Workshop on Combinatorial Scientific Computing (CSC), 2016
- IEEE Workshop on High Performance Computational Biology (HICOMB), 2016, 2020-21
- Intl. Symposium on Computer Architecture and HPC (SBAC-PAD), 2015
- WWW Workshop on Big Graph Mining, 2014
- Workshop on High Perf Big Graph Data Mgmt, Analysis & Mining (at BigData'14)
- BAŞARIM, Ankara, Turkey, 2012, 2015
- International Conference on Contemporary Computing (IC3), 2011
- Graduate Student Workshop on Computing, UCSB, 2009.
- *Steering Committee:*
 - IEEE Workshop on Graph Algorithms Building Blocks (GABB) at IPDPS, 2014-2018.
 - Workshop on Graphs, Architectures, Programming, and Learning (GrAPL) at IPDPS, 2019-
- *Society Service:*
 - SIAM Membership Committee, 2022-2024
- *Conference Tutorial (Co-)organizer:*
 - “Building Graph Algorithms with the GraphBLAS”, Minitutorial at SIAM Conference on Computational Science and Engineering, 2021.
 - “A Hands-On Introduction to GraphBLAS” at the International Conference on Supercomputing (ICS), 2021.
- *Minisymposia (Co-)organizer:*
 - “Large Scale Graph Analytics”, at SIAM Conference on Parallel Processing for Scientific Computing, 2024.
 - “Sparse Computations in Science and Engineering”, at SIAM Conference on Computational Science and Engineering, 2023.
 - “GraphBLAS: Tools, Algorithms, and Applications”, at SIAM Conference on Computational Science and Engineering, 2021.
 - “Combinatorial Scientific Computing” at International Congress on Industrial and Applied Mathematics (ICIAM), 2019.
 - “GraphBLAS: Graph Algorithms in the Language of Linear Algebra” at SIAM Annual Meeting, 2016
 - “Combinatorial Scientific Computing” at SIAM Conference on Parallel Processing for Scientific Computing, 2016.
 - “Graph Analysis for Scientific Discovery” at SIAM Conference on Parallel Pro-

- censing for Scientific Computing, 2014.
- “Minimizing Communication in Scientific Computing” at SIAM Conference on Computational Science and Engineering, 2013.
- “Reducing Communication in Linear Algebra” at SIAM Conference on Applied Linear Algebra, 2012.
- “Optimization Challenges in the Energy Sector” at SIAM Conference on Optimization, 2011.
- *Poster Committee Member:*
 - ACM/IEEE Intl. Conf. for High Performance Computing, Networking, Storage and Analysis (SC), 2016
- *Grant Review Panels:* NSF CISE (many times), DOE ASCR Early Career program, DOE BER Genomics program.
- *Reviewer:*
 - *Books:* MIT Press.
 - *Journals:* Communications of the ACM (2), SIAM Review, SIAM Journal of Scientific Computing (6), Nature Computational Science (2), Bioinformatics (3), SIAM Journal on Matrix Analysis and Applications, IEEE Transactions on Parallel and Distributed Systems (4), ACM Transactions on Mathematical Software (5), The International Journal of High Performance Computing Applications (3), SIAM Journal on Mathematics of Data Science (SIMODS), Parallel Computing (11+), ACM Transactions on Parallel Computing, Computers & Electrical Engineering, ACM Transactions on Knowledge Discovery from Data (TKDD), Journal of Parallel and Distributed Computing (2), IEEE Transactions on Knowledge and Data Engineering (2), ACM Journal on Experimental Algorithmics (2), ACM Transactions on Architecture and Code Optimization, Concurrency and Computation: Practice and Experience, Computers and Electrical Engineering, Current Genomics
 - *Conferences:* SPAA 2014/2016/2019, IPDPS 2010/2012, Euro-Par 2011/2013, ICPP 2011, SC 2011/2012, ICS 2012.
 - *Grants:* DOE Office of Science (ASCR) (since 2013), DOE/ASCR SBIR/STTR program (2023)
- *Founding President*, Mathematics Club, Sabancı University
- *LBNL Service:*
 - Computing Sciences Area Staffing Committee, 2023-
 - Luis W. Alvarez Postdoctoral Fellowship selection committee, 2014-16, 2018-19
- *DOE Service:*
 - *2019 Round Table on Data for Artificial Intelligence Workshop*, Rockville, MD, June 2019
 - *Scientific Machine Learning Workshop*, North Bethesda, MD, January 2018
 - *Exascale Requirements Review for Biological and Environmental Research*, Rockville, MD, March 2016
 - *ASCR Data Management, Visualization, and Analysis of Experimental and Observational Data (EOD) Workshop*, Bethesda, MD, September 2015.
 - *DOE Data Council Meeting*, Rockville, MD, September 2014

