

## Aydın Buluç

---

CONTACT INFORMATION	Lawrence Berkeley National Laboratory, Mailstop 50F1650 1 Cyclotron Road Berkeley, CA 94720 USA	Office: +1 (510)-486-5197 Mobile: +1 (510)-325-9659 e-mail: <a href="mailto:abuluc@lbl.gov">abuluc@lbl.gov</a> <a href="https://people.eecs.berkeley.edu/~aydin/">https://people.eecs.berkeley.edu/~aydin/</a>
RESEARCH INTERESTS	Parallel Computing, Combinatorial Scientific Computing, Sparse Matrices, Computational Biology, High-Performance Computing, Graph Algorithms, Machine Learning	
EDUCATION AND TRAINING	<b>UC Center for Executive Education</b> , Berkeley, CA, USA LBNL Leadership Development Program for Emerging Leaders, 2015-16  <b>University of California</b> , Santa Barbara, CA, USA Ph.D. in Computer Science, 2010 <i>Thesis: <a href="#">Linear Algebraic Primitives for Parallel Computing on Large Graphs</a></i> <i>Advisor: John R. Gilbert</i> M.S. in Computer Science, 2009  <b>Massachusetts Institute of Technology</b> , Cambridge, MA, USA Visiting Student at the Department of Mathematics, 2008  <b>Sabanci University</b> , Istanbul, Turkey B.S. in Computer Science, 2005 (with Minors in Mathematics)	
RESEARCH AND PROFESSIONAL EXPERIENCE	<b>CRD, Lawrence Berkeley National Laboratory</b> , Berkeley, CA - Computational Staff Scientist (Career), <i>March 2016 - Present</i> - Computational Research Scientist (Career), <i>April 2014 - March 2016</i> - Computational Research Scientist (Career-track), <i>April 2012 - March 2014</i> - <a href="#">Luis W. Alvarez Postdoctoral Fellow</a> , <i>April 2010 - April 2012</i> (a) Parallel algorithms and software for graph computations (b) Communication-avoiding algorithms and autotuning for sparse-matrix kernels (c) High-performance computational genomics  <b>EECS Department, University of California</b> , Berkeley, CA - Adjunct Assistant Professor, <i>January 2017 - Present</i> - Affiliated Researcher, <i>April 2010 - Present</i> <a href="#">Bebop (BERkeley Benchmarking and OPTimization) group</a>  <b>Simons Institute for the Theory of Computing</b> , Berkeley, CA Visiting Scientist (Long Term Participant), <i>August 2013 - December 2013</i> <a href="#">Program on Theoretical Foundations of Big Data Analysis</a>  <b>CSRI, Sandia National Laboratory</b> , Albuquerque, NM Graduate Research Intern, <i>June 2008 - August 2008</i>  <b>Citrix Online</b> , Santa Barbara, CA Software Engineer Intern, <i>June 2006 - September 2006</i>  <b>Garanti Technology</b> , Istanbul, Turkey Software Engineer Intern, <i>July 2004 - September 2004</i>  <b>Computer Science Department, University of California</b> , Santa Barbara, CA Research Assistant, <a href="#">Combinatorial Scientific Computing Laboratory</a> , <i>2007-2010</i>	

Teaching Assistant, 2005-2007

SCHOLARSHIPS &  
AWARDS

- Distinguished Paper, 24th International European Conference on Parallel and Distributed Computing (Euro-Par 2018)
- *IEEE TCSC Award for Excellence for Early Career Research* by the IEEE Technical Committee on Scalable Computing, 2015
- *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
- *DOE Early Career Award* by Office of Science, Department of Energy, 2013
- *Luis W. Alvarez Postdoctoral Fellowship in Computational Science* by Lawrence Berkeley National Laboratory, 2010
- *J. H. Wilkinson Fellowship in Scientific Computing* by Argonne National Laboratory, 2010 (declined)
- *Lawrence Berkeley National Laboratory Spot Award* For keeping NERSC at the top 2 of Graph500 list for three consecutive terms, 2012.
- SIAM Travel Grant award to attend SIAM ALA 2012
- TCPP Travel Grant award to attend the IPDPS conference, 2009
- Full financial support (tuition, fees, monthly stipend) for graduate studies, Computer Science Department, UCSB, 2005-2010
- Merit scholarship support (tuition, fees, accommodation) for undergraduate studies, Sabancı University, 2001-2005.

