

Communication-Avoiding Linear-Algebraic Primitives for Graph Analytics

Aydın Buluç (Lawrence Berkeley National Laboratory)

Joint work with Grey Ballard, James Demmel, John Gilbert, Laura Grigori, Ben Lipshitz, Adam Lugowski, Oded Schwartz, Edgar Solomonik, and Sivan Toledo

Graph algorithms typically have very low computational intensities, hence their execution times are bounded by their communication requirements. In addition to improving the running time drastically, reducing communication will also help improve the energy consumption of graph algorithms. Many of the positive results for communication-avoiding algorithms come from numerical linear algebra. This suggests an immediate path forward for developing communication-avoiding graph algorithms in the language of linear algebra. Unfortunately, the algorithms that achieve communication optimality for asymptotically more available memory are the so-called 3D algorithms, yet the existing software for graph analytics is either 1D or 2D. In this talk, I will describe two new communication-avoiding kernels for graph computations, discuss how they can be integrated into an existing library like the Combinatorial BLAS and how they can be incorporated into the future Graph BLAS standard.

Sparse matrix-matrix multiplication (SpGEMM) enables efficient parallelization of various graph algorithms. It is the workhorse of a scalable distributed-memory implementation of betweenness centrality, an algorithm that finds influential entities in networks. Existing parallel algorithms for SpGEMM spend the majority of their time in inter-node communication on large concurrencies. We investigated communication-optimal algorithms for SpGEMM. Our theoretical paper [1] proves new communication lower bounds, presents two new communication-optimal algorithms, and provides a unified communication analysis of existing and new algorithms. Here, I will also discuss the implications of input/output sparsity on the choice of stationary and replicated matrices [2].

All-pairs shortest paths is a computationally intensive graph algorithm. For dense graphs, we developed a distributed memory algorithm that is communication optimal. This work-efficient algorithm is based on the recursive version of Kleene's algorithm (as opposed to the more popular Floyd-Warshall algorithm), and uses the ideas from 3D algorithms that replicate the input data c times to reduce communication by a factor of \sqrt{c} . The algorithm enables much better strong scaling, and we were able to solve a 65K vertex dense APSP problem in about two minutes [3].

References

- [1] G. Ballard, A. Buluç, J. Demmel, L. Grigori, B. Lipshitz, O. Schwartz, and S. 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.
- [2] M. Schatz, J. Poulson, and R. van de Geijn. Parallel matrix multiplication: 2D and 3D, FLAME Working Note #62. Technical Report TR-12-13, The University of Texas at Austin, Department of Computer Sciences, June 2012.
- [3] E. Solomonik, A. Buluç, and J. Demmel. Minimizing communication in all-pairs shortest paths. In *Proceedings of the IPDPS*. IEEE Computer Society, 2013.