A Numerical Evaluation of Sparse Direct Solvers for the Solution of Large Sparse Symmetric Linear Systems of Equations

NICHOLAS I. M. GOULD and JENNIFER A. SCOTT
Rutherford Appleton Laboratory
and
YIFAN HU
Wolfram Research

In recent years a number of solvers for the direct solution of large sparse symmetric linear systems of equations have been developed. These include solvers that are designed for the solution of positive definite systems as well as those that are principally intended for solving indefinite problems. In this study, we use performance profiles as a tool for evaluating and comparing the performance of serial sparse direct solvers on an extensive set of symmetric test problems taken from a range of practical applications.

Categories and Subject Descriptors: G.1.0 [**Numerical Analysis**]: General—*Numerical algorithms*; G.1.3 [**Numerical Analysis**]: Numerical Linear Algebra—*Sparse, structured, and very large systems (direct and iterative methods)*

General Terms: Algorithms, Performance

Additional Key Words and Phrases: Sparse matrices, symmetric linear systems, Gaussian elimination, direct solvers, software

ACM Reference Format:

Gould, N. I. M., Scott, J. A., and Hu, Y. 2007. A numerical evaluation of sparse direct solvers for the solution of large sparse symmetric linear systems of equations. ACM Trans. Math. Softw., 33, 2, Article 10 (June 2007), 32 pages. DOI = 10.1145/1236463.1236465 http://doi.acm.org/10.1145/1236463.1236465

1. INTRODUCTION

Solving linear systems of equations lies at the heart of many problems in computational science and engineering. In many cases, particularly when discretizing

This work was partly funded by the EPSRC Grants GR/R46641 and GR/S42170.
Authors' addresses: N. I. M. Gould and J. A. Scott, Computational Science and Engineering Department, Atlas Centre, Rutherford Appleton Laboratory, Oxon OX11 0QX, England; email: j.a.scott@rl.ac.uk; Y. Hu, Wolfram Research, Inc., 100 Trade Center Drive, Champaign, IL 61820.

continuous problems, the system is large and the associated matrix A is sparse. Furthermore, for many applications, the matrix is symmetric; sometimes, such as in some finite-element applications, A is positive definite, while in other cases, including constrained optimization and problems involving conservation laws, it is indefinite.

A direct method for solving a sparse linear system $Ax = b$ involves the explicit factorization of the system matrix A (or more usually, a permutation of A) into the product of lower and upper triangular matrices L and U. In the symmetric case, for positive definite problems $U = L^T$ (Cholesky factorization) or, more generally, $U = DL^T$, where D is a block diagonal matrix with 1×1 and 2×2 blocks. Forward elimination followed by backward substitution completes the solution process for each given right-hand side b. Direct methods are important because of their generality and robustness. Indeed, for the 'tough' linear systems arising from some applications, they are currently the only feasible solution methods. In many other cases, direct methods are the method of choice because finding and computing a good preconditioner for an iterative method can be computationally more expensive than using a direct method. Furthermore, direct methods provide an effective means of solving multiple systems with the same A but different right-hand sides b because the factorization only needs to be performed once.

Since the early 1990s, many new algorithms and a number of new software packages that are designed for the efficient solution of sparse symmetric systems, have been developed. Because a potential user may be bewildered by such choice, our intention in this article is to compare the alternatives on a significant set of large test examples from many different application areas, and as far as is possible, to make recommendations concerning the efficacy of the various algorithms and packages. This study is an extension of a recent comparison by Gould and Scott [2004] of sparse symmetric direct solvers in the mathematical software library HSL [2004]. This earlier study concluded that the best general-purpose HSL package for solving sparse symmetric systems is currently `MA57` [Duff 2004]. Thus the only HSL direct solver included here is `MA57`, but the reader should be aware that, for some classes of problems, other HSL codes may be more appropriate. For full details and results for the HSL symmetric solvers, the reader is referred to Gould and Scott [2003].

For ease of reference, all the sparse solvers used in this study are listed in Table I. The release date of the version of the code used in our experiments is given. The codes will be discussed in more detail in Sections 2 and 3. We remark that a number of the packages offer versions for complex symmetric and/or Hermitian matrices, and some can be used for unsymmetric systems. Our experiments are limited to real symmetric matrices. Some of the packages have parallel versions (and may even have been written primarily as parallel codes); this study considers only serial codes and serial versions of parallel solvers. In addition, we have excluded solvers that are integrated parts of more general application software (such as optimization packages) but appear not to be callable on their own.

Table I. Solvers Used in our Numerical Experiments. An '&' Indicates Both Languages are Used in the Source Code; 'F77/F90' Indicates There is a F77 Version and a F90 Version

Code	Date/Version	Language	Authors
BCSLIB-EXT	11.2001, v4.1	F77	The Boeing Company
CHOLMOD	08.2005, v0.4	C	T. Davis
MA57	06.2005, v3.0.1	F77/F90	I.S. Duff, HSL
MUMPS	10.2005, v4.5.5	F90	P.R. Amestoy, I.S. Duff, J.-Y. L'Excellent, J. Koster, A. Guermouche and S. Pralet
Oblio	12.2003, v0.7	C++	F. Dobrian and A. Pothen
PARDISO	04.2005, v1.2.3	F77 & C	O. Schenk and K. Gärtner
SPOOLES	1999, v2.2	C	C. Ashcraft and R. Grimes
SPRSBLKLLT	1997, v0.5	F77	E.G. Ng and B.W. Peyton
TAUCS	08.2003, v2.2	C	S. Toledo
UMFPACK	04.2003, v4.1	C	T. Davis
WSMP	02.2006, v6.3.15	F90 & C	A. Gupta, IBM

Table II. Availability and Contact Details of The Solvers Used in our Numerical Experiments

Code	Free to Academics	Web Page / Email Contact
BCSLIB-EXT	×	www.boeing.com/phantom/bcslib-ext
CHOLMOD	√	www.cise.ufl.edu/∼davis/
MA57	×	www.cse.clrc.ac.uk/nag/hsl
MUMPS	√	www.enseeiht.fr/lima/apo/MUMPS/
Oblio	√	dobrian@cs.odu.edu or pothen@cs.odu.edu
PARDISO	√	www.computational.unibas.ch/cs/scicomp/software/pardiso
SPOOLES	√	www.netlib.org/linalg/spooles/spooles.2.2.html
SPRSBLKLLT	√	EGNg@lbl.gov
TAUCS	√	www.cs.tau.ac.il/∼stoledo/taucs/
UMFPACK	√	www.cise.ufl.edu/research/sparse/umfpack/
WSMP	√	www-users.cs.umn.edu/∼agupta/wsmp.html

Some of the solvers are freely available to academics while it is necessary to purchase a licence to use others. This information is provided in Table II. For each code a Web page address is also given (or, if no Web page is currently available, an email contact is provided that may be used to obtain further information). Note that for nonacademic users, some of the solvers are available without cost but the conditions for obtaining and using a solver vary among the different packages, so we advise interested users to refer to the Web page or contact the code's authors directly for full details. We remark that while the use of the HSL code MA57 requires that a licence be purchased (with preferential rates for academic users), its forerunner MA27 is freely available for all noncommercial use; details may be found at hsl.rl.ac.uk/archive/hslarchive.html.

We carried out our study of sparse solvers by writing, for each package, a separate driver that reads our test data (see Section 4.1), runs the different phases of the solver (see Section 2), and stores the output statistics (see Section 4.2) in a series of files. The drivers are written in the language of the solver (Table I). Once all the runs were complete, we read the data from the output files and used it to plot the performance profiles (see Section 4.2). Our findings are presented in Section 5.

2. AN INTRODUCTION TO SPARSE SYMMETRIC SOLVERS

Sparse direct methods solve systems of linear equations by factorizing the coefficient matrix A, generally employing graph models to try to minimize both the storage needed and work performed. Sparse direct solvers have a number of distinct phases. Although the exact subdivision depends on the algorithm and software being used, a common subdivision is given by:

(1) An ordering phase that exploits structure.
(2) An analyze phase (which is sometimes referred to as the symbolic factorization step) that analyzes the matrix structure to optionally determine a pivot sequence and data structures for efficient factorization. A good pivot sequence significantly reduces both memory requirements and the number of floating-point operations required.
(3) A factorization phase that uses the pivot sequence to factorize the matrix (some codes scale the matrix prior to the factorization).
(4) A solve phase that performs forward elimination followed by back substitution using the stored factors. The solve phase may include iterative refinement.

In a serial implementation, the factorization is usually the most time-consuming phase, while the solve phase is generally significantly faster. In many software packages, the first two phases are combined into a single user-callable subprogram. An introduction to sparse direct solvers is given in Duff et al. [1986]. Another useful reference for symmetric positive definite systems is George and Liu [1981].

2.1 Ordering Choices

There are a number of different approaches to the problem of obtaining a good pivot sequence. An important class of ordering methods is based upon the minimum degree (MD) algorithm, first proposed by Tinney and Walker [1967]. Variants include the multiple minimum degree (MMD) algorithm [Liu 1985] and the approximate minimum degree (AMD) algorithm [Amestoy et al. 1996, 2004]. QAMD is a version of AMD that is designed to avoid the problems caused by quasi-dense rows (details are available at `www.netlib.org/linalg/amd/`), while COLAMD is a column approximate minimum degree algorithm [Davis et al. 2004a, 2004b]. Other methods are based on nested dissection (ND), a term introduced by George [1973]. Many of the recent packages include an explicit interface to the multilevel nested-dissection routine `METIS_NodeND` (or a variant of it) from the METIS package [Karypis and Kumar 1998, 1999]. Other orderings include multisection [Ashcraft and Liu 1998] and orderings based on local minimum fill [Tinney and Walker 1967]. The ordering options offered by the codes in this study are summarized in Table III. An entry marked with * indicates the default (or recommended) ordering. Note that for `MUMPS`, the default is dependent on the size of the linear system while the packages `SPOOLES` and `WSMP` perform two orderings by default and select the better. By default, version 3.0.1 of `MA57` automatically chooses whether to use QAMD or METIS depending

Table III. Ordering Options and Factorization Algorithm for the Solvers Used in our Numerical Experiments. MD = Minimum Degree; AMD = Approximate Minimum Degree, MMD = Multiple Minimum Degree; ND = Nested Dissection; METIS = Explicit Interface to METIS_NodeND (or Variant of it); MS = Multisection; MF = Minimum Fill. $\surd\surd$ Indicates AMD and One or More Variants of AMD are Available (Further Details are Given in Section 3).* Indicates the Default

	Ordering Options							Factorization
Code	MD	AMD	MMD	ND	METIS	MS	MF	Algorithm
BCSLIB-EXT	×	×	$\surd$	×	$\surd^*$	×	×	Multifrontal
CHOLMOD	×	$\surd\surd^*$	×	×	$\surd^*$	×	×	Left-looking
MA57	$\surd$	$\surd\surd^*$	×	×	$\surd^*$	×	×	Multifrontal
MUMPS	$\surd$	$\surd\surd^*$	×	×	$\surd^*$	$\surd$	$\surd$	Multifrontal
Oblio	×	×	$\surd$	×	$\surd^*$	×	×	Left-looking, right-looking, multifrontal
PARDISO	×	×	$\surd$	×	$\surd^*$	×	×	Left-right looking
SPOOLES	×	×	$\surd$	$\surd^*$	×	$\surd^*$	×	Left-looking
SPRSBLKLLT	×	×	$\surd^*$	×	×	×	×	Left-looking
TAUCS	$\surd$	$\surd$	$\surd$	×	$\surd^*$	×	×	Left-looking, multifrontal
UMFPACK	×	$\surd^*$	×	×	×	×	×	Unsymmetric multifrontal
WSMP	×	×	×	$\surd^*$	×	×	$\surd^*$	Multifrontal

on the order of the system and the characteristics of the sparsity pattern; in some cases it will perform both orderings and use the one with the smaller predicted level of fill (see Duff and Scott [2005] for details). CHOLMOD tries an AMD variant first and then, depending on the quality of this ordering, automatically decides whether or not to try METIS. We note that all the solvers also allow the user to supply his or her own ordering. For the package SPRSBLKLLT, this can only be done if the matrix is preordered before entry; the other packages perform any necessary permutations on the input matrix using the supplied ordering.