JOURNAL  
PUBLICATIONS  
(REFEREED)

1. Ariful Azad, Georgios A. Pavlopoulos, Christos A. Ouzounis, Nikos C. Kyrpidis, and **Aydin Buluç**. HipMCL: A high-performance parallel implementation of the Markov clustering algorithm for large-scale networks. *Nucleic Acids Research (NAR)*, 2018.
2. Ariful Azad, **Aydin 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.
3. Hasan Metin Aktulga, Md. Afibuzzaman, Samuel Williams, **Aydin Buluç**, Meiyue Shao, Chao Yang, Esmond 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.
4. Ariful Azad, Grey Ballard, **Aydin 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.
5. Ariful Azad and **Aydin Buluç**. A matrix-algebraic formulation of distributed-memory maximal cardinality matching algorithms in bipartite graphs. *Parallel Computing*, 2016.
6. Jarrod Chapman, Martin Mascher, **Aydin 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.
7. Adam Lugowski, Shoaib Kamil, **Aydin 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.

8. **Aydın 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.
9. **Aydın 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.
10. **Aydın 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. 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.
2. Carl Yang, **Aydın Buluç**, and John D Owens. Implementing push-pull efficiently in GraphBLAS. In *47th International Conference on Parallel Processing (ICPP)*, 2018.
3. 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.
4. 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.
5. 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.
6. 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.
7. Ariful Azad and **Aydın Buluç**. A work-efficient parallel sparse matrix-sparse vector multiplication algorithm. In *Proceedings of the IPDPS*, 2017.
8. Ariful Azad, Mathias Jacquelin, **Aydın Buluç**, and Esmond G. Ng. The reverse Cuthill-McKee algorithm in distributed-memory. In *Proceedings of the IPDPS*, 2017.
9. 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.
10. 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.
11. Ariful Azad and **Aydın Buluç**. Distributed-memory algorithms for maximum cardinality matching in bipartite graphs. In *Proceedings of the IPDPS*, 2016.
12. 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 *Proceedings of the IPDPS*, 2016.

13. 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.
14. Ariful Azad and **Aydın Buluç**. Distributed-memory algorithms for maximal cardinality matching using matrix algebra. In *IEEE International Conference on Cluster Computing (CLUSTER)*, 2015.
15. 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.
16. Ariful Azad, **Aydın Buluç**, and Alex Pothén. A parallel tree grafting algorithm for maximum cardinality matching in bipartite graphs. In *Proceedings of the IPDPS*, 2015.
17. Evangelos Georganas, **Aydın Buluç**, Jarrod Chapman, Leonid Oliker, Daniel Rokhsar, and Katherine Yelick. meraligner: A fully parallel sequence aligner. In *Proceedings of the IPDPS*, 2015.
18. 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.
19. 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.
20. H. Metin Aktulga, **Aydın Buluç**, Samuel Williams, and Chao Yang. Optimizing sparse matrix-multiple vectors multiplication for nuclear configuration interaction calculations. In *Proceedings of the IPDPS*. IEEE Computer Society, 2014.
21. 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.
22. **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 *Proceedings of the IPDPS*. IEEE Computer Society, 2013.
23. Edgar Solomonik, **Aydın Buluç**, and James Demmel. Minimizing communication in all-pairs shortest paths. In *Proceedings of the IPDPS*. IEEE Computer Society, 2013.
24. **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.
25. 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.
26. **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.

27. **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.
28. **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.
29. **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.
30. **Aydın Buluç** and John R. Gilbert. On the representation and multiplication of hyper-sparse 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. 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.
2. **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.
3. Ariful Azad and **Aydın Buluç**. Towards a GraphBLAS library in Chapel. In *Intl. Parallel & Distributed Processing Symposium Workshop (IPDPSW)*, 2017.
4. 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.
5. 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.
6. 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.
7. 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 (to appear).
2. **Aydın Buluç**, Scott Beamer, Kamesh Madduri, Krste Asanović, and David Patterson. Distributed-memory breadth-first search on massive graphs. In D. Bader, editor, *Parallel Graph Algorithms*. CRC Press / Taylor-Francis, 2017 (to appear).

3. **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.
4. **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.
5. **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.

