YIBIN ZHAO, University of Toronto
A Simple and Efficient Parallel Laplacian Solver

A symmetric matrix is called a Laplacian if it has nonpositive off-diagonal entries and zero row sums. Laplacian linear systems naturally appears in various optimization problems, e.g. [Daitch and Spielman, 2008, Madry, 2013, Lee and Sidford, 2014, Dong, Gao, Goranci, Lee, Peng, Sachdeva, and Ye, 2021]. Since the seminal work of Spielman and Teng [2004] on solving Laplacian linear systems in nearly linear time, several algorithms have been designed for the task. Yet, the work of Kyng and Sachdeva [2016] remains the simplest and most practical sequential solver. They presented a solver purely based on random sampling and without graph-theoretic constructions such as low-stretch trees and sparsifiers.

In this work, we extend the result of Kyng and Sachdeva [2016] to a simple parallel Laplacian solver with $O(m \log^3 n \log\log n)$ work and $O(\log^2 n \log\log n)$ depth using the ideas of block Cholesky factorization from Kyng, Lee, Peng, Sachdeva, and Spielman [2016]. Our proof is simple and based on standard matrix concentration inequalities. Compared to the best known parallel Laplacian solvers that achieve polylogarithmic depth due to Lee, Peng, and Spielman [2015], our solver achieves both better depth and, for sparse graphs, better work.

This talk is based on a joint work with Sushant Sachdeva.