2.2 Factorization Algorithms

Following the analyze phase, the factorization can be performed in many different ways, depending on the order in which matrix entries are accessed and/or updated. Possible variants include left-looking, right-looking, and multifrontal algorithms. The (supernodal) right-looking variant computes a (block) row and column at each step and uses them to immediately update all rows and columns in the part of the matrix that has not yet been factored. In the supernodal left-looking variant, the updates are not applied immediately; instead, before a block column k is eliminated, all updates from previous columns of L are applied together to the block column k of A. Hybrid left-right looking algorithms have also been proposed [Schenk et al. 2000]. The multifrontal method was first introduced by Duff and Reid [1983]. It accumulates the updates and they are propagated from a descendant column j to an ancestor column k via all intermediate nodes on the elimination tree path from j to k. Further details of these

variants may be found, for example, in the survey article of Heath et al. [1991] and in Dongarra et al. [1998]. A useful overview of the multifrontal method is given by Liu [1992]. The algorithm used by each of the codes involved in our tests is given in Table II. Note that a number of the solvers (in particular, Oblio and TAUCS) offer more than one algorithm.

2.3 Pivoting for Stability

For symmetric matrices that are positive definite, the pivot sequence may be chosen using the sparsity pattern alone, and so the analyze phase involves no computation on real numbers and the factorization phase can use the chosen sequence without modification. Moreover, the data structures are determined by the analyze phase and can be static throughout the factorization phase. For symmetric indefinite problems, using the pivot sequence from the analyze phase may be unstable or impossible because of (near) zero diagonal pivots. The disadvantage of using standard partial pivoting for stability is that symmetry is destroyed. Different codes try to address this problem in different ways. The simplest approach is to terminate the computation if a zero (or very small) pivot is encountered. Alternatively, the computation may be continued by perturbing near-zero pivots. This allows the data structures chosen by the analyze phase to be used, but may lead to large growth in the entries of the factors. The hope is that accuracy can be restored through the use of iterative refinement but, with no numerical pivoting, these simple static approaches are only suitable for a restricted set of indefinite problems.

A larger set of problems may be solved by selecting only numerically stable 1×1 pivots from the diagonal, that is, a pivot on the diagonal is only chosen if its magnitude is at least u times the largest entry in absolute value in its column, where $0 < u \leq 1$ is a threshold parameter set by the user. Potentially unstable pivots (those that do not satisfy the threshold test) will be delayed, and the data structures chosen during the analyze phase may have to be modified.

To preserve symmetry and maintain stability, pivots may be generalized to 2×2 blocks. Again, different packages use different 2×2 pivoting strategies. The approach of PARDISO is to use Bunch-Kaufmann pivoting [Bunch and Kaufmann 1977] on the dense diagonal blocks that correspond to supernodes, and if a zero or nearly zero pivot occurs, it is perturbed. Since pivots are only chosen from within the supernodal block, numerical stability is not guaranteed but because there is no searching or dynamic reordering during the factorization, it is anticipated that this static pivoting strategy will have a substantial performance advantage over more robust approaches that search for a stable pivot and force the delay of any that are unstable. The stable approach is followed by MA57, which uses a modified version of the algorithm of Bunch et al. [1976]; details are given in Duff [2004]. A threshold parameter $u \in (0, 0.5]$ must be selected. Values close to zero will generally result in a faster factorization with fewer entries in the factors, but values close to 0.5 are likely to result in a more stable factorization. MUMPS and Oblio follow a similar approach to MA57. BCSLIB-EXT and WSMP also use 1×1 and 2×2 block pivots, again with a threshold parameter under the user's control. BCSLIB-EXT gives preference to 2×2 pivots;

the algorithm is described in Ashcraft et al. [1998]. This article also proposed using $k \times k$ block pivots to improve performance but none of the solvers in our study currently employs pivot blocks with $k > 2$. For the software developer, the main disadvantage of including full 2×2 pivoting is that it adds significantly to the complexity of the code, particularly in a parallel implementation.

In recent years, significant effort has been put into improving the choice of pivot sequence made for indefinite problems during the analyze phase so that it can be used with a minimum number of modifications during the factorization phase. Most codes select a tentative pivot sequence based upon the sparsity pattern alone. Then during the factorization they either return an error message if the sequence cannot be used, or modify it to allow the factorization to continue. By default, the versions of MUMPS and WSMP that we tested in this study use the numerical values when selecting the pivot sequence [Duff and Pralet 2004], and PARDISO also has an option to use the numerical values. In both cases, the pivot sequence is again passed from the analyze phase to the factorization phase, and it may still need modifying during the factorization to retain stability. An alternative approach is to combine the analyze and factorization phases. In such cases, the code is sometimes described as an analyze-factorize code. The software library HSL includes the analyze-factorize code MA67, which is primarily designed for the solution of symmetric indefinite problems. The results of our earlier study [Gould and Scott 2004] found that overall MA67 was slower than Version 1.0.0 of the multifrontal code MA57, but MA67 was successful in efficiently solving some "tough" (highly ill-conditioned and singular) indefinite problems that MA57 struggled with. However, it is common to encounter the need to factorize and solve a sequence of sparse linear systems where the coefficient matrices change, but their sparsity pattern remains fixed. A key advantage of designing a solver with separate analyze and factorization phases is that the work of choosing a pivot sequence may not have to be repeated (of course, if the analyze phase uses the numerical values, the analyze may have to be repeated if the numerical values change significantly).

The pivoting strategies offered by the codes used in this study are summarized in Table IV. Further details are given in Section 3. Although each of the codes may be used to solve positive definite problems, some have an option that allows the user to indicate that the matrix is positive definite and, in this case, the code follows a logically simpler path. A '$\surd$' in the column headed 'Positive Definite' indicates that the code either has such an option or is designed principally for positive definite systems. A '$\times$' in the 'Indefinite' column indicates that the documentation available with the code states it is designed for solving positive definite problems and is thus not intended for indefinite examples. In our numerical experiments, the latter codes will only be used to solve the positive definite problems.

2.4 Out-of-Core Factorization

To solve very large problems using a direct solver it is usually necessary to work out-of-core. By holding the matrix and/or its factor in files, the amount of main memory required by the solver can be substantially reduced. In this study,

Table IV. Default Pivoting Strategies Offered by the Solvers Used in our Numerical Experiments. √ Indicates the Positive Definite Case is Treated Separately. × Denotes Package Not Intended for Indefinite Problems. $\times^*$ Denotes an Indefinite Version is under Development.

Code	Positive Definite	Indefinite
BCSLIB-EXT	√	Numerical pivoting with 1×1 and 2×2 pivots.
CHOLMOD	√	×
MA57	√	Numerical pivoting with 1×1 and 2×2 pivots.
MUMPS	√	Numerical pivoting with 1×1 pivots.
Oblio	√	Numerical pivoting with 1×1 and 2×2 pivots.
PARDISO	√	Supernode Bunch-Kaufmann within diagonal blocks.
SPOOLES	√	Fast Bunch-Parlett.
SPRSBLKLLT	√	×
TAUCS	√	$\times^*$
UMFPACK		Partial pivoting with preference for diagonal pivots.
WSMP	√	Numerical pivoting with 1×1 and 2×2 pivots.

only the solvers BCSLIB-EXT, Oblio, and TAUCS include an option for holding the matrix factor out-of-core. Oblio also allows the stack used in the multifrontal algorithm to be held in a file. BCSLIB-EXT is the most flexible. It offers the option of holding the matrix data and/or the stack in direct access files and, if a front is too large to reside in memory, it is temporarily held in a direct access file. In addition, information from the ordering and analyze phases may be held in sequential access files. We anticipate that the facility for out-of-core working and out-of-core storage of the matrix factor will allow the solution of problems that are too large for the other codes to successfully solve with the memory available in our test environment. The penalty of out-of-core working is possibly slower factorization and solve times because of I/O overheads.

2.5 Other Key Features

We conclude this section by briefly highlighting some of the other key features of sparse direct algorithms that are offered by some or all of the solvers in this study. All the codes employ supernodal techniques that enable dense linear algebra routines to be used to improve the efficiency of the factorization phase. All the packages except SPOOLES use high level Basic Linear Algebra Subprograms (BLAS) [Dongarra et al. 1990] and a number also employ LAPACK routines. Once the factors have been computed, they may be used to solve repeatedly for different right-hand sides b. Some codes offer the option of solving for more than one right-hand side at once because this enables them to take advantage of Level 3 BLAS in the solve phase (see Table V).

A number of codes offer options for automatically scaling the matrix and/or automatically performing iterative refinement to improve the quality of the computed solution and to help assess its accuracy (again, see Table V).

When solving problems that arise from finite-element applications, it is often convenient not to assemble the matrix A but to hold the matrix as a sum of element matrices. The only code in this study that allows A to be input in element form is MUMPS (but note that the HSL code MA62 which was included in our earlier study, is designed exclusively for positive definite problems in

Table V. Summary of Other Key Features of the Sparse Direct Solvers Used in this Study. (†) Denotes Default; (††) Denotes Default for Indefinite Problems

Code	Element Entry	Scaling	Out-of -core	Iterative Refinement	Multiple Rhs	Complex Symmetric	Hermitian
BCSLIB-EXT	×	×	√(†)	×	√	√	√
CHOLMOD	×	×	×	×	√	×	√
MA57	×	√(†)	×	√	√	√	×
MUMPS	√	√	×	√	√	√	×
Oblio	×	×	√	√	√	√	×
PARDISO	×	×	×	√(††)	√	√	√
SPOOLES	×	×	×	×	√	√	√
SPRSBLKLLT	×	×	×	×	√	×	×
TAUCS	×	×	√	×	×	√	√
UMFPACK	×	√(†)	×	√	×	√	√
WSMP	×	√(†)	×	√	√	√	√

unassembled element format), although a number of packages (in particular, BCSLIB-EXT) offer the user more than one input format for the assembled matrix A.

A summary of the key features of the solvers in this study that have not already been included in earlier tables is given in Table V.

3. SPARSE SYMMETRIC SOLVERS USED IN THIS STUDY

In this section, we give a very brief description of the software packages listed in Table I. We highlight some of the main features, with particular reference to our introductory discussion.

3.1 BCSLIB-EXT

BCSLIB-EXT is a library of mathematical software modules for solving large sparse linear systems and large sparse eigenvalue problems. It includes multifrontal solvers that are designed both for positive definite and indefinite symmetric systems.

When factorizing indefinite problems, the sequence may be modified and both 1×1 and 2×2 block pivots are used with a user-controlled threshold parameter u (with default value 0.01). Modifying the pivot sequence may cause additional fill-in in the matrix factor beyond that predicted by the analyze phase. BCSLIB-EXT allows the user to set a parameter that will cause the factorization to abort if this fill-in exceeds a prescribed level. By default, the factorization also terminates if a zero pivot is encountered. Alternatively, a parameter may be set to allow the package to perturb a nearly zero pivot and continue the factorization. The user can also request that the computation terminates as soon as a negative pivot is found. The size of the blocks used by the Level 3 BLAS routine GEMM during the factorization is controlled by parameters that may be reset by the user.

As already mentioned, a key feature of BCSLIB-EXT is its use of files to reduce the amount of main memory required by the package. The user can choose to hold the original matrix and/or the matrix factors in files. If there is not enough memory to hold the multifrontal stack and the current frontal matrix,

the code will store the stack out-of-core. It will also perform an out-of-core frontal assembly and factorization step if the current frontal matrix does not fit in memory. The user can choose a minimum core processing option that forces out-of-core storage. In our tests, we provide the amount of storage recommended in the documentation and provide positive stream numbers for each of the files used by the code. In this case, if the amount of main storage we have provided is insufficient, the code will use sequential files for holding information from the ordering and analyze phases and may use one or more files during the factorization and solve phase.

3.2 CHOLMOD

CHOLMOD is a set of ANSI C routines for solving sparse symmetric linear systems that is being developed by Tim Davis of the University of Florida. This is a very new package and is still being beta tested. In this study, we tested version 0.4; version 0.7 is now available but Davis advises us that both versions have the same performance.