EXTENDED  
ABSTRACTS  
WITHOUT  
PROCEEDINGS  
(REFEREED)

1. 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.
2. Veronika Strnadova-Neeley, Aydın Buluc, 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 (LSRS16), in conjunction with RecSys*, 2016.
3. 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.
4. **Aydın Buluç**, John Gilbert, and Sivan Toledo. Computing shortest paths using sparse gaussian elimination. In *SIAM Workshop on Network Science (NS14)*, 2014. (Poster presentation).
5. **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.
6. **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.
7. 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.
8. **Aydın Buluç**. Parallel primitives for computation with large graphs. In *TCPD 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. Ariful Azad, Aydın Buluc, Xiaoye S Li, Xinliang Wang, and Johannes Langguth. A distributed-memory approximation algorithm for maximum weight perfect bipartite matching. *arXiv preprint arXiv:1801.09809*, 2018.
2. Evangelos Georganas, Marquita Ellis, Rob Egan, Steven Hofmeyr, **Aydın 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.
3. **Aydın Buluç**, Timothy Mattson, Scott McMillan, José Moreira, and Carl Yang. The GraphBLAS C API Specification, version 1.0.0. Technical report, The GraphBLAS Signatures Subgroup, May 2017. [http://graphblas.org/aydin/GraphBLAS\\_API.C.pdf](http://graphblas.org/aydin/GraphBLAS_API.C.pdf).

4. 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.
5. 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).
6. **Aydın 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.
7. **Aydın 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 as of May 2018 (first publication appeared in 2008)

- Total: 2098
- H-Index: 23

#### FUNDING AND OTHER GRANTS

- External Funding
  - *PI: Energy-efficient Parallel Graph and Data Mining Algorithms*  
DOE Early Career Research Award, July 2013-2018, \$2.5M
  - *Inst. PI: ExaGraph: Combinatorial Methods for Enabling Exascale Applications*  
DOE ECP Co-Design Center, FY17-18: \$900K for LBNL. (Lead PI: Halappanavar)
  - *Co-PI: Exascale Solutions for Microbiome Analysis*  
DOE ECP, FY17-18: \$1.3M. (PI: Yelick)
  - *Co-PI: THOR: Throughput-Oriented Runtimes*  
DOD, FY16-17: \$450K (PI: Iancu)
  - *Co-PI: Scalable Statistics and Machine Learning for Data-Centric Science*  
DOE Applied Math Grant, Oct 2013-2016, \$2M (PI: Prabhat)
- 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)
- Compute Time
  - NERSC production allocation for 2.35M compute hours, 2018
  - NERSC production allocation for 3.25M compute hours, 2017
  - NERSC production allocation for 1.9M compute hours, 2016
  - NERSC production allocation for 1.4M compute hours, 2015
  - NSF Teragrid research allocation for 285K compute hours, 2009-2010

#### ADVISING

##### *Current*

- Carl Yang (UC Davis, ECE, Ph.D.), co-advisor, 2015- (w/ John Owens)
- Ben Brock (EECS, Ph.D.), co-advisor, 2017- (w/ Kathy Yelick)
- Giulia Guidi (EECS, Ph.D.), co-advisor, 2018- (w/ Kathy Yelick)

##### *Past*

- Yusuke Nagasaka (Tokyo Institute of Technology), Graduate Summer Intern, 2017 (w/ John Shalf)
- Veronika Strnadova-Neeley (UCSB), GSRA, 2013-2016 (w/ Lenny Oliker, campus advisor: John Gilbert)

- Ariful Azad, Postdoctoral Fellow, 2014-2016
- Patrick Flick (Georgia Tech), Graduate Summer Intern, 2016
- Sang-Yun Oh, Simons Postdoctoral Fellow, 2013-15 (w/ Lenny Olike)
- Harsha Vardhan Simhadri, Postdoctoral Fellow, 2013-14 (w/ Kathy Yelick)
- 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. Olike & S. Williams)

*Ph.D. Committee Membership*

- Fazle Sadi (ECE, CMU), 2017- (advisors: Larry Pileggi and Franz Franchetti)

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.

*HipMer*

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

*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 SpDM<sup>3</sup>*

SpDM<sup>3</sup> 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.



## TEACHING

- **Applications of Parallel Computers (CS267), Spring 2017** (w/ Kathy Yelick).
- **Applications of Parallel Computers (CS267), Spring 2018** (w/ Jim Demmel and Kathy Yelick).

