Algorithm 740: Fortran Subroutines to Compute Improved Incomplete Cholesky Factorizations

MARK T. JONES and PAUL E. PLASSMANN
Argonne National Laboratory

Efficient and reliable code to compute incomplete Cholesky factors of sparse matrices for use as preconditioners in a conjugate gradient algorithm is described. This code implements two recently developed, improved incomplete factorization algorithms An efficient implementation of the standard incomplete Cholesky factorization is also included.

Categories and Subject Descriptors: G.1.3 [**Numerical Analysis**]: Numerical Linear Algebra—*linear systems; sparse and very large systems*

General Terms: Algorithms

Additional Key Words and Phrases: Incomplete Cholesky, incomplete factorization, preconditioners, sparse matrices

1. INTRODUCTION

The Fortran functions JPICC and JPICR are efficient implementations of the column-based and row-based incomplete Cholesky factorizations introduced by Jones and Plassmann [1991]. For comparison purposes, we include the Fortran function ISTDIC, an implementation of the usual incomplete Cholesky factorization [Meijerank and Van Der Vorst 1977]. The subroutines generate incomplete factorizations of symmetric, positive definite sparse matrices for use as preconditioners for a conjugate gradient algorithm. A preconditioned conjugate gradient subroutine is not included; however, the preconditioners generated can easily be incorporated into an existing routine, such as the DCG routine in SLAP [Seager 1988], or a routine may be written by the user.

2. ROUTINES

The software conforms to the 1977 ANSI standards for Fortran and has no system dependencies. Comments in each user-callable module describe the calling sequences. All parameter names conform to the Fortran typing default. Only a double-precision implementation is included. The modules are as follows:

ISTDIC: An integer function that returns the standard column-oriented incomplete Cholesky factorization.

JPICC: An integer function that returns the improved column-oriented incomplete Cholesky factorization.

JPICR: An integer function that returns the improved row-oriented incomplete Cholesky factorization.

IBSORT: A subroutine that returns the K largest nonzeros from a double-precision array that is indirectly addressed by an integer array.

DBSORT: A subroutine which performs a bubble sort on an integer array.

IHSORT: A subroutine that performs heap sort on a double-precision array that is indirectly addressed by an integer array.

DHSORT: A subroutine that performs heap sort on an integer array.

For testing and demonstration purposes, a small sample driver is included with the subroutines. The expected output for the driver is also included. The user may find it beneficial to replace the sorting subroutines with ones that have been optimized for a particular architecture.

REFERENCES

JONES, M. T. AND PLASSMANN, P. E. 1991. An improved incomplete Cholesky factorization. Preprint MCS-P206-0191, MCS Division, Argonne National Laboratory, Argonne, Ill.

MEIJERINK, J. AND VAN DER VORST, H. A. 1977. An iterative solution method for linear systems of which the coefficient matrix is a symmetric M-Matrix. *Math. Comput. 31*, 137 (Jan.), 148–162.

SEAGER, M. 1988. A SLAP for the masses. Tech. Rep. UCRL-100267, Lawrence Livermore National Laboratory.