CHOLMOD uses a left-looking supernodal factorization [Ng and Peyton 1993]. After a matrix is factorized, its factors can be updated or downdated using the techniques described by Davis and Hager [1999, 2001, 2005]. Interfaces to a number of ordering algorithms are provided, including approximate minimum degree, COLAMD ([Davis et al. 2004a, 2004b]), constrained minimum degree (CSYMAMD, CCOLAMD), and nested dissection using a modified version of METIS. By default, the code automatically chooses between an AMD variant and METIS. An AMD ordering is first computed. This is used if either the number of entries in the factor is less than 5 times the number of entries in the lower triangular part of A or the number of flops is less than 500 times the number of entries in the factor. Otherwise, the METIS ordering is computed and the better of the two then selected.

The intention is that CHOLMOD will be included in MATLAB as the default solver for sparse symmetric positive definite systems.

3.3 MA57

MA57 is part of the HSL mathematical software library [HSL 2004] and was designed by Duff [2004] to supersede the earlier well-known HSL multifrontal code MA27 [Duff and Reid 1983] for the solution of symmetric indefinite systems. Our earlier study [Gould and Scott 2004] compared the performance of version 1.0.0 using an AMD ordering with that of the nested dissection ordering from the METIS package. Our findings were that for very large, positive definite test problems (typically those of order $> 50,000$), it is often advantageous to use the METIS ordering but for small and very sparse problems and also for many indefinite problems, using an AMD ordering with quasi-dense row detection (QAMD) is preferable [Amestoy 1997]. Based on our findings and experiments by, among others, Duff and Scott [2005], MA57 has now been modified so that since version 3.0.0 the default has been for the code to automatically select the use of either the QAMD ordering or the METIS ordering based on the order

of the system and characteristics of the sparsity pattern; for some indefinite problems, it computes both orderings and chooses the one with the smaller predicted level of fill. Details of the current strategy are given in Duff and Scott [2005].

During the factorization phase, when diagonal 1×1 pivots would be numerically unstable, 2×2 diagonal blocks are used. Note that, for a given threshold parameter u, the test for stability of 2×2 pivots is less severe than the test used in the earlier MA27 code (details are given by Duff [2004]). If the problem is known to be positive definite, the user can set a control parameter that switches off threshold pivoting. In this case, if a sign change or a zero is detected among the pivots, an error exit will optionally occur. Alternatively, the user can choose to terminate the computation if any pivot is found to be of modulus less than a user-defined value.

Parameters that are under the user's control determine the size of the blocks used by the Level 3 BLAS routines during the factorization and solve phases. The use of iterative refinement is optional. The iterative refinement offered within the package is based on that of Arioli et al. [1989]. Estimates of the error may optionally be computed.

We have already observed that the default ordering in the version of the code used in this study differs from the one used in our earlier study [Gould and Scott 2004]. Since the earlier study there have been a number of other key changes to MA57. In particular, by default the new version scales the matrix using a symmetrized version of the HSL code MC64 [Duff and Koster 1999]. The aim is to put large entries in the diagonal blocks so as to restrict the number of pivots that are rejected for stability reasons during the factorization; details are given by Duff and Pralet [2004]. The matrix is explicitly scaled within the package as are the right-hand side and the solution so that the user need not be concerned with this. Iterative refinement, if requested, is based on the original unscaled matrix. Static pivoting is now an option so that the factorization can be performed using the storage predicted by the analysis phase even if the matrix is not positive definite. Because static pivoting is not the default strategy, it is not used in our tests.

There is little difference between the speed of the Fortran 90 version of MA57 and the Fortran 77 version, because the former is essentially a Fortran 90 encapsulation of the latter. However, the Fortran 90 version does offer some additional facilities, and the user interface is simplified through the use of dynamic storage allocation. In our numerical experiments, the Fortran 77 version is used.

3.4 MUMPS

The MUMPS (MUltifrontal Massively Parallel Solver) package is designed and developed by Amestoy et al. [2001] (see also Amestoy et al. [2000]). It is a multifrontal code for solving both symmetric and unsymmetric systems. Both C and Fortran 90 interfaces to MUMPS are available; in our numerical experiments, the Fortran 90 interface is used. MUMPS has been developed primarily as a parallel solver (originally targeted at distributed memory computers); in this study we use the sequential version. We note that between July 2005 and October 2005

there were five releases of MUMPS (versions 4.5.1 to 4.5.5). Version 4.5.1 represented a substantial upgrade, with a number of important new features (notably, improved algorithms for symmetric indefinite problems); the more recent versions offer minor upgrades and bug fixes. We use version 4.5.5 in our tests.

MUMPS offers the user a wide range of options for choosing the pivot sequence (see Table III). These include a version of QAMD, that is, AMD with automatic quasi-dense row detection [Amestoy 1997] and an approximate minimum fill-in algorithm. The multisection ordering is implemented using the code PORD [Schulze 2001] and nested dissection using one of the graph partitioning packages METIS or SCOTCH (www.labri.fr/Perso/~pelegrin/scotch/). By default, MUMPS automatically chooses the ordering algorithm depending on the packages installed, the size of the matrix, the density of the matrix, and the number of processors available. On a single processor, QAMD is used for problems of size $n < 10^4$ and for larger but very sparse matrices with quasi-dense rows. Otherwise, METIS_NodeND is used.

Version 4.5.1 introduced the use of 2×2 pivots for indefinite problems. By default, the analyze phase chooses the pivot sequence using a compressed graph, which is constructed using the numerical values and allows 2×2 pivots to be selected. The intention is that, by choosing 2×2 pivots in the analyze phase, the pivot sequence will require fewer modifications to maintain numerical stability during the factorization [Duff and Pralet 2005]. Because the features for indefinite problems are very new, we felt it would be useful to run the symmetric and unsymmetric versions of MUMPS (the unsymmetric version includes pivoting for stability) and so results for both versions are included for indefinite problems.

Other features of the MUMPS package include facilities for use in domain decomposition, error analysis, optional iterative refinement using the approach of Arioli et al. [1989], and estimation of rank deficiency. An option exists that allows the user to input sparse right-hand sides.

3.5 Oblio

Oblio is a sparse symmetric direct solver library developed by Dobrian and Pothen as an experimental tool [Dobrian et al. 2000]. Their goal was to create a "laboratory for quickly prototyping new algorithmic innovations, and to provide efficient software on serial and parallel platforms." The code is written in C++ using object-oriented techniques and is still being actively developed.

The most recent version (0.7) is able to solve both positive definite and indefinite systems. For indefinite problems, the user is offered so-called static LDL^T or dynamic LDL^T. In the former case, if a small pivot is encountered, it is perturbed to a value under the user's control, allowing the computation to continue. The default (which we use in our tests) is dynamic LDL^T. This employs a combination of 1×1 and 2×2 pivots. When a diagonal 1×1 pivot would be unstable, a search is made for a suitable 2×2 pivot. Thus searches for 1×1 and 2×2 pivots are interlaced.

For flexibility, Oblio implements three different sparse factorizations: left-looking, right-looking and multifrontal. For 2-dimensional problems the multifrontal option is recommended but for large 3-dimensional problems the

user documentation reports the multifrontal factorization can be outperformed by the other two algorithms. The default algorithm is the multifrontal algorithm and this is used in our tests. The multifrontal version includes an out-of-core option. This allows the matrix factor and/or the stack to be held in files.

3.6 PARDISO

The PARDISO package of Schenk and Gärtner offers serial and parallel solvers for the direct solution of unsymmetric and symmetric sparse linear systems on shared memory multiprocessors. In this study, only the serial version for symmetric systems is used. PARDISO employs a combination of left- and right-looking Level 3 BLAS supernode techniques [Schenk et al. 2000; Schenk and Gärtner 2004b] and is written using a combination of Fortran 77 and C source code. PARDISO is included in Intel Math Kernel Library (see www.intel.com/software/products/mkl/features/dss.htm).

The default ordering is a modified version of METIS; if the user does not wish to use this ordering, a fill-reducing ordering may be input. The user must set a parameter to indicate whether a Cholesky factorization or an LDL^T factorization is required. For indefinite problems, the current version includes Bunch-Kaufmann pivoting [Bunch and Kaufmann 1977] applied to the dense diagonal supernode blocks. A modified version of the LAPACK routine dsytf2 is used for factorizing these blocks. Pivots that are zero or nearly zero are perturbed so that pivots are not delayed beyond the current block. The amount by which pivots are perturbed in this static pivoting strategy is determined by a parameter under the user's control. The current version includes an option to perform preprocessing based on symmetric weighted matchings. The documentation states that this is very robust, but because it incurs an overhead and involves the numerical values of entries of the matrix, so that a new analyze may be required if the entries change, even if the sparsity pattern is unaltered, it is not the default and is not used in this study. Iterative refinement is offered, with the maximum number of steps controlled by a parameter set by the user. In our tests the default value of 0 is used. This means that iterative refinement is only used if pivots have been perturbed during the factorization. In this case, two steps of iterative refinement are performed.

We note that when calling PARDISO it is assumed that zero diagonal entries are stored explicitly in the list of matrix entries. For many indefinite examples, one or more of the diagonal entries are often not present within the sparsity pattern and the user must add explicit zeros. PARDSIO also requires that the upper triangular part of the matrix is entered by rows with the entries within each row ordered by increasing column index.

3.7 SPOOLES

SPOOLES is a library for solving sparse real and complex linear systems of equations, and may be used for both symmetric and unsymmetric problems. The package is written in C using an object-oriented design. Both serial and parallel versions are available. The serial version for real symmetric systems is used in our tests.

SPOOLES uses the Crout reduction variant of Gaussian elimination, which is a left-looking algorithm. In addition to MMD and generalized ND, the analyze phase offers a multisection ordering algorithm [Ashcraft and Liu 1998]. The default is to use the better of the nested dissection and multisection methods (although the user reference manual does comment on situations where the user may find it beneficial to select another choice).

To try and ensure stability of the factorization for indefinite problems, the entries of the triangular factor L are bounded by a user-supplied tolerance (in our tests we use the recommended value of 100 for this tolerance). The fast Bunch-Parlett algorithm described by Ashcraft et al. [1998] is used to choose 1×1 or 2×2 pivot blocks. SPOOLES is the only package tested that does not use the high level BLAS kernels; instead it performs operations within the factorization phase using multiple dot products.

We note that SPOOLES requires that the sparsity pattern of the input matrix includes the diagonal. For many indefinite examples, one or more of the diagonal entries is often not present within the sparsity pattern. In such cases, the user must include an explicit zero.

3.8 SPRSBLKLLT

SPRSBLKLLT was developed by Esmond Ng and Barry Peyton at Oak Ridge National Laboratory in the early 1990s for the solution of sparse symmetric positive definite systems. The pivot sequence is selected using the MMD algorithm; the implementation used is taken from the Waterloo sparse matrix package SPARSPAK (see www.cs.uwaterloo.ca/~jageorge/Sparspak/sparspak.html). The symbolic factorization subroutines are independent of any ordering algorithms.

SPRSBLKLLT implements a supernodal left-looking Cholesky factorization algorithm (details are given by Ng and Peyton [1993]). The symbolic factorization algorithm uses the results of Gilbert et al. [1994], which allow storage requirements to be determined in advance, regardless of the ordering strategy used. The performance of the package has been enhanced since it was first released by exploiting the memory hierarchy: it splits supernodes into sub-blocks that fit into the available cache; and it unrolls the outer loop of matrix-vector products in order to make better use of available registers. A parameter that must be set by the user determines the maximum supernode size. The storage requirements depend on this parameter (large values increase the storage). Based on the limited documentation provided with the code, in our tests this parameter is set to 100.

3.9 TAUCS

TAUCS has been developed since 2001 by Sivan Toledo's research group in the Department of Computer Science at Tel-Aviv University as a platform for research on sparse linear solvers. TAUCS is designed to support the development of research codes by providing a library of fundamental algorithms and services, and to facilitate the maintenance and distribution of the resulting research codes. Toledo and his colleagues are still developing the package; a version for

indefinite problems will be available in the future. TAUCS is currently used in Mathematica 5.