## SERVICES AND LEADERSHIP SKILLS

- *Founding Associate Editor (2013-):* ACM Transactions on Parallel Computing
- *Guest Editor (2015):* Parallel Computing (Elsevier), 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).
- *Program Committee Member:*
  - The 21st Meeting on Algorithm Engineering and Experiments (ALENEX), 2019
  - ACM Symposium on Parallelism in Algorithms and Architectures (SPAA), 2018
  - ACM Symposium on High-Performance Parallel and Distributed Computing (HPDC), 2018
  - ACM/IEEE Intl. Conf. for High Performance Computing, Networking, Storage and Analysis (SC), 2013, 2014, 2016
  - IEEE Intl. Parallel & Dist. Processing Symp. (IPDPS), 2013-2017
  - ACM Conf. on Bioinformatics, Comp. Biology, & Health Informatics (BCB), 2017
  - IEEE Cluster, 2017
  - Workshop on Irregular Applications: Architectures and Algorithms (IA<sup>3</sup>), 2017
  - PLDI ARRAY Workshop on Libraries, Languages & Compilers for Prog., 2017
  - IEEE High Performance Extreme Computing Conference (HPEC), 2016
  - 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
  - Intl. Symposium on Computer Architecture and HPC (SBAC-PAD), 2015
  - International Conference on Parallel Processing (ICPP), 2013
  - 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.
- *Grant Review Panels:* NSF (2015, twice).
- *Minisymposia (Co-)organizer:*
  - “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 Processing 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
- *Reviewer:*
  - *Books:* MIT Press.
  - *Journals:* Communications of the ACM (2), SIAM Review, SIAM Journal of Scientific Computing (5), Bioinformatics, SIAM Journal on Matrix Analysis and Applications, IEEE Transactions on Parallel and Distributed Systems (4), The International Journal of High Performance Computing Applications (2), Parallel Computing (10+), ACM Transactions on Mathematical Software, ACM Transactions on Parallel Computing, 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.
  - *Conferences:* SPAA 2014/2016, IPDPS 2010/2012, Euro-Par 2011/2013, ICPP 2011, SC 2011/2012, ICS 2012.
  - *Grants:* DOE Office of Science (ASCR) 2013-2017
- *Founding President*, Mathematics Club, Sabancı University
- *LBNL Service:*
  - Luis W. Alvarez Postdoctoral Fellowship selection committee, 2014-2016, 2018
- *DOE Service:*
  - *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

INVITED  
TALKS/SEMINARS

*Scaling Parallel Graph Analysis & Machine Learning using Sparse Matrix Operations*

- Michigan State University (Host: H. Metin Aktulga), April 27, 2018
- NERSC Data Seminar, May 25, 2018

*Faster Parallel Graph BLAS Kernels and New Graph Algorithms in Matrix Algebra*

- Google Research (Host: Jack Poulson), Nov 1, 2016
- EECS, University of California, Berkeley, Oct 14, 2016
- *IEEE High Performance Extreme Computing Conference*, Sep 16, 2015
- HP Labs (Host: Robert Schreiber), August 18, 2015

*A tour of contemporary genome assembly algorithms and software*

- Keynote at the workshop on Parallel Software Libraries for Sequence Analysis (pSALSA), October 2, 2016

*Parallel de novo Assembly of Complex Genomes via HipMER*

- *IEEE International Workshop on High Performance Computational Biology (HiCOMB)*, May 23, 2016

*Scalable parallel algorithms for de novo assembly of complex (meta)genomes*

- MANTISSA Day (LBL), August 7, 2015
- Joint Genome Institute, August 4, 2015

*Parallel Graph Algorithms & Primitives*

- Swift Navigation Inc (Host: Timothy Harris), May 26, 2015

*Scalable algorithms for complex genome assembly, alignment, and genetic mapping*

- School of CSE (Host: Srinivas Aluru), Georgia Tech, Jan 16, 2015

*Distributed-Memory Parallel Algorithms for Graph Traversal & Genome Assembly*

- Dept. CS&E (Host: Murat Demirbas), University at Buffalo, Dec 5, 2014
- Dept. CS (Host: Petko Bogdanov), University at Albany, Dec 4, 2014
- Dept. CS (Host: Leman Akoglu), Stony Brook University, Dec 3, 2014

*Three Goals in Parallel Graph Computations: High Performance, High Productivity, and Reduced Communication*

- College of Engineering (Host: Mike Wittie), Montana State University, Sep 13, 2013

