Matrix Computations and Scientific Computing Seminar

Organizer: J. Demmel and M. Gu

Wednesday, 12:00–1:00 pm, 380 Soda Hall

Nov 30 **Alex Gittens**, UC Berkeley

Terabyte-scale linear algebra: Spark vs MPI, and Spark with MPI

Distributed linear algebra is an indispensable primitive in scientific data analysis; matrix decompositions in particular are ubiquitous. The traditional approach to distributed linear algebra is to write MPI codes in Fortran or C for maximal performance, but it is difficult to integrate these codes into larger analytics pipelines. On the other hand, Spark can easily integrate linear algebraic operations into longer analytics pipelines at the cost of reduced performance. To quantify the gap in performance between MPI and Spark, we compute principal component analyses and non-negative matrix factorizations of multi-terabyte datasets drawn from two real scientific analytics problems, with concurrencies of up to 1500 nodes on a Cray XC40 supercomputer. The results point to fundamental limits in using Spark alone to perform distributed linear algebra. To help bridge this performance/usability gap, we introduce Alchemist, an in-memory framework for interfacing Spark with arbitrary MPI codes that take matrices as input. We report preliminary results on the performance of Alchemist.