Both a multifrontal algorithm and a left-looking algorithm are implemented; the documentation states the latter is slower than the former but requires less memory. As well as MD, AMD, MMD, and METIS_NodeND, a no-fill ordering code for matrices whose graphs are trees is available. This is a special case of MD but is faster. METIS is recommended for large problems and was used in our tests. The current version of TAUCS is designed for positive definite symmetric problems, so numerical pivoting is not incorporated (although the package does include a general sparse LU factorization code with partial pivoting). An option exists to compute an incomplete LL^T factorization.

TAUCS is able to factorize a matrix whose factor is larger than the main memory by holding the factor out-of-core. The factor is held in multiple files, each at most 1 Gbyte in size (see Rotkin and Toledo [2004] for details). The user must use a different interface for the out-of-core working, so that switching between in-core and out-of-core working is not automatic. Since the out-of-core option is not the default, we have not included it in our tests results but our experience has been that it does allow larger problems to be solved than is otherwise possible.

3.10 UMFPACK

The principal author of the sparse direct solver UMFPACK is Tim Davis of the University of Florida [Davis 2003a, 2003b]. The tested version (version 4.1) is written in C; the original code was developed by Davis and Duff [1993] in Fortran 77. It is primarily written for unsymmetric matrices, that is, it requires the sparsity pattern of the whole matrix A and computes an LU factorization (there is no option to compute a Cholsky factorization). However, for symmetrically, or nearly symmetrically, structured matrices it offers a symmetric pivoting strategy and for this reason we were encouraged by Davis to include the package in this study. This also serves as a benchmark to illustrate how symmetric solvers compare to a state of the art unsymmetric package.

UMFPACK combines a column ordering strategy with a right-looking unsymmetric-pattern multifrontal numerical factorization. All pivots with zero Markowitz cost are eliminated first and placed in the LU factors. The analyze phase then automatically selects one of three ordering and pivoting strategies (*unsymmetric*, *2-by-2*, and *symmetric*). For symmetric matrices with a zero-free diagonal, the symmetric strategy is used. This computes a column ordering using AMD. No modification of the column ordering is made during the numerical factorization. A nonzero diagonal entry is selected as a suitable pivot if in magnitude it is at least u_1 times the largest entry in its column. Otherwise, an off-diagonal pivot is selected with magnitude at least u_2 times the largest entry in its column. The parameters u_1 and u_2 are under the user's control with default values of 0.001 and 0.1, respectively. Thus strong preference is given to pivoting on diagonal entries. For symmetric indefinite problems with zeros on the diagonal, the so-called 2-by-2 strategy is attempted. This looks for a

row permutation that puts nonzero entries onto the diagonal. The symmetric strategy is applied to the permuted matrix.

MATLAB, C and Fortran interfaces are offered. Version 4.3 appears as a built-in routine in MATLAB 7.0 and Mathematica 5 uses Version 4.1. An earlier version (2.2.1) by Davis and Duff is available as routine MA38 within the software library HSL. Versions prior to 4.1 only offer the unsymmetric pivoting strategy and are thus not well suited for matrices with a symmetric nonzero pattern. The current version is 4.6; however Davis advised that the performance of the version used in this study (version 4.1) is comparable; subsequent versions have added extra features and fixed minor bugs.

3.11 WSMP

The Watson Sparse Matrix Package (WSMP) was developed by Anshul Gupta of the IBM T. J. Watson Research Center. The package is written using Fortran 90 and C and includes direct solvers for both symmetric and unsymmetric systems. WSMP was primarily developed as a highly scalable parallel code that can be used in either a shared-memory multiprocessor or a message-passing environment. A serial version is available and is used in this study.

The analyze phase offers a minimum local fill ordering and an ordering based on recursive bisection. By default, both orderings are computed and the one that will result in the least fill-in is selected. The factorization phase implements a modified multifrontal algorithm. WSMP supports three types of factorization for symmetric matrices. The user can choose between LL^T, LDL^T without numerical pivoting, or LDL^T with 1×1 and 2×2 pivots. For the LDL^T factorization with pivoting, by default WSMP performs a preprocessing step that includes scaling and choosing a pivot sequence using the numerical values of the matrix. Thus the values of the matrix must be supplied to the analyze phase. If subsequent matrices with the same sparsity pattern but different numerical values are factorized, WSMP tracks the level of fill in the factors resulting from pivoting and automatically chooses whether or not to perform a new analysis.

WSMP includes a routine that may be called to perform iterative refinement, with an option of using extended precision arithmetic. WSMP requires that the sparsity pattern of the input matrix includes the diagonal. If one or more of the diagonal entries is not present, the user must add an explicit zero. The input matrix must be supplied by columns, with the diagonal entry at the start of each column list.

Further details of WSMP are given in Gupta et al. [1997, 2001]. Currently, WSMP is available for use on AIX, SunOS, Tru64, HP-UX, and Linux platforms. Although WSMP libraries contain multithreaded code, the libraries are not thread-safe.

4. THE TEST ENVIRONMENT

4.1 The Test Set

Our aim in this study is to test the solvers on a wide range of test problems from as many different application areas as possible. In collecting test data we

imposed only two conditions:

—The matrix must be of order greater than 10, 000.
—The data must be available to other users.

The first condition was imposed because our interest in this study is in large problems. The second condition was to ensure that our tests could be repeated by other users, and furthermore, it enables other software developers to test their codes on the same set of examples and thus to make comparisons with other solvers. Provided the above conditions are satisfied, we have included all square real symmetric matrices of order exceeding 10, 000 that were available in June 2003 in the Matrix Market (math.nist.gov/MatrixMarket/), the Harwell-Boeing Collection [Duff et al. 1989], and the University of Florida Sparse Matrix Collection (www.cise.ufl.edu/research/sparse/matrices), as well as a number of problems that were supplied to us by colleagues. The test set comprises 88 positive definite problems and 61 numerically indefinite problems. We note that some of the indefinite problems are highly ill-conditioned and 5 are structurally singular. Of these matrices, those of order 50,000 or more are further classed as being in the *subset* of larger examples (there are 43 positive definite and 30 indefinite examples in this category). Any matrix for which we only have the sparsity pattern available is included in the positive definite set, and appropriate numerical values have been generated (see Section 4.6). Application areas represented by our test set include linear programming, structural engineering, computational fluid dynamics, acoustics, and financial modelling. A full list of the test problems together with a brief description of each is given by Gould et al. [2005]. The problems are all available from ftp://ftp.numerical.rl.ac.uk/pub/matrices/symmetric (and are also now part of the University of Florida Sparse Matrix Collection).

4.2 The Performance Profile

Benchmark results are generated by running a solver on a set $\mathcal{T}$ of problems and recording information of interest, such as the computing time and memory used. In this study, we use a performance profile as a means to evaluate and compare the performance of the solvers on our test set $\mathcal{T}$.

Let $\mathcal{S}$ represent the set of solvers that we wish to compare. Suppose that a given solver $i \in \mathcal{S}$ reports a statistic $s_{ij} \geq 0$ when run on example j from the test set $\mathcal{T}$, and that the smaller this statistic the better the solver is considered to be. For example, s_{ij} might be the CPU time required to solve problem j using solver i. For all problems $j \in \mathcal{T}$, we want to compare the performance of solver i with the performance of the best solver in the set $\mathcal{S}$.

For $j \in \mathcal{T}$, let $\hat{s}_j = \min\{s_{ij}; i \in \mathcal{S}\}$. Then for $\alpha \geq 1$ and each $i \in \mathcal{S}$ we define

$$k(s_{ij}, \hat{s}_j, \alpha) = \begin{cases} 1 \text{ if } s_{ij} \leq \alpha \hat{s}_j \\ 0 \text{ otherwise.} \end{cases}$$

The *performance profile* [Dolan and Moré 2002] of solver i is then given by the function

$$p_i(\alpha) = \frac{\sum_{j \in \mathcal{T}} k(s_{ij}, \hat{s}_j, \alpha)}{|\mathcal{T}|}, \quad \alpha \geq 1.$$

Thus $p_i(1)$ gives the fraction of the examples for which solver i is the most effective (according to the statistic s_{ij}), $p_i(2)$ gives the fraction for which it is within a factor of 2 of the best, and $\lim_{\alpha \longrightarrow \infty} p_i(\alpha)$ gives the fraction for which the algorithm succeeded. We note that software for creating performance profiles is available at `www-unix.mcs.anl.gov/~more/cops/`.

In this study, the statistics used are:

—The CPU times required to perform the analyze, factorize, and solve phases.

—The number of nonzero entries in the matrix factor.

—The total memory used by the solver.

The range of α illustrated is chosen in each case to highlight the dominant trends in the data. We note that for the number of entries in the factor we rely on the statistics that are reported by the solvers. For the package `BCSLIB-EXT` this statistic is not available.

4.3 Computing Platform

The numerical results were all obtained on a Compaq DS20 Alpha server with a pair of EV6 CPUs; in our experiments only a single processor with 3.6 Gbytes of RAM was used. We compiled the codes with full optimization; the vendor-supplied BLAS were used where applicable. One processor operates at 500MHz with a peak performance of 1 Gflop. For `GEMM`, the peak is 415 Mflops, while the LAPACK Cholesky factorization routine `POTRF` performs at 342 Mflops. All CPU reported times are in seconds, and where appropriate, include all I/O costs involved in holding the factors in direct-access files. A CPU limit of 30 minutes was imposed for each code on each problem; any code that had not completed after this time was recorded as having failed.

In all the experiments, double precision reals were used. Thus storage for a real was 8 bytes and for an integer was 4 bytes. Memory is measured using the C utility function `getrusage`. In particular, the maximum resident set size of the current process is measured. Extra memory required for setting up the test is subtracted.

4.4 Control Parameters

Each of the sparse solvers used in our numerical experiments has a number of parameters that control the action. These are either assigned default values through a call to an initialization subroutine or the values recommended in the user documentation are used. Unless otherwise stated, we use these defaults in each case, even if different codes sometimes choose a different value for essentially the same parameter. The main exception is the stability threshold parameter u (see Section 2.3). We remark that we decided early on in this study not to try and fine tune the input parameters for each solver on each problem. In some solvers, on some problems, performance would have been improved (possibly substantially) by tweaking. However, we felt that to do this for each individual code and all the problems in the test set would be an almost impossible task, and more importantly, our aim is to compare the codes from a common standpoint, that is using the control settings chosen by the authors

of the packages. Our experience is that many users (in particular, those who would regard themselves as non-experts) rely on the default settings and are reluctant to try other values (possibly because they do not feel confident about making other choices).

When testing the solvers on positive definite problems the threshold parameter u is set to zero. This results in no numerical pivoting being performed. For our tests on numerically indefinite problems, for the codes that employ a stability threshold parameter, we run both with the code's default u value and with u set to 10^{-10}. Such a value is frequently used in optimization applications [Saunders 1994; Gould and Toint 2002], where speed is of the essence, and any instability is countered either by iterative refinement, or ultimately by refactorization with a larger value of u.

MA57, MUMPS and BCSLIB-EXT use a default threshold $u = 0.01$, while UMFPACK has two threshold parameters with default values of 0.001 and 0.1 (see Section 3.10). When testing with a small threshold, both UMFPACK parameters are set to 10^{-10}.

4.5 Out-of-Core Factorization

Out-of-core options are offered by the packages BCSLIB-EXT, Oblio and TAUCS. In our tests, the out-of-core facilities are only used if this is the default. For Oblio and TAUCS, the user must decide explicitly if the out-of-core option is required; as it is not the default, we do not select this option. By default, BCSLIB-EXT switches automatically to out-of-core working if it finds that the user has provided insufficient workspace for the code to run in-core (see Section 3.1). We therefore anticipate that out-of-core working will be used by BCSLIB-EXT for some of our largest test examples.

4.6 Numerical Values and Scaling

Some of our test examples are not supplied with numerical values (only the sparsity pattern is available). For these cases, appropriate numerical values are generated. Reproducible pseudo-random off-diagonal entries in the range $(0, 1)$ are generated using the HSL routine FA14, while the i-th diagonal entry is set to $\max(100, 10\rho_i)$, where ρ_i is the number of off-diagonal entries in row i of the matrix, thus ensuring that the generated matrix is numerically positive definite.