CONTRIBUTED ORAL PRESENTATIONS

1. The ubiquitous sparse matrix-matrix products. In *Joint Mathematics Meetings (JMM)*, San Francisco, California, 2024. (talk at special session on Combinatorics for Science).
2. The ubiquitous sparse matrix-matrix products. In *JRG70: Workshop celebrating the scientific contributions of John R. Gilbert*, Santa Barbara, California, 2023.
3. The ubiquitous sparse matrix-matrix products. In *SIAM Conference on Computational Science and Engineering*, Amsterdam, Netherlands, 2023. (Minisymposium Talk).

4. Sparse matrices in biology and machine learning. In *1st ACDA Workshop at Aussois*, Aussois, France, 2022.
5. Sparse matrices for scientific data analysis and machine learning. In *Alvarez Seminar Series*, Virtual, 2022.
6. Sparse matrices for scaling data-intensive problems on distributed-memory systems. In *SIAM Conference on Computational Science and Engineering*, Virtual, 2021. (Minisymposium Talk).
7. Communication-avoiding and memory-constrained sparse matrix-matrix multiplication at extreme scale. In *The IEEE International Symposium on Parallel and Distributed Processing*, Virtual, 2021. (Conference Talk).
8. Parallelism in deep neural network training, with an emphasis on graph neural networks. In *LBNL Computing Sciences Summer Student Program*, Virtual, 2021.
9. Sparse matrices for scaling data-intensive problems on distributed-memory systems. In *Exascale Computing Project (ECP) 5th Annual Meeting*, Virtual, 2021. Exagraph Tutorial Session.
10. Sparse matrices beyond solvers: Graphs, biology, and machine learning. In *LBNL Computing Sciences Summer Student Program*, Virtual, 2020.
11. GraphBLAST: A high-performance C++ GPU library implementing GraphBLAS. In *Exascale Computing Project (ECP) 4th Annual Meeting*, Houston, TX, 2020. Exagraph Tutorial Session.
12. GraphBLAS on various problem domains. In *International Congress on Industrial and Applied Mathematics (ICIAM)*, Valencia, Spain, 2019. (Minisymposium Talk).
13. Graph algorithms, computational motifs, and GraphBLAS. In *Exascale Computing Project (ECP) 3rd Annual Meeting*, Houston, TX, 2019. Exagraph Tutorial Session.
14. Communication-avoiding sparse-matrix primitives for parallel machine learning. In *Sparse Days Meeting*, Toulouse, France, 2018.
15. Matrices are all you need. In *NECSTLab/LBNL/UCB Workshop*, Berkeley, CA, 2018.
16. Graph algorithms, computational motifs, and GraphBLAS. In *Exascale Computing Project (ECP) 2nd Annual Meeting*, Knoxville, TN, 2018. Exagraph Tutorial Session.
17. The reverse Cuthill-McKee algorithm in distributed-memory. In *The IEEE International Symposium on Parallel and Distributed Processing*, Orlando, FL, 2017. (Conference Talk).
18. Developing a PGAS graph library using GraphBLAS primitives. In *SIAM Annual Meeting*, Boston, MA, 2016. (Minisymposium Talk).
19. Faster and more scalable sparse matrix-matrix multiplication. In *SIAM Conference on Parallel Processing for Scientific Computing*, Paris, France, 2016. (Minisymposium Talk).
20. Scalable parallel algorithms for de novo assembly of complex genomes. In *The Resurgence of Reference Quality Genome Sequence, Plant & Animal Genome (PAG) Conference*, San Diego, CA, 2016.
21. The GraphBLAS effort and new parallel algorithms for bipartite graph matching. In *LBNL-Tsukuba Joint Meeting*, Berkeley, CA, 2015.
22. Scalable parallel algorithms for de novo assembly of complex (meta)genomes. In *MANTISSA Day*, Berkeley, CA, 2015.
23. The Graph BLAS: building blocks for graph algorithms in the language of linear algebra. In *Seminar on High-performance Graph Algorithms and Applications in Computational Science*, Dagstuhl, Germany, 2014.

