Reduced and Mixed Precision QR Decomposition Techniques

Eda Oktay

Max Planck Institute for Dynamics of Complex Technical Systems,
Magdeburg, Germany
OKTAY@MPI-MAGDEBURG.MPG.DE

Abstract

Modern high-performance computing (HPC) architectures increasingly rely on reduced-precision arithmetic to improve performance, reduce energy consumption, and more efficiently use limited memory bandwidth. These developments are particularly relevant for large-scale PDE solvers, where many algorithmic components ultimately rely on linear algebra operations. In particular, QR factorizations appear in least-squares formulations of PDE-constrained problems, in domain-decomposition methods that require orthogonalization, and in Krylov subspace methods where orthogonality and stability are essential. However, using uniformly low precision in these settings can introduce rounding errors that accumulate during orthogonalization, leading to numerical instability, loss of orthogonality, and slower or unreliable convergence of the overall PDE solver.

In this talk, we will cover several QR decomposition techniques, including Householder and (shifted) Cholesky factorizations, and study their behavior under reduced- and mixed-precision arithmetic. We show which parts of these algorithms can safely be executed in low precision and which require higher precision to maintain stability. We also discuss how combining multiple precisions within a single algorithm can balance efficiency and numerical reliability. Performance results on modern HPC architectures, together with numerical experiments, demonstrate the achievable speedups and clarify how precision choices influence stability, accuracy, and convergence in practice.