In all our tests, right-hand side vectors b are computed so that the exact solution x (of the unscaled system) is $x = e \stackrel{\text{def}}{=} (1, 1, \ldots, 1)^T$.

If the input matrix has entries differing widely in magnitude, then an inaccurate solution may be obtained in the indefinite case and the accuracy may be difficult to assess in all cases. A number of the packages tested include an option for scaling the input matrix. We do not use these options unless scaling is performed by default (this is the case for MA57 and UMFPACK and, for indefinite problems, WSMP). To examine the effects of scaling on the codes that do not perform scaling by default, for each value of the threshold parameter u used, we run both with and without scaling of the matrix A; and the corresponding right-hand side b using the HSL scaling routine MC30. For our positive definite

problems, scaling was found to make an insignificant difference, hence we report on the effects of scaling only for the indefinite examples.

4.7 Residuals and Iterative Refinement

A number of the solvers include routines for automatically performing iterative refinement. Unless the solver's default is to perform iterative refinement, we have not used these routines in this study (by default for indefinite problems `PARDISO` performs up to two steps of iterative refinement if pivots have been perturbed). Instead, once we have computed the approximate solution x, we perform one step of iterative refinement by computing the residual $r = Ax - b$ and then recalling the solve routine to solve $A\delta x = r$ for the correction δx.

For each right-hand side b and corresponding solution x, we compute the scaled residual

$$\|b - Ax\|_\infty / (\|A\|_\infty \|x\|_\infty + \|b\|_\infty).$$

A check is made after one step of iterative refinement that this residual is sufficiently small. In our tests, a residual greater than 0.0001 causes an error message to be returned. We fully recognize that in many applications greater accuracy is required but this value was chosen so that we could flag cases where a solver has failed completely; in general, the computed residuals are considerably smaller. Indeed, for our positive definite problems, the computed residuals were at most $\mathcal{O}(10^{-15})$ for each of the solvers on each of the problems that was successfully solved within our CPU time limit (the computed residuals are reported in full by Gould et al. [2005]). Note that the residual of the unscaled system is computed.

For nonsingular A, we also check the accuracy of the computed solution. Some of the systems are highly ill-conditioned and for these the norm of the error $x - e$ was large for some solvers. A positive warning flag is set in this case, but we do not count this as a failure provided the scaled residual is small.

5. RESULTS

Since some of the solvers we are examining are specifically designed for positive definite problems (and may be unreliable, or even fail, on indefinite ones), we will discuss the positive definite and indefinite cases separately. Moreover, as the competing algorithms have different design goals, we consider it worth examining each of the solution phases (analyze, factorize, solve) both separately and ultimately together.

Full details of the statistics generated by each solver are given in an accompanying technical report [Gould et al. 2005].

Note that in the remainder of this article, a failure means either the complete solution (analyze plus factorize plus solve) failed or the residual was unacceptably large. The exact reason for a code's failure on a particular problem is reported by Gould et al. [2005].

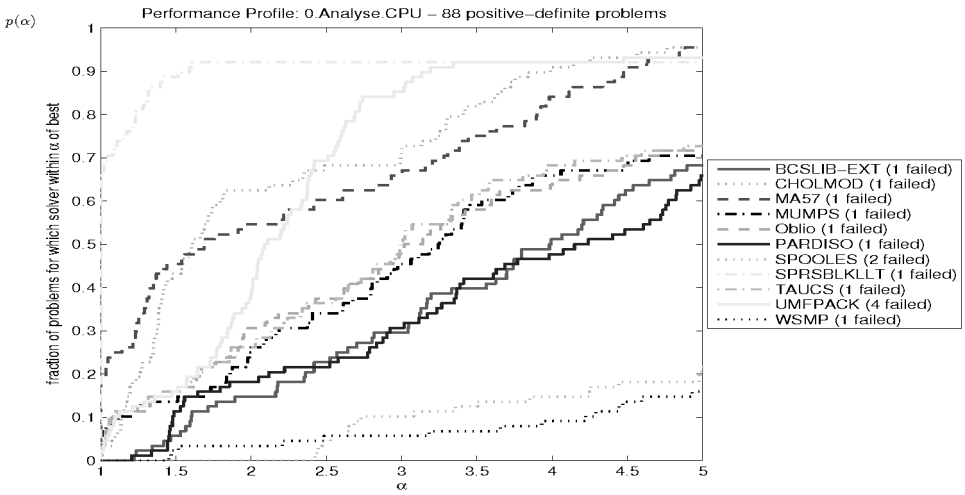

Fig. 1. Performance profile, $p(\alpha)$: CPU time for the analyze phase (positive definite problems).

5.1 Positive Definite Examples

Overall, the reliability of the solvers for positive definite examples was excellent. With the restrictions we imposed on the numerical experiments, all the solvers failed to solve the problem `audikw_1` because of a lack of space required to hold its factors,[1] but for the majority this was the only failure. `UMFPACK` was the solver with the largest number of failures, caused either by the CPU time limit being exceeded or by a lack of space. This is the only solver for which extra precautions must be taken to guarantee stability, because it permits off-diagonal pivoting.

In Figure 1 we present the performance profile for the analyze time for the eleven solvers. It is immediately apparent that in general the solvers that use (or select) variants of the minimum degree strategy (`SPRSBLKLLT`, `MA57`, `CHOLMOD` and `UMFPACK`) have a faster ordering than those that employ a dissection-based strategy (`MUMPS` also selects minimum degree for some problems but it appears to be no faster than some of the codes that use only nested dissection). The most expensive strategies are those employed by `SPOOLES` and `WSMP`, both of which compute two orderings and then select the better (they are also the only two codes that implement their own nested dissection algorithm while the other solvers offering nested dissection use METIS).

When it comes to the factorization, we see in Figure 2 that the careful analysis strategy adopted by `WSMP` pays off. Over the complete set of positive definite examples, the other codes (with the exception of `UMFPACK` and `SPOOLES`) are broadly comparable. Interestingly, the differences between left/right-looking and multifrontal factorizations do not seem as significant as might have

[1]In the case of `BCSLIB-EXT`, which permits out-of-core factorization, the run was terminated because of excessive CPU time. However, subsequent experiments showed that `BCSLIB-EXT` was able to solve the problem if sufficient time (roughly 2.5 CPU hours) was allowed. `TAUCS` is also able to solve this problem if the out-of-core option is selected.

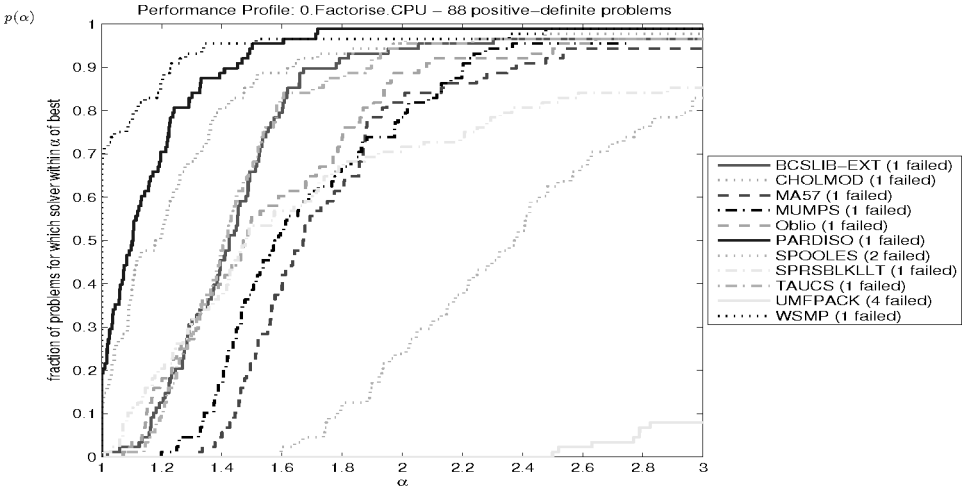

Fig. 2. Performance profile, $p(\alpha)$: CPU time for the factorization phase (positive definite problems).

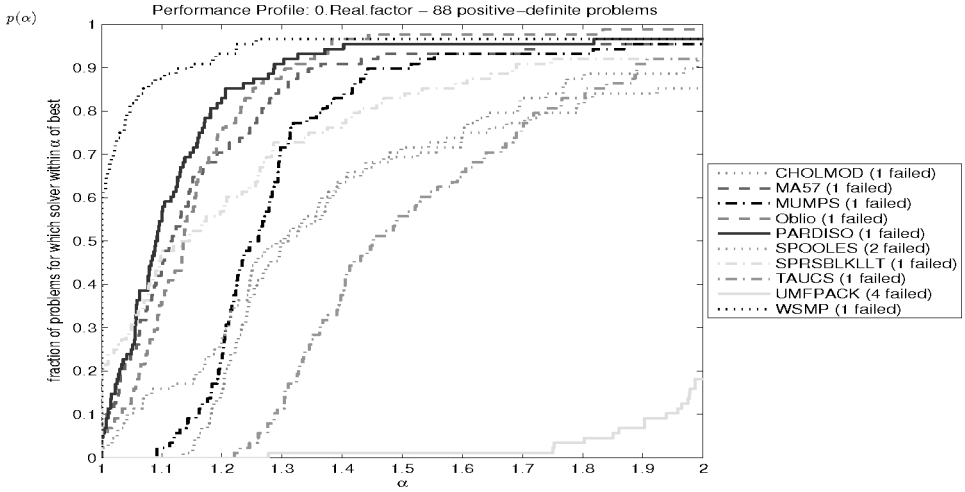

Fig. 3. Performance profile, $p(\alpha)$: Number of entries in the factors (positive definite problems).

been anticipated. UMFPACK is slower because it is essentially an unsymmetric solver and this agrees with our observations in our earlier article [Gould and Scott 2004] concerning the unsymmetric HSL code MA48. We believe that SPOOLES is not competitive because it does not use high level BLAS. We also see in Figure 3 that in our computing environment, the fastest factorization is generally closely tied to the number of nonzeros in the generated factors (for BCSLIB-EXT statistics on the number of entries in the factors are not available).

Having computed the factors, the performance profiles for solving for a single right-hand side are illustrated in Figure 4. Here there is a reasonable correlation between sparsity in the factors and time taken, with PARDISO, BCSLIB-EXT,

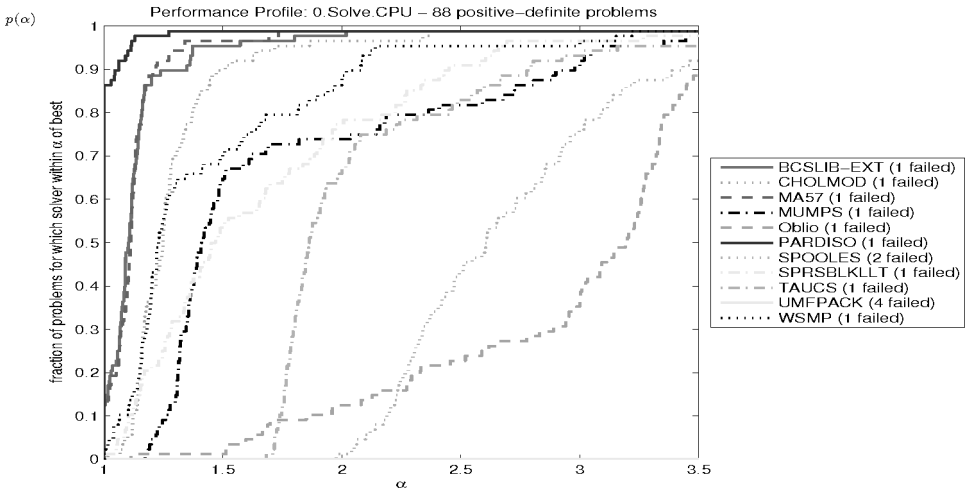

Fig. 4. Performance profile, $p(\alpha)$: CPU time for the solution phase (positive definite problems).