24. Libraries and algorithms for graph analytics at scale. In *ACS Productivity Workshop*, Baltimore, MD, 2014.
25. The Graph BLAS effort and its implications for Exascale. In *SIAM Workshop on Exascale Applied Mathematics Challenges and Opportunities (EX14)*, Chicago, IL, 2014.
26. Communication-avoiding linear-algebraic primitives for graph analytics. In *Graph Algorithms Building Blocks (GABB'2014), IPDPS Workshops*, Phoenix, AZ, 2014.
27. High-productivity and high- performance analysis of filtered semantic graphs. In *SIAM Conference on Parallel Processing for Scientific Computing*, Portland, OR, 2014. (Minisymposium Talk).
28. High-productivity and high- performance analysis of filtered semantic graphs. In *Seminar on Automatic Application Tuning for HPC Architectures*, Dagstuhl, Germany, 2013.
29. Reducing communication in parallel graph computations. In *DOE ASCR Applied Math PI Meeting*, Albuquerque, NM, 2013.
30. High-performance filtered queries in attributed semantic graphs. In *SIAM Conference on Computational Science and Engineering*, Boston, MA, 2013. (Minisymposium Talk).
31. Communication-avoiding sparse matrix-matrix multiplication. In *SIAM Conference on Applied Linear Algebra*, Valencia, Spain, 2012. (Minisymposium Talk).
32. Parallel algorithms for sparse matrix product, indexing, and assignment. In *Scientific Computing and Matrix Computations (LAPACK) Seminar*, UC Berkeley, CA, 2012.
33. Parallel breadth-first search on distributed memory systems. In *SC'11: The 24th International Conference for High Performance Computing, Networking, Storage and Analysis*, Seattle, WA. (Conference Talk).
34. Tuning and communication reduction for graph and sparse matrix computations. In *CScADS Workshop on Libraries and Autotuning for Extreme Scale Applications*, Tahoe City, CA, 2011.
35. High-performance combinatorial algorithms for the analysis of the electric power grid. In *ICIAM 2011: 7th International Congress on Industrial and Applied Mathematics*, Vancouver, BC, Canada, 2011. (Invited Minisymposium Talk).
36. Parallel sparse matrix indexing and assignment. In *CSC'11: The 5th SIAM Workshop on Combinatorial Scientific Computing*, Darmstadt, Germany. (Contributed Lecture).
37. A parallel framework for identifying vulnerabilities in the electric power grid. In *SIAM Conference on Optimization*, Darmstadt, Germany, 2011. (Minisymposium Talk).
38. Building blocks for scalable graph and data mining software. In *SIAM Conference on Parallel Processing for Scientific Computing*, Seattle, WA, 2010. (Minisymposium Talk).
39. Parallel sparse matrix-vector and matrix-transpose-vector multiplication using compressed sparse blocks. In *SPAA 2009: The 21st ACM Symposium on Parallelism in Algorithms and Architectures*, Calgary, Canada. (Conference Talk).
40. Parallel Combinatorial BLAS and applications in graph computation. In *SIAM Annual Meeting*, Denver, CO, 2009. (Minisymposium Talk).
41. Parallel primitives for computation with large graphs. In *T CPP PhD Forum at the IEEE International Parallel and Distributed Processing Symposium*, Rome, Italy, 2009. (Refereed Doctoral Colloquium Poster Presentation).

42. Sparse matrix-matrix multiplication for accelerating parallel graph computations. In *SIAM Conference on Computational Science and Engineering*, Miami, FL, 2009. (Invited Minisymposium Talk).
43. Challenges and advances in parallel sparse matrix-matrix multiplication. In *International Conference on Parallel Processing*, Portland, OR, 2008. (Conference Talk).
44. Gaussian elimination based algorithms on the GPU. In *PMAA'08: The 5th International Workshop on Parallel Matrix Algorithms and Applications*, Neuchatel, Switzerland. (Contributed Talk).
45. On the representation and multiplication of hypersparse matrices. In *The IEEE International Symposium on Parallel and Distributed Processing*, Miami, FL, 2008. (Conference Talk).
46. Parallel primitives for computation with large graphs. In *SIAM Conference on Parallel Processing for Scientific Computing*, Atlanta, GA, 2008. (Minisymposium Talk).

