Enhancing Incomplete Cholesky Decompositions

Jiří Kopal, Jennifer Scott, Miroslav Tůma, and Miroslav Rozložník

Abstract

Incomplete Cholesky decompositions represent an important component in the solution of large sparse symmetric positive-definite systems of equations. Such decompositions arise in a wide range of practical applications. Preconditioners based on the decomposition combined with a Krylov-space accelerator are routinely used in a number of production codes.

Many variants of the incomplete Cholesky decomposition and of its refinements have been proposed and used to solve practical problems. The enhancements of the basic procedure have varied from new mathematical ideas and algorithmic simplifications to more sophisticated implementations. If we try to systematize the decades of development, we could classify them as offering improvements either in the accuracy of the LL^T decomposition measured by a norm of the distance from the system matrix A or in the stability of the computed factors.

Neither of these goals can be separated from the way in which the Cholesky decomposition is implemented. Consider, for simplicity, sequential implementations; we could state that for small $k \geq 1$, a useful and robust decomposition can afford to spend k-times more time in the factorization than the simplest no-fill procedure. With this flexibility in mind, a space-efficient incomplete Cholesky factorization based on the concept of intermediate memory introduced by Tismenetsky was recently discussed in [2] (with the resulting software made available to the community [3]). As we observed from extensive numerical experiments, the use of intermediate memory can improve preconditioner accuracy. While a fixed bound on the memory used during the decomposition seems to be a practical must, another important challenge is to decide on the distribution of the entries in the factor columns. There have been a lot of important contributions on this that include, for example, exploiting structural information provided by the symbolic factorization of a complete factorization or using dropping rules based on estimated or computed inverses of the factors. In this talk, we introduce a new way to split the memory between that required for the factor and the intermediate memory used only in its construction. It is based on estimating norms of Schur complement updates and we consider it as a natural extension of the ideas presented in [2].

We also propose a new dropping approach. It is theoretically motivated and it extends the concept of a posteriori filtering from the approximate inverse decomposition in [1] to the incomplete Cholesky decomposition. The new approach drops entries dynamically such that the significance of factor entries in different columns is balanced. We believe that both our proposals may provide steps on the way to achieving more robust incomplete Cholesky decompositions.

This work was partially supported by the projects P201/13-06684S and 108/11/0853 of the Grant Agency of the Czech Republic and by EPSRC grant EP/I013067/1.

References

[1] M. Rozložník J. Kopal and M. Tůma. Approximate inverse preconditioners with adaptive dropping. *submitted to International Journal of Advances in Engineering Software*, 2013.

[2] J. A. Scott and M. Tůma. On positive semidefinite modification schemes for incomplete Cholesky factorization. Technical Report RAL-TR-2013-P-005, 2013.

[3] J. A. Scott and M. Tůma. `HSL_MI28`: an efficient and robust limited memory incomplete Cholesky factorization code. Technical Report RAL-TR-2013-P-004, 2013, to appear in ACM Trans. Math. Software.