Opportunities for Mixed Precision in Preconditioned Iterative Methods

Erin Carson[1]

Mixed precision hardware is now commercially available and is expected to be a crucial aspect of supercomputers going forward. Accompanying this trend is a renewed interest in developing mixed precision algorithms for numerical linear algebra. The goal is to determine how to selectively use low precision in parts of the computation and higher precision in others, in a way that provides both performance improvements and acceptable guarantees on attainable accuracy. One popular approach is mixed precision GMRES-based iterative refinement schemes for solving linear systems $Ax = b$. Such algorithms form the basis for the new mixed precision HPL-AI benchmark, which already exceeds exaflop performance on the world's top supercomputers.

Most existing analyses of iterative refinement schemes based on Krylov subspace methods assume that the preconditioner, typically an LU factorization, is constructed such that it is an exact inverse of A in exact arithmetic. In practical applications involving large sparse matrices A, one typically uses an approximate preconditioner such as an incomplete LU factorization or a sparse approximate inverse. In this talk, we discuss recent work on analyzing the quality of approximate preconditioners constructed in low precision and their use within iterative refinement schemes. We also give a general overview of the landscape of research in mixed precision preconditioned iterative methods for solving linear systems and discuss open problems.

[1]Charles University
carson@karlin.mff.cuni.cz