GUEST LECTURING

- High Performance Parallel Computing (CSE 6230), Georgia Tech, Spring 2024
 - Computational journeys in a sparse universe
- High Performance Parallel Computing (CSE 6230), Georgia Tech, Spring 2023
 - HPC and Graph Algorithms
- Software Engineering for Scientific Comp. (CS294-73), UC Berkeley, Fall 2017
 - Graph algorithms, motifs, and a little bit of parallelism
- Applications of Parallel Computers (CS267), UC Berkeley, Spring 2012-16
 - Parallel Graph Algorithms
- Communication Avoiding Algorithms (CS294), UC Berkeley, Fall 2011
 - Communication in sequential and parallel BFS
- Parallel Scientific Computing (CS140), UC Santa Barbara, Winter 2009-10
 - Parallel matrix multiplication
 - Shared-memory and multicore programming
 - Divide and conquer examples in Cilk Plus
 - Numerical examples and hyperobjects in Cilk Plus
- Sparse Matrix Algorithms (CS219), UC Santa Barbara, Fall 2009.
 - Parallel matrix-vector multiplication
 - Parallel sparse matrix and graph computation
- Applied Parallel Computing (CS240), UC Santa Barbara, Spring 2009
 - Shared-memory and multicore programming
 - Breadth-first search in Cilk Plus
 - Divide and conquer examples using Cilk Plus
 - Numerical examples using Cilk Plus
- Programming Methods (CS20), UC Santa Barbara, Fall 2006.
 - Heapsort and Mergesort

COMPUTER SKILLS

Programming Languages

C, C++ (preferred), C#, Java, Prolog, Scheme

Parallel Programming

MPI, UPC, Intel Cilk Plus, Chapel, OpenMP, Pthreads, CUDA.

Scripting Languages

MATLAB, Perl, Python, Javascript, PHP.

TEACHING ASSISTANTSHIP	Graduate: Advanced Operating Systems Upper Division: Programming Languages, Formal Languages and Automata Lower Division: Discrete Mathematics, Programming Methods
UNDERGRADUATE EMPLOYMENT	<ul style="list-style-type: none"> • <i>Student Peer Assistant, Academic Support Program</i>, Sabancı University <ul style="list-style-type: none"> - Lead Coordinator for the Linear Algebra course. - Assistant for Calculus and Science of Nature courses. • <i>Undergraduate Researcher</i>, Sabancı University <ul style="list-style-type: none"> - Human Language and Speech Technologies Laboratory, 2003 <i>Building morphological and statistical models for Turkic languages</i> - Computer Vision and Pattern Analysis Laboratory, 2004 <i>Optical character recognition using decision trees</i>
PROFESSIONAL SOCIETY MEMBERSHIPS	<ul style="list-style-type: none"> • Society of Industrial and Applied Mathematics (SIAM) <ul style="list-style-type: none"> - SIAG on Supercomputing - SIAG on Computational Science and Engineering - SIAG on Data Mining and Analytics - SIAG on Linear Algebra • Association of Computing Machinery (ACM) <ul style="list-style-type: none"> - SIGHPC: Special Interest Group on High-Performance Computing • Institute of Electrical and Electronics Engineers (IEEE)
SELECTED OTHER PARTICIPATIONS	<ul style="list-style-type: none"> - <i>Emerging Solutions in Petabyte Scale Sequence Search Workshop</i>, jointly organized by DOE and NIH. - <i>Broader Engagement Program</i>, Guided Affinity Group Leader, SIAM CSE, 2021 - <i>Early Career Panel at SIAM CSE'19</i>, Panelist, Spokane, WA, 2019 - <i>IA³ Workshop at SC'18</i>, Panelist, Dallas, TX, 2018 - <i>HPC for Undergrads at SC'18</i>, Panelist, Dallas, TX, 2018 - <i>CCC Workshop: Digital Computing Beyond Moore's Law</i>, San Francisco, CA, 2018 - <i>BOF on HPC Graph Toolkits and GraphBLAS Forum</i>, Co-organizer, SC, 2017-18. - <i>Partitioned Global Address Space (PGAS) BOF</i>, Panelist, SC, 2016 - <i>Panel discussion moderator</i>, GABB@IPDPS, 2016. - <i>BOF (birds of a feather) session leader of GraphBLAS forum</i>, HPEC, 2014-15. - <i>High-Performance Computing on Graphs</i>, Panelist, SIAM PP, 2014 - <i>Intel workshop on parallel algorithms for non-numeric computing</i>, Invited participant, Santa Clara, CA, 2011/2012. - <i>Workshop on scalable graph libraries</i>, Invited participant, Atlanta, GA, 2011 - <i>Workshop on optimization in energy systems, Institute for computing in science</i>, Invited participant, Snowbird, UT, 2010. - <i>CSCAPES workshop</i>, Invited participant, Santa Fe, NM, 2008. - Summer School of Turkish Mathematical Society, Amasra, Turkey, 2005. - Civic involvement project: Kayisdagi elderly people project, 2002-2003
LANGUAGES	Turkish (native), English (fluent), Spanish (intermediate)