Compressed threshold pivoting for sparse symmetric indefinite systems

Jonathan Hogg and Jennifer Scott
Scientific Computing Department
STFC Rutherford Appleton Laboratory
Harwell Oxford, OX11 0QX, England.

Abstract

We are interested in the efficient and stable factorization of large sparse symmetric indefinite matrices of full rank. In this talk, we propose two new pivoting strategies that are designed to significantly reduce the amount of communication required when selecting pivots.

Most algorithms for the factorization of large sparse symmetric indefinite matrices employ supernodes, that is, a set of consecutive columns having the same (or similar) sparsity pattern in the factor. By storing only those rows that contain nonzeros, each supernode may be held as a dense $n \times p$ trapezoidal matrix. At each stage, a search is made for a pivot from the $p \times p$ leading block. If a candidate pivot is found to be unsuitable, its elimination is delayed to a later supernode, with a guarantee that all pivots will be eliminated in the final supernode. Such delayed pivots generate additional floating-point operations and storage requirements. Good scalings and orderings can reduce the number of delayed pivots but not remove the need for testing pivots for numerical stability.

With the advent of manycore processors and the growing gap between the speed of communication and computation, many algorithms need to be rewritten to reflect the changing balance in resource. As pivoting decisions must be taken in a serial fashion, they are highly sensitive to the latency and speed of any communication or bandwidth costs incurred. With current algorithms that take into account the entire candidate pivot column below the diagonal, all threads working on a supernode must endure stop-start parallelism for every column of the supernode.

We seek to develop effective pivoting strategies that significantly reduce the amount of communication required by compressing the necessary data into a much smaller matrix that is then used to select pivots. A provably stable algorithm and a heuristic algorithm are presented; we refer to these algorithms as *compressed threshold pivoting* algorithms. The heuristic algorithm is faster than the provably stable alternative and it more accurately approximates the behaviour of traditional threshold partial pivoting in terms of modifications to the pivot sequence. While it can demonstrably fail to control the growth factor for some pathological examples, in practice, provided it is combined with appropriate scaling and ordering, it achieves numerical robustness even on the most difficult practical problems. Further details are given in [1].

This work was funded by EPSRC grant EP/I013067/1.

[1] J. D. Hogg and J. A. Scott. Compressed threshold pivoting for sparse symmetric indefinite systems. Technical Report RAL-TR-2013-P-007, Rutherford Appleton Laboratory, 2013.