*Parallel software for high-performance and high-productivity graph analysis*

- CS Dept. (Host: Semih Salihoglu), Stanford University, Feb 13, 2013

*A sustainable software stack for parallel graph analysis*

- School of CSE (Host: David Bader), Georgia Tech, March 30, 2012

*Scalable Parallel Primitives for Massive Graph Computation*

- EECS (Host: Ahmet Bulut), İstanbul Şehir University, March 21, 2011
- Dept. CS (Host: Alex Pothén), Purdue University, April 20, 2010
- Math. & CS Div. (Host: Mihai Anitescu), Argonne National Laboratory, Jan 25, 2010
- CSRI (Host: Danny Rintoul), Sandia Labs Albuquerque, Jan 20, 2010
- CRD (Host: Juan Meza), Lawrence Berkeley National Laboratory, Jan 12, 2010
- Sandia Labs Livermore (Host: Tammy Kolda), Nov 18, 2009

ORAL  
PRESENTATIONS

1. Communication-avoiding algorithms for large-scale graph and machine learning problems. In *SIAM Annual Meeting*, Portland, OR, 2018. (Minisymposium Talk).
2. Matrices are all you need. In *NECSTLab/LBNL/UCB Workshop*, Berkeley, CA, 2018.
3. Genomics, graphs and the GraphBLAS. In *Graphs Across Domains (GraphXD) Workshop*, Berkeley, CA, 2018. Berkeley Institute for Data Science (BIDS).
4. Graph algorithms, computational motifs, and GraphBLAS. In *Exascale Computing Project (ECP) 2nd Annual Meeting*, Knoxville, TN, 2018. Exagraph Tutorial Session.
5. The reverse Cuthill-McKee algorithm in distributed-memory. In *The IEEE International Symposium on Parallel and Distributed Processing*, Orlando, FL, 2017. (Conference Talk).
6. Parallel algorithms across the GraphBLAS stack. In *ACS HPC and Data Analytics Workshop*, Baltimore, MD, 2017.
7. Developing a PGAS graph library using GraphBLAS primitives. In *SIAM Annual Meeting*, Boston, MA, 2016. (Minisymposium Talk).
8. Faster and more scalable sparse matrix-matrix multiplication. In *SIAM Conference on Parallel Processing for Scientific Computing*, Paris, France, 2016. (Minisymposium Talk).
9. The GraphBLAS effort: Kernels, API, and parallel implementations. In *SIAM Conference on Parallel Processing for Scientific Computing*, Paris, France, 2016. (Minisymposium Talk).

10. 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.
11. The GraphBLAS effort and new parallel algorithms for bipartite graph matching. In *LBNL-Tsukuba Joint Meeting*, Berkeley, CA, 2015.
12. 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.
13. Libraries and algorithms for graph analytics at scale. In *ACS Productivity Workshop*, Baltimore, MD, 2014.
14. The Graph BLAS effort and its implications for Exascale. In *SIAM Workshop on Exascale Applied Mathematics Challenges and Opportunities (EX14)*, Chicago, IL, 2014.
15. Reducing communication in parallel graph computations. In *Workshop on Algorithms for Modern Massive Data Sets (MMDS)*, Berkeley, CA, 2014.
16. Communication-avoiding linear-algebraic primitives for graph analytics. In *Graph Algorithms Building Blocks (GABB'2014), IPDPS Workshops*, Phoenix, AZ, 2014.
17. High-productivity and high- performance analysis of filtered semantic graphs. In *SIAM Conference on Parallel Processing for Scientific Computing*, Portland, OR, 2014. (Minisymposium Talk).
18. 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.
19. High-productivity and high- performance analysis of filtered semantic graphs. In *Seminar on Automatic Application Tuning for HPC Architectures*, Dagstuhl, Germany, 2013.
20. Reducing communication in parallel graph computations. In *DOE ASCR Applied Math PI Meeting*, Albuquerque, NM, 2013.
21. High-performance filtered queries in attributed semantic graphs. In *SIAM Conference on Computational Science and Engineering*, Boston, MA, 2013. (Minisymposium Talk).
22. 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.
23. A sustainable software stack for parallel graph analysis. In *Discovery 2015: HPC and Cloud Computing Workshop*, Berkeley, CA, 2012.
24. Communication-avoiding sparse matrix-matrix multiplication. In *SIAM Conference on Applied Linear Algebra*, Valencia, Spain, 2012. (Minisymposium Talk).
25. Parallel graph libraries: Where do we go from here? In *KDT Spring Mind Meld*, Santa Barbara, CA, 2012.
26. Parallel algorithms for sparse matrix product, indexing, and assignment. In *Scientific Computing and Matrix Computations (LAPACK) Seminar*, UC Berkeley, CA, 2012.
27. 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).
28. 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.

