

Conferences and Workshops

Carl Yang, Aydın Buluç, and John D. Owens. **Design Principles for Sparse Matrix Multiplication on the GPU**. In *Proceedings of the 24th International European Conference on Parallel and Distributed Computing*, Euro-Par 2018. Distinguished Paper. 57/197 acceptance.

Carl Yang, Aydın Buluç, and John D. Owens. **Implementing Push-Pull Efficiently in GraphBLAS**. In *Proceedings of the International Conference on Parallel Processing*, ICPP 2018.

Yuechao Pan, Yangzihao Wang, Yuduo Wu, Carl Yang, and John D. Owens. **Multi-GPU Graph Analytics**. In *Proceedings of the 31st IEEE International Parallel and Distributed Processing Symposium*, IPDPS 2017.

Tim 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.

Aydın Buluç, Tim Mattson, Scott McMillan, José Moreira, and Carl Yang. **Design of the GraphBLAS API for C**. In *Proceedings of the 31st IEEE International Parallel and Distributed Processing Symposium Workshops*, IPDPSW 2017.

Jeremy Kepner, Peter Aaltonen, David Bader, Aydın Buluç, Franz Franchetti, John Gilbert, Dylan Hutchison, Manoj Kumar, Andrew Lumsdaine, Henning Meyerhenke, Scott McMillan, Jose Moreira, John D. Owens, Carl Yang, Marcin Zalewski, and Timothy Mattson. **Mathematical Foundations of the GraphBLAS**. In *Proceedings of the IEEE High Performance Extreme Computing Conference*, HPEC 2016.

Leyuan Wang, Yangzihao Wang, Carl Yang, and John D. Owens. **A Comparative Study on Exact Triangle Counting Algorithms on the GPU**. In *Proceedings of the 1st High Performance Graph Processing Workshop*, HPGP 2016.

Yuduo Wu, Yangzihao Wang, Yuechao Pan, Carl Yang, and John D. Owens. **Performance Characterization of High-Level Programming Models for GPU Graph Analytics**. In *IEEE International Symposium on Workload Characterization*, IISWC-2015. Best Paper finalist. 20/61 acceptance.

Carl Yang, Yangzihao Wang, and John D. Owens. **Fast Sparse Matrix and Sparse Vector Multiplication Algorithm on the GPU**. In *Proceedings of the 29th IEEE International Parallel and Distributed Processing Symposium Workshops*, IPDPSW 2015.

Other work

Yangzihao Wang, Yuechao Pan, Andrew Davidson, Yuduo Wu, Carl Yang, Leyuan Wang, Muhammad Osama, Chenshan Yuan, Weitang Liu, Andy T. Riffel, and John D. Owens. **Gunrock: GPU Graph Analytics**. *ACM Transactions on Parallel Computing*, TOPC 2017.

Software artifact for Carl Yang, Aydın Buluç, and John D. Owens. **Design Principles for Sparse Matrix Multiplication on the GPU**. In *Proceedings of the 24th International European Conference on Parallel and Distributed Computing*, Euro-Par 2018: <https://github.com/owensgroup/merge-spmmm>

Software artifact for Carl Yang, Aydın Buluç, and John D. Owens. **Implementing Push-Pull Efficiently in GraphBLAS**. In *Proceedings of the International Conference on Parallel Processing*, ICPP 2018: <https://github.com/owensgroup/push-pull>