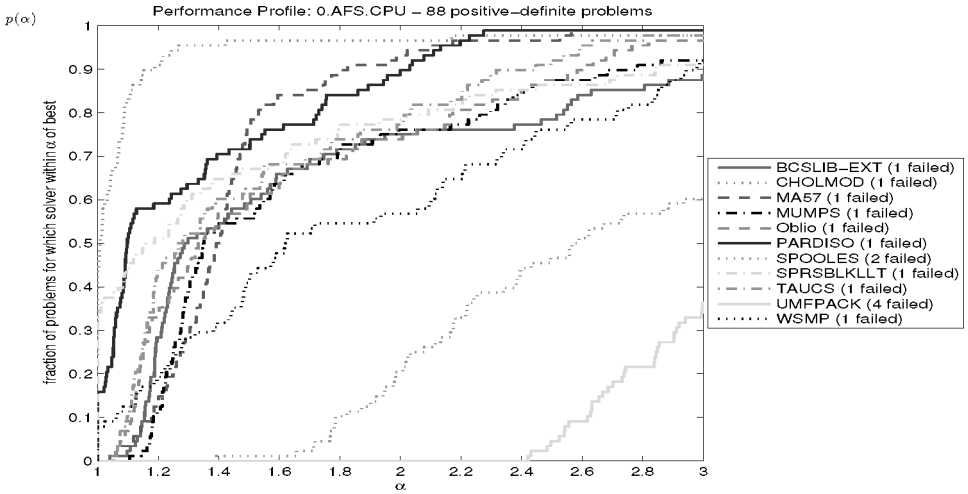

Fig. 5. Performance profile, $p(\alpha)$: CPU time for the complete solution (positive definite problems).

and MA57 generally the faster codes. The only slight surprise is that, although WSMP produces the sparsest factors, its solve time is longer than most of the other codes.

In Figure 5 we present the performance profile for the CPU time for a single solution (that is, the CPU time for analyzing, factorizing and solving for a single right-hand side) for the eleven solvers under consideration. It appears that the newest code CHOLMOD offers the best balance between the three solution phases and so gives the best overall performance. Of the remaining solvers, PARDISO and MA57 appear to perform marginally better than the rest, but SPRSBLKLLT, TAUCS, OBLIO, MUMPS, and BCSLIB-EXT are all close. The slow analyze and solve

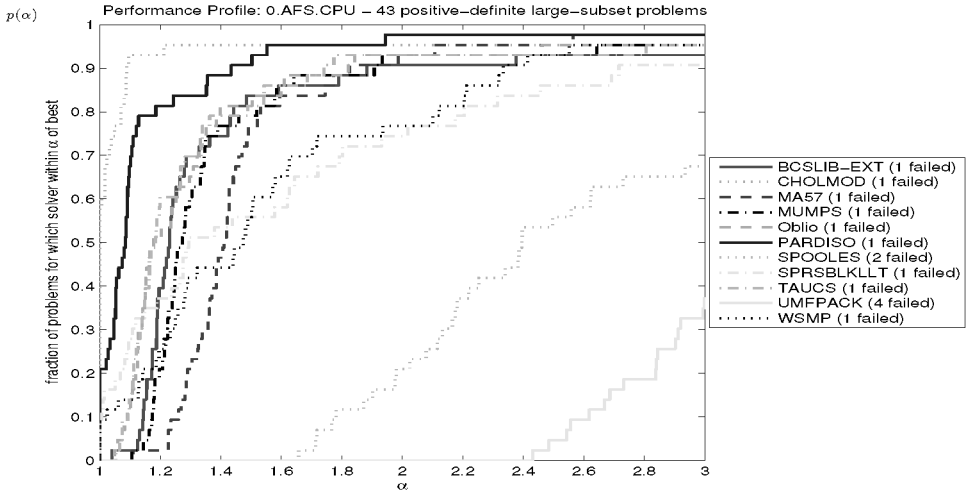

Fig. 6. Performance profile, $p(\alpha)$: CPU time for the complete solution (large positive definite subset problems).

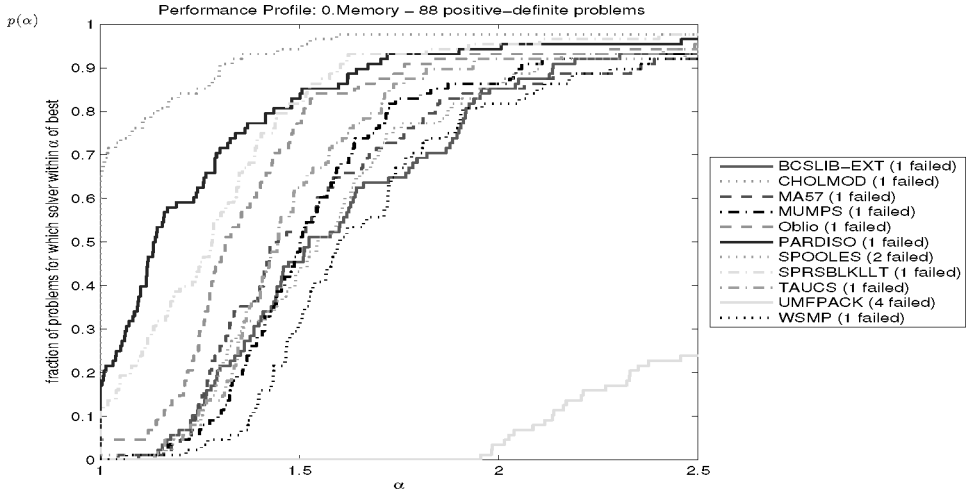

Fig. 7. Performance profile, $p(\alpha)$: Memory used (positive definite problems).

times clearly affect WSMP. Only SPOOLES and UMFPACK are uncompetitive. For the subset of larger problems illustrated in Figure 6, all the codes except SPOOLES and UMFPACK perform within a factor of two of the fastest code on around three quarters of the large test problems.

In Figure 7 we also compare the total memory used. We had expected that the multifrontal solvers would require significantly more memory than the other codes but our results suggest that there is generally little to distinguish between any of the symmetric solvers from this perspective, although CHOLMOD generally requires the least memory.

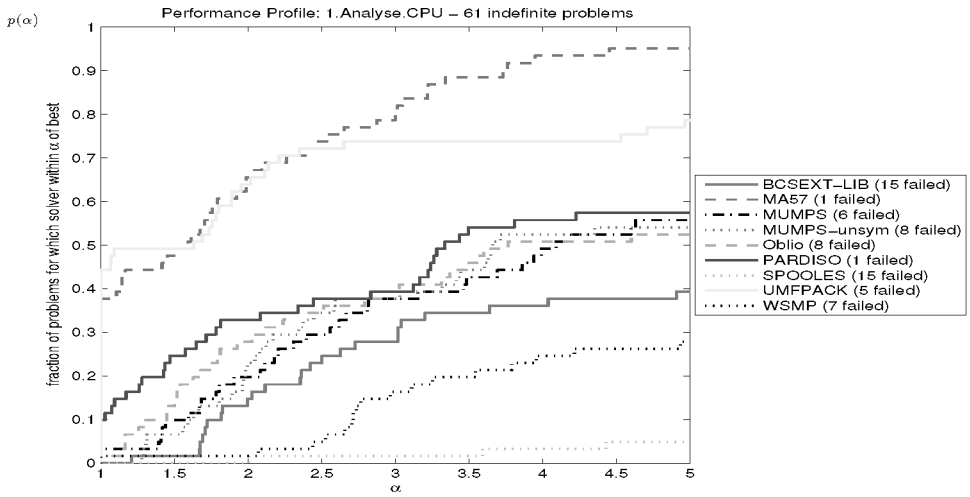

Fig. 8. Performance profile, $p(\alpha)$: CPU time for the analyze phase (indefinite problems).

5.2 Indefinite Examples

We now turn to indefinite problems, for which numerical pivoting is important. We need to assess the effects of the different ordering and pivoting strategies. Note that CHOLMOD, SPRSBLKLLT and TAUCS were not designed for indefinite problems and thus are omitted from these tests. Moreover, as discussed in Section 2.3, many of the other solvers only offer limited forms of pivoting, and thus give no stability guarantees. At its authors' suggestion, we include results for both the symmetric (here denoted by MUMPS) and unsymmetric (MUMPS_unsym) versions of MUMPS (note that the unsymmetric version includes off-diagonal pivoting).

Although, in our companion article [Gould et al. 2005], we report on the results of four different pre scaling/pivoting strategies, here we largely restrict our attention to the default strategy. The first thing to note is that the general reliability for indefinite problems is far below that for the definite case. Only MA57 and PARDISO had just one failure, while some of the solvers failed on 20% or more of the problems—admittedly some of the latter issued strong warnings in their documentation about possible limitations (including not being able to factorize singular systems and not performing numerical pivoting). All the solvers failed on the problem SPARSINE either because of a lack of space or because they exceeded our 30 minute CPU limit.

We start by presenting in Figure 8 the performance profile for the analyze times. The conclusions are broadly as for the definite case, with those solvers that use (or select) variants of the minimum degree strategy being faster than those opting by default for dissection orderings.

Now examining the factorize times (see Figure 9), we see a significant gap between PARDISO and the remaining solvers. Recall that PARDISO employs static pivoting and thus does not need to alter the ordering suggested by the analyze phase to complete its factorization. One might anticipate a lack of robustness

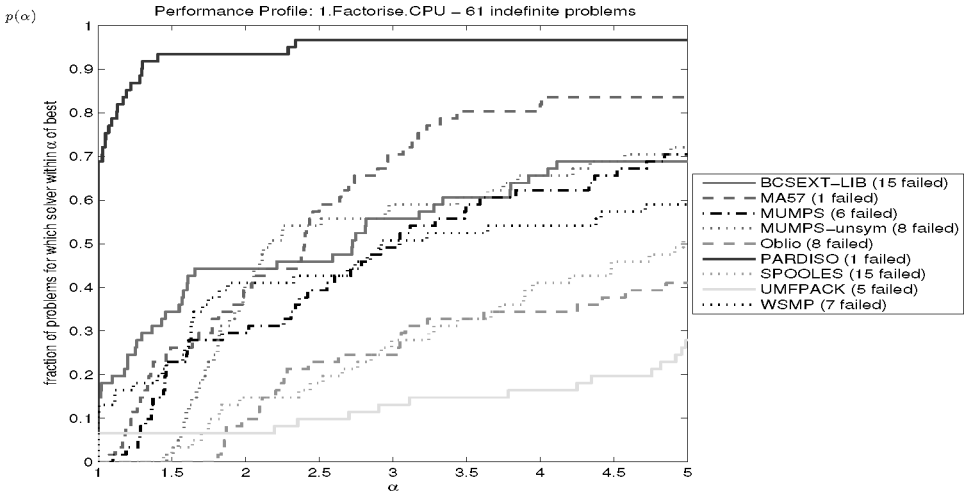

Fig. 9. Performance profile, $p(\alpha)$: CPU time for the factorization phase (indefinite problems).