29. 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. (Minisymposium Talk).
30. Parallel sparse matrix indexing and assignment. In *CSC'11: The 5th SIAM Workshop on Combinatorial Scientific Computing*, Darmstadt, Germany. (Contributed Lecture).
31. A parallel framework for identifying vulnerabilities in the electric power grid. In *SIAM Conference on Optimization*, Darmstadt, Germany, 2011. (Minisymposium Talk).
32. Building blocks for scalable graph and data mining software. In *SIAM Conference on Parallel Processing for Scientific Computing*, Seattle, WA, 2010. (Minisymposium Talk).
33. 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).
34. Parallel Combinatorial BLAS and applications in graph computation. In *SIAM Annual Meeting*, Denver, CO, 2009. (Minisymposium Talk).
35. Parallel primitives for computation with large graphs. In *TCPD PhD Forum at the IEEE International Parallel and Distributed Processing Symposium*, Rome, Italy, 2009. (Refereed Doctoral Colloquium Poster Presentation).
36. Sparse matrix-matrix multiplication for accelerating parallel graph computations. In *SIAM Conference on Computational Science and Engineering*, Miami, FL, 2009. (Minisymposium Talk).
37. Challenges and advances in parallel sparse matrix-matrix multiplication. In *International Conference on Parallel Processing*, Portland, OR, 2008. (Conference Talk).
38. Gaussian elimination based algorithms on the GPU. In *PMAA'08: The 5th International Workshop on Parallel Matrix Algorithms and Applications*, Neuchatel, Switzerland. (Contributed Talk).
39. On the representation and multiplication of hypersparse matrices. In *The IEEE International Symposium on Parallel and Distributed Processing*, Miami, FL, 2008. (Conference Talk).
40. Parallel primitives for computation with large graphs. In *SIAM Conference on Parallel Processing for Scientific Computing*, Atlanta, GA, 2008. (Minisymposium Talk).

#### GUEST LECTURING

- 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

- Advanced Operating Systems (Graduate),
- Programming Languages, Formal Languages and Automata (Junior),
- Discrete Mathematics, Programming Methods (Sophomore)

UNDERGRADUATE EMPLOYMENT

- *Student Peer Assistant, Academic Support Program*, Sabancı University
  - Lead Coordinator for the Linear Algebra course.
  - Assistant for Calculus and Science of Nature courses.
- *Undergraduate Researcher*, Sabancı University
  - Human Language and Speech Technologies Laboratory, 2003  
*Building morphological and statistical models for Turkic languages*
  - Computer Vision and Pattern Analysis Laboratory, 2004  
*Optical character recognition using decision trees*

PROFESSIONAL SOCIETY MEMBERSHIPS

- Society of Industrial and Applied Mathematics (SIAM)
  - SIAG on Supercomputing
  - SIAG on Computational Science and Engineering
  - SIAG on Data Mining and Analytics
  - SIAG on Linear Algebra
- Association of Computing Machinery (ACM)
  - SIGHPC: Special Interest Group on High-Performance Computing
- Institute of Electrical and Electronics Engineers (IEEE)

SELECTED OTHER PARTICIPATIONS

- *BOF on HPC Graph Toolkits and GraphBLAS Forum*, Co-organizer, SC, 2017.
- *Partitioned Global Address Space (PGAS) BOF*, Panelist, SC, 2016
- *Panel discussion moderator*, GABB@IPDPS, 2016.
- *BOF (birds of a feather) session leader of GraphBLAS forum*, HPEC, 2014-15.
- *High-Performance Computing on Graphs*, Panelist, SIAM PP, 2014
- *Intel workshop on parallel algorithms for non-numeric computing*, Invited participant, Santa Clara, CA, 2011/2012.
- *Workshop on scalable graph libraries*, Invited participant, Atlanta, GA, 2011
- *Workshop on optimization in energy systems, Institute for computing in science*, Invited participant, Snowbird, UT, 2010.
- *CSCAPES workshop*, 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)