Towards an optimal-order approximate sparse factorization exploiting data-sparseness in separators

Xiaoye S. Li

Lawrence Berkeley National Laboratory

Nested dissection ordering and its graph partitioning generalization give rise to an ordered sequence of separators of (roughly) geometrically decreasing size. The fill-ins in the LU factors are confined in the parts of the matrix associated with the separators. In particular, the diagonal blocks associated with the separators are fully dense in the factors and they contribute to the dominant terms in the costs of the storage for the factors and the flops for computing the factors. Employing the data-sparse representations or compressions, such as low-rank approximation, for these separator blocks can drastically lower the overall factorization costs both in memory and flops. The low rankness can appear in many matrices from discretized PDEs.

Recently, we have been investigating fast and stable algorithms for one type of data-sparse representation called hierarchically semi-separable (HSS) structure, and using them in the sparse LU factorization. In this talk, we will show that both in theory and in practice, the HSS-embedded saprse LU factorization has much lower complexity than the traditional factorization algorithm. The complexity of this class of algorithms is closely related to the numerical ranks of the separator blocks which vary with the PDEs. For elliptic problems, we can achieve large amount of compression and the resulting factorization can be used as a nearly optimal-order direct solver. For wider classes of problems, the approximate factorization can be used as a preconditioner. We will show performance results from direct solvers as well as preconditioners for a wide range of problems. We will also illustrate the potential of such solvers/preconditioners being used for the extreme-scale computers and problem size.

This is joint work with A. Napov, F.-H. Rouet, S. Wang, and J. Xia.