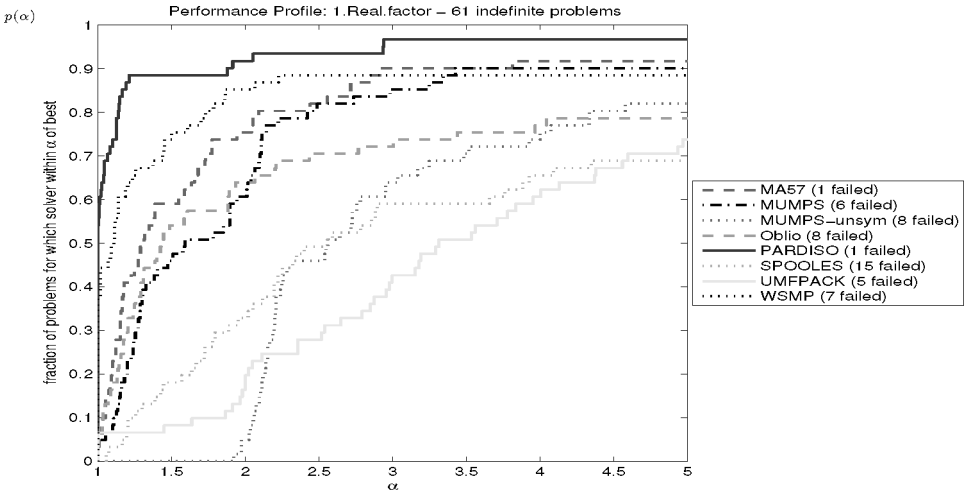

Fig. 10. Performance profile, $p(\alpha)$: Number of entries in the factors (indefinite problems).

with such an approach; what we find is that, by using the default iterative refinement, all the problems pass our residual test (see Section 4.7), but we must emphasise that for a small number of problems the scaled residuals from PARDISO are significantly larger than those obtained using other solvers. Specifically, for CONT-300, crystk02 and crystk03, they are $\mathcal{O}(10^{-9})$, $\mathcal{O}(10^{-7})$ and $\mathcal{O}(10^{-6})$, respectively; for all the other test examples, the residuals from PARDISO after one step of iterative refinement are at most $\mathcal{O}(10^{-10})$. Of course, in some applications, these larger residuals from PARDISO may be unacceptably large. Interestingly, the gap in performance is less pronounced when comparing the numbers of entries in the factors (see Figure 10), with WSMP, the runner up—again the statistics for BCSLIB-EXT are not available. As one might predict,

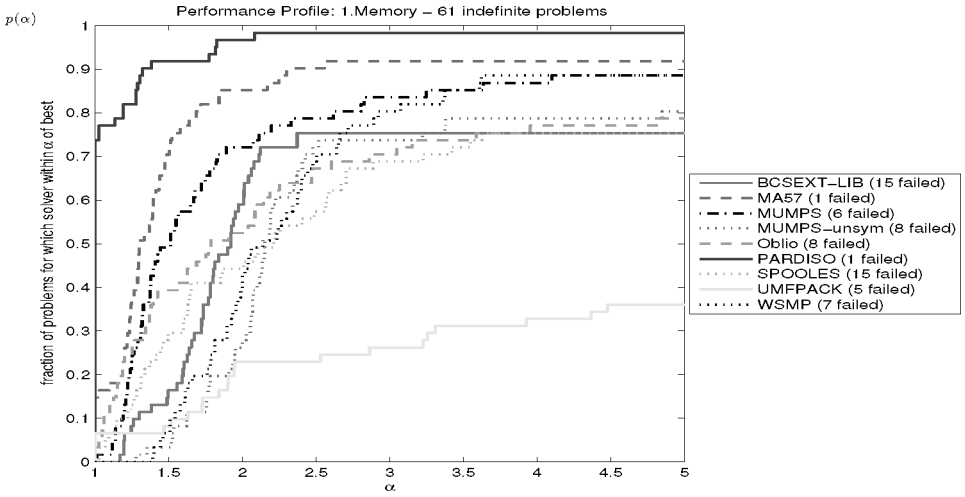

Fig. 11. Performance profile, $p(\alpha)$: Memory used (indefinite problems).

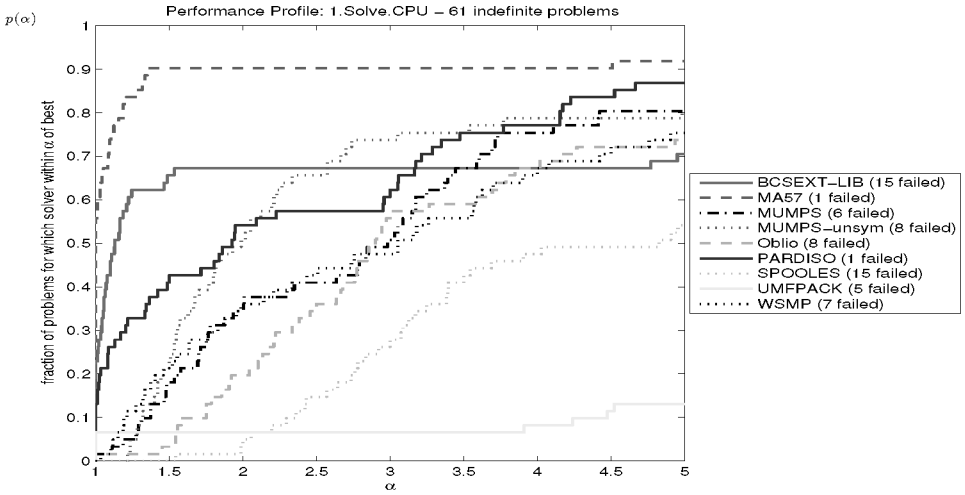

Fig. 12. Performance profile, $p(\alpha)$: CPU time for the solution phase (indefinite problems).

Figure 11 indicates there is generally a good correlation between the total memory used and the numbers of nonzeros in the factors (see Figure 10), with PARDISO requiring the least memory, followed by MA57.

Of course, there is some penalty to be paid for using a potentially less stable factorization, and that is that iterative refinement is a necessary precaution when using the generated factors to solve $Ax = b$. This is apparent in Figure 12. Now MA57 is a clear winner (with BCSLIB-EXT also performing well on the problems it solved within the CPU time limit), while PARDISO, which performs iterative refinement when pivots have been perturbed, is slower. A closer investigation of the detailed results shows that, if pivots have been perturbed during the PARDISO factorization, the corresponding solve can be up to three

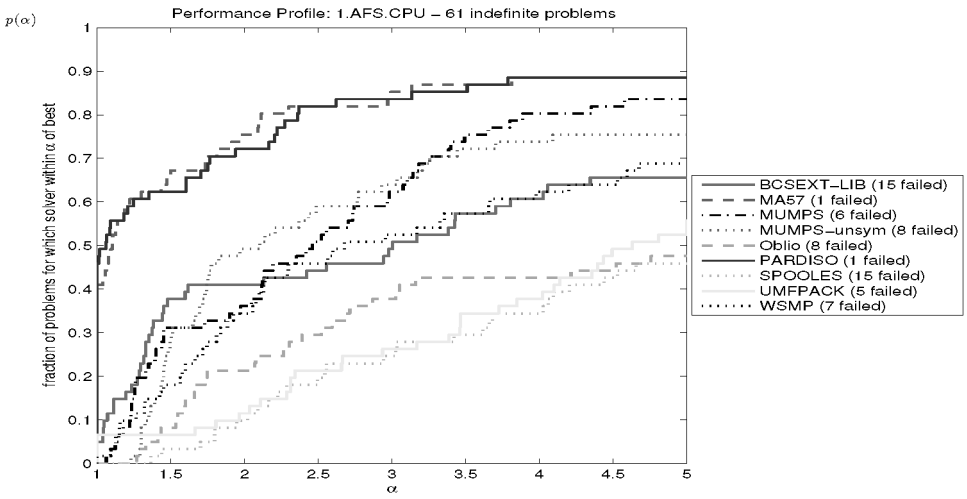

Fig. 13. Performance profile, $p(\alpha)$: CPU time for the complete solution (indefinite problems).

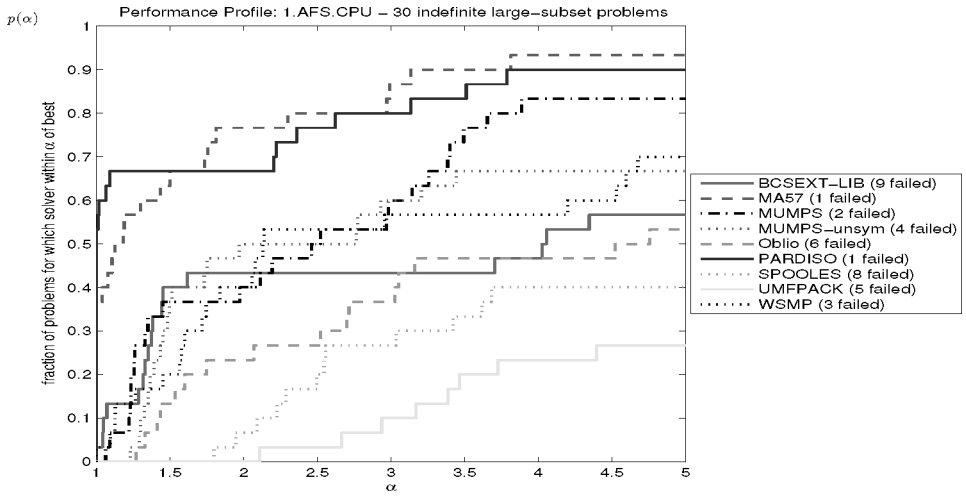

Fig. 14. Performance profile, $p(\alpha)$: CPU time for the complete solution (large indefinite subset problems).

times slower than the comparable MA57 solve precisely because of the possible two extra "refinement" steps taken. The solve time for WSMP is less competitive than might have been anticipated from the sparsity of its factors.

If a complete solution (analyze-factorize-solve) is the primary concern, Figure 13 indicates a clear preference for MA57 and PARDISO. In terms of CPU time, there is little to choose between the two. Of the remaining solvers, BCSLIB-EXT, the two variants of MUMPS, and WSMP perform best. This trend is reinforced when the subset of larger problems is considered (see Figure 14).

We also investigated using a small stability threshold parameter (see Section 4.4). In some cases, this improved the quality of the factorization (reduced both the CPU time and numbers of nonzeros in the factors), but at the cost

of lower overall reliability for some solvers (for example, without employing iterative refinement, MA57 did not solve three additional problems with the required accuracy and for UMFPACK a further twelve failures occurred). But for other solvers (BCSLIB-EXT and MUMPS) there were fewer failures overall either because, with the smaller threshold, the solver completed within our time limit or because less space was required because of fewer delayed pivots. But again, for some problems, the residuals were significantly larger than those obtained with the default threshold parameter. Using external scaling (see Section 4.6) did not appear to offer a significant or consistent advantage.

6. CONCLUDING REMARKS

In this article, we have compared a number of stand alone software packages for the direct solution of real symmetric large sparse linear systems of equations. Our tests were restricted to using the default settings. Although there are detailed differences, all the methods we have considered broadly comprise three phases: an analysis of the sparsity pattern with a view to reordering the variables to reduce fill-in, a (static or dynamic) factorization of the reordered matrix, and a solution of the given system using forward- and back-substitution. The interaction between all three phases is crucial for a reliable and fast solution. Thus although minimum degree based analysis phases generally appear to be faster than dissection-based ones, for many very large problems (typically of order greater than $50,000$) the resulting factors are generally less sparse, and this negatively influences the speed of both the subsequent factorization and solve phases.

For positive definite systems, we find in general that there is little in terms of reliability and efficiency to distinguish between the leading competitors (BCSLIB-EXT, CHOLMOD, MA57, MUMPS, Oblio, PARDISO, SPRSBLKLLT, TAUCS and WSMP), with a slight edge overall for CHOLMOD. Nevertheless, if many factorizations of matrices with identical sparsity patterns but differing values are required, WSMP, PARDISO, and, CHOLMOD are the strongest candidates, while if many solutions for a given matrix are needed BCSLIB-EXT, MA57, and PARDISO can be recommended. For indefinite problems, the leading contenders here are arguably MA57 and PARDISO. The former is more cautious with its factorization phase (and consequently the latter is faster), but such caution pays off in a faster solution phase as there is less need to resort to iterative refinement to correct for poor residuals. Both of these codes are being actively developed; indeed, both codes have been significantly improved since we started work on this study, partly as a result of feedback from us. The careful use of static pivoting within PARDISO [Schenk and Gärtner 2004a] is surprisingly effective and currently under investigation by the authors of other packages (see, for example, Duff and Pralet [2005]). Other codes that are still being actively developed include CHOLMOD, MUMPS and WSMP. The later two have introduced 2×2 pivoting and numerical analyze phases in recent releases, which have significantly enhanced their performnce.

As we discussed in Section 4.4, we have limited our experiments to running each of the packages with its default (or recommended) settings. Clearly, for

many problems it may be possible to get an individual code to run significantly faster and produce sparser factors by tuning the control parameters to the problem (indeed, a particular parameter choice may enable a code to succeed where we report a failure). From the brief descriptions of the codes and their key features given in Sections 2 and 3, it should be apparent that some of the codes offer the user a large number of parameters that can be used to tune the code for particular applications. Notable examples of packages that offer the user many options are `BCSLIB-EXT` and `MA57`, while `MUMPS` has the largest number of ordering options. Uniquely, `Oblio` offers the user the possibility of trying different factorization algorithms.

