Mixed precision iterative refinement for the solution of large sparse linear systems

Bastien VIEUBLE

University of Manchester, United Kingdom

Abstract

The increasing availability of low precision (e.g., fp16, bfloat16, fp8) in hardware tends to lead modern high performance computation to use mixed precision standards. In that context, iterative refinement (IR), focused initially on improving linear system solution accuracy, has embraced approximate computing. In this presentation, we will cover the recent works related to IR that led to the finalization of the Ph.D. thesis entitled "Mixed precision iterative refinement for the solution of large sparse linear systems" and which has been defended in November 2022. It includes principally theoretical analyses and practical implementations of state-of-the-art mixed precision iterative refinement algorithms combined with the MUMPS direct solver for the solution of parallel sparse problems. We used single precision computation, block low-rank, and static pivoting to improve the performance of the LU direct solver, while keeping double precision accuracy on the solution. We demonstrated that these algorithms could achieve significant gains in both memory consumption and execution time on parallel machines while providing accurate solutions for numerically difficult problems. In this presentation, we will also briefly tackle the use of IR for the improvement of preconditioned iterative methods, such as GMRES, and propose an extended error analysis that covers algorithms of up to six precisions.