Towards an Optimal Parallel Approximate Sparse Factorization Algorithm Using Hierarchically Semi-separable Structures

Funded Project: CACHE - Math/CS Institute (Algorithms and Software for Communication Avoidance and Communication Hiding at the Extreme Scale)

X. Sherry Li
Lawrence Berkeley National Laboratory
1 Cyclotron Road, Berkeley, CA 94720.

Abstract

Many DOE simulation codes involve discretization of PDEs that lead to highly indefinite and ill-conditioned linear systems, such as the extended MHD equations in plasma fusion and the Maxwell equations in accelerator cavity modeling. These codes heavily rely on sparse factorization based algorithms either as direct solvers or as preconditioners in the iterative solvers. Traditional sparse factorization algorithms are not scalable in operation count nor in memory. We are developing a class of novel superfast and memory-efficient sparse factorization methods, which can be used as general-purpose, robust direct solvers as well as preconditioners. The novelty is to exploit low-rank structures arising in the intermediate submatrices throughout factorization. In particular, we use the *Hierarchically Semi-separable (HSS) matrices* in our representation to enable low-rank approximations [1]. The new factorization algorithms possess nearly linear complexity in time, memory and communication for solving both 2D and 3D elliptic PDE problems. The similar complexity holds when they are used as approximate factorization preconditioners for a wider class of problems. More importantly, much less data is needed to communicate than conventional algorithms; this is particularly superior in the extreme scale computing.

The HSS matrix structure is very useful in the development of superfast direct solvers for both dense and sparse linear systems. For example, it serves as a kernel embedded in a massively parallel structured Helmholtz solver for large-scale frequency domain wave equation modeling prevailing in the oil and gas industry. Recently, we designed and implemented highly efficient parallel algorithms for the key HSS operations that are required in the linear solvers. These include parallel HSS constructions, parallel rank-revealing QR factorizations, parallel HSS ULV factorizations, and parallel HSS solutions. We also designed a specialized HSS-embedded multifronal solver for the standard stencils on regular meshes and obtained significant speedup and memory saving over the classical multifrontal solver when applied to the anisotropic Helmholtz equations for seismic imaging. We were able to solve a linear system with 6.4 billion unknowns using 4096 processors, in less than 20 minutes [2]. The classical multifrontal factorization simply failed due to high demand of memory.

In this talk we will analyze the algorithm scalability and parallel efficiency of the new parallel HSS operations and the HSS-embedded sparse solver, and describe how it can be used as a preconditioner in a communication-avoiding Krylov solver.

References

[1] M. Bebendorf. *Hierarchical Matrices*, volume 63 of *Lecture Notes in Computational Science and Engineering*. Springer, Berlin Heidelberg, 2008.

[2] Shen Wang, Xiaoye Li, Jianlin Xia, Yingchong Situ, and Maarten V. de Hoop. Efficient parallel algorithms for hierarchically semiseparable matrices. *SIAM J. Scientific Computing*, 2011, submitted.