A further limitation of this study is that all our experiments were performed using a single computing platform. Although our main concern is how the codes perform relative to each other rather than individual CPU timings, clearly there could be some variation in performance on different computing platforms. Of course, some of our reported statistics (such as the number of entries in the factors and memory usage) are independent of the platform.

We readily concede that this article is merely a snapshot of an evolving field, and that perhaps a different picture will emerge in the not-too-distant future. Nevertheless, since the solution of large sparse linear systems is a vital component in many areas of scientific computation, we believe that our article will be useful to both software developers and potential users as a guide to the current state of the art of sparse direct symmetric solvers.

ACKNOWLEDGMENTS

We would like to thank the authors of the solvers used in this study who supplied us with copies of their codes and documentation, helped us to use the software, answered our queries, and commented on a draft of this article. In particular, we are grateful to Patrick Amestoy, Cleve Ashcraft, Tim Davis, Florin Dobrian, Iain Duff, Jean-Yves L'Excellent, Anshul Gupta, John Lewis, Esmond Ng, Alex Pothen, Stephane Pralet, Olaf Schenk, Sivan Toledo, and David Wah. Our thanks also to those who supplied test problems, including Mario Arioli, Christian Damhaug, Tim Davis, Anshul Gupta, Alison Ramage, Olaf Schenk, Miroslav Tuma, and Andy Wathen. We gratefully acknowledge the constructive criticism of three anonymous referees.

REFERENCES

AMESTOY, P. 1997. Recent progress in parallel multifrontal solvers for unsymmetric sparse matrices. In *Proceedings of the 15th World Congress on Scientific Computation, Modelling and Applied Mathematics, IMACS 97*, Berlin.

AMESTOY, P., DAVIS, T., AND DUFF, I. 1996. An approximate minimum degree ordering algorithm. *SIAM J. Matrix Anal. Appl. 17*, 886–905.

AMESTOY, P., DAVIS, T., AND DUFF, I. 2004. Algorithm 837: AMD, an approximate minimum degree ordering algorithm. *ACM Trans. Math. Softw. 30,* 3, 381–388.

AMESTOY, P., DUFF, I., AND L'EXCELLENT, J. 2000. Multifrontal parallel distributed symmetric and unsymmetric solvers. *Comput. Methods in Appl. Mech. Eng 184*, 501–520.

AMESTOY, P., DUFF, I., L'EXCELLENT, J., AND KOSTER, J. 2001. A fully asynchronmous multifrontal solver using distributed dynamic scheduling. *SIAM J. Matrix Anal. Appl. 23*, 15–41.

ARIOLI, M., DEMMEL, J., AND DUFF, I. 1989. Solving sparse linear systems with sparse backward error. *SIAM J. Matrix Anal. Appl. 10*, 165–190.

ASHCRAFT, C., GRIMES, R., AND LEWIS, J. 1998. Accurate symmetric indefinite linear equation solvers. *SIAM J. Matrix Anal. Appl. 20*, 513–561.

ASHCRAFT, C. AND LIU, J. 1998. Robust ordering of sparse matrices using multisection. *SIAM J. Matrix Anal. Appl. 19*, 816–832.

BUNCH, J. AND KAUFMANN, L. 1977. Some stable methods for calculating inertia and solving symmetric linear systmes. *Math. Comput. 31*, 163–179.

BUNCH, J., KAUFMANN, L., AND PARLETT, B. 1976. Decomposition of a symmetric matrix. *Numerische Mathematik 27*, 95–110.

DAVIS, T. 2003a. Algorithm 832: UMFPACK, an unsymmetric-pattern multifrontal method. *ACM Trans. Math. Softw. 30,* 2, 196–199.

DAVIS, T. 2003b. A column preordering strategy for the unsymmetric-pattern multifrontal method. *ACM Trans. Math. Softw. 34*, 2, 165–195.

DAVIS, T. AND DUFF, I. 1993. An unsymmetric-pattern multifrontal method for sparse LU factorization. Tech. Rep. RAL-93-036, Rutherford Appleton Laboratory.

DAVIS, T. A., GILBERT, J. R., LARIMORE, S. I., AND NG, E. G. 2004a. Algorithm 836: COLAMD, a column approximate minimum degree ordering algorithm. *ACM Trans. Math. Softw. 30*, 3, 377–380.

DAVIS, T. A., GILBERT, J. R., LARIMORE, S. I., AND NG, E. G. 2004b. A column approximate minimum degree ordering algorithm. *ACM Trans. Math. Softw. 30*, 3, 353–376.

DAVIS, T. A. AND HAGER, W. W. 1999. Modifying a sparse Cholesky factorization. *SIAM J. Matrix Anal. Appl. 20*, 3, 606–627.

DAVIS, T. A. AND HAGER, W. W. 2001. Multiple-rank modifications of a sparse Cholesky factorization. *SIAM J. Matrix Anal. Appl. 22*, 4, 997–1013.

DAVIS, T. A. AND HAGER, W. W. 2005. Row modifications of a sparse Cholesky factorization. *SIAM J. Matrix Anal. Appl. 26*, 3, 621–639.

DOBRIAN, F., KUMFERT, G., AND POTHEN, A. 2000. The design of sparse direct solvers using object-oriented techniques. In *Advances in Software Tools in Scientific Computing*, H. Langtangen, A. Bruaset, and E. Quak, Eds. Lecture Notes in Computational Science and Engineering, vol. 50. Springer-Verlag, 89–131.

DOLAN, E. AND MORÉ, J. 2002. Benchmarking optimization software with performance profiles. *Mathematical Programming 91*, 2, 201–213.

DONGARRA, J., DUCROZ, J., DUFF, I., AND HAMMARLING, S. 1990. A set of level 3 basic linear algebra subprograms. *ACM Trans. Math. Softw. 16*, 1, 1–17.

DONGARRA, J., DUFF, I., SORENSEN, D., AND VAN DER VORST, H. 1998. *Numerical Linear Algebra for High-Performance Computers*. SIAM.

DUFF, I. 2004. MA57—a new code for the solution of sparse symmetric definite and indefinite systems. *ACM Trans. Math. Softw. 30*, 118–154.

DUFF, I., ERISMAN, A., AND REID, J. 1986. *Direct Methods for Sparse Matrices*. Oxford University Press, England.

DUFF, I., GRIMES, R., AND LEWIS, J. 1989. Sparse matrix test problems. *ACM Trans. Math. Softw. 15*, 1–14.

DUFF, I. AND KOSTER, J. 1999. The design and use of algorithms for permuting large entries to the diagonal of sparse matrices. *SIAM J. Matrix Anal. Appl. 20*, 889–901.

DUFF, I. AND PRALET, S. 2004. Strategies for scaling and pivoting for sparse symmetric indefinite problems. Tech. Rep. RAL-TR-2004-020, Rutherford Appleton Laboratory. To appear in SIAM J. Matrix Analysis and Applications.

DUFF, I. AND PRALET, S. 2005. Towards a stable static pivoting strategy for the sequential and parallel solution of sparse symmetric indefinite systems. Tech. Rep. RAL-TR-2005-007, Rutherford Appleton Laboratory.

DUFF, I. AND REID, J. 1983. The multifrontal solution of indefinite sparse symmetric linear systems. *ACM Trans. Math. Softw. 9*, 302–325.

DUFF, I. AND SCOTT, J. 2005. Towards an automatic ordering for a symmetric sparse direct solver. Tech. Rep. RAL-TR-2006-001, Rutherford Appleton Laboratory.

GEORGE, A. 1973. Nested dissection of a regular finite-element mesh. *SIAM J. Numer. Anal. 10*, 345–363.

GEORGE, A. AND LIU, J.-H. 1981. *Computer Solution of Large Sparse Positive Definite Systems*. Prentice-Hall, New Jersey, U.S.A.

GILBERT, J., NG, E., AND PEYTON, B. 1994. An efficient algorithm to compute row and column counts for sparse Cholesky factorization. *SIAM J. Matrix Anal. Appl. 15*, 1075–1091.

GOULD, N., HU, Y., AND SCOTT, J. 2005. Complete results for a numerical evaluation of sparse direct solvers for the solution of large, sparse, symmetric linear systems of equations. Numerical Analysis Internal Report 2005-1, Rutherford Appleton Laboratory. Available from www.numerical.rl.ac.uk/reports/reports.shtml.

GOULD, N. AND SCOTT, J. 2003. Complete results for a numerical evaluation of HSL packages for the direct solution of large sparse, symmetric linear systems of equations. Numerical Analysis Internal Report 2003-2, Rutherford Appleton Laboratory. Available from www.numerical.rl.ac.uk/reports/reports.shtml.

GOULD, N. AND SCOTT, J. 2004. A numerical evaluation of HSL packages for the direct solution of large sparse, symmetric linear systems of equations. *ACM Trans. Math. Softw. 30*, 300–325.

GOULD, N. AND TOINT, P. 2002. An iterative working-set method for large-scale non-convex quadratic programming. *Applied Numer. Math. 43,* 1–2, 109–128.

GUPTA, A., JOSHI, M., AND KUMAR, V. 2001. WSMP: A high-performance serial and parallel sparse linear solver. Tech. Rep. RC 22038 (98932), IBM T.J. Watson Reserach Center. www.cs.umn.edu/~agupta/doc/wssmp-paper.ps.

GUPTA, A., KARYPIS, G., AND KUMAR, V. 1997. Highly scalable parallel algorithms for sparse matrix factorization. *IEEE Trans. on Para. Distrib. Syst. 8,* 5, 502–520.

HEATH, M., NG, E., AND PEYTON, B. 1991. Parallel algorithms for sparse linear systems. *SIAM Rev. 33*, 420–460.

HSL. 2004. A collection of Fortran codes for large-scale scientific computation. See http://www.cse.clrc.ac.uk/nag/hsl/.

KARYPIS, G. AND KUMAR, V. 1998. METIS: A software package for partitioning unstructured graphs, partitioning meshes and computing fill-reducing orderings of sparse matrices—version 4.0. See http://www-users.cs.umn.edu/~karypis/metis/.

KARYPIS, G. AND KUMAR, V. 1999. A fast and high quality multilevel scheme for partitioning irregular graphs. *SIAM J. Sci. Comput. 20*, 359–392.

LIU, J. 1985. Modification of the minimum-degree algorithm by multiple elimination. *ACM Trans. Math. Softw. 11*, 2, 141–153.

LIU, J. 1992. The multifrontal method for sparse matrix solution: theory and practice. *SIAM Revi. 34*, 82–109.

NG, E. AND PEYTON, B. 1993. Block sparse Cholesky algorithms on advanced uniprocessor computers. *SIAM J. Sci. Comput. 14*, 5, 1034–1056.

ROTKIN, V. AND TOLEDO, S. 2004. The design and implementation of a new out-of-core sparse Cholesky factorization method. *ACM Trans. Math. Softw. 30*, 19–46.

SAUNDERS, M. 1994. Sparse matrices in optimization. Presented at Sparse Days at St Girons, International meeting on Sparse Matrix Methods, St Girons, France. See http://www.stanford.edu/group/SOL/talks/saunders-stgirons.ps.

SCHENK, O. AND GÄRTNER, K. 2004a. On fast factorization pivoting methods for sparse symmetric indefinite systems. Tech. Rep. CS-2004-004, Department of Computer Science, University of Basel, Switzerland.

SCHENK, O. AND GÄRTNER, K. 2004b. Solving unsymmetric sparse systems of linear equations with PARDISO. *J. Future Gen. Compu. Syst. 20*, 3, 475–487.

SCHENK, O., GÄRTNER, K., AND FICHTNER, W. 2000. Efficient sparse LU factorization with left-right looking strategy on shared memory multiprocessors. *BIT 40*, 1, 158–176.

SCHULZE, J. 2001. Towards a tighter coupling of bottom-up and top-down sparse matrix ordering methods. *BIT 41*, 800–841.

TINNEY, W. AND WALKER, J. 1967. Direct solutions of sparse network equations by optimally ordered triangular factorization. *Proc. IEEE 55*, 1801–1809.

Received May 2005; revised November 2005, March 2006; accepted May 2006