A brief bibliography of recent research and software for the parallel solution of large sparse linear equations

I. S. Duff*

Rutherford Appleton Laboratory, Chilton, Didcot, Oxon OX11 0QX, England.
Also at CERFACS, Toulouse, France.
email:isd@rl.ac.uk

Abstract

We give some pointers to recent work on the parallel solution of sparse linear equations. We consider both iterative and direct methods and combinations of these two approaches and list some Web sites where software may be found.

1. Introduction

After a discussion with the organizers and participants at the Cattle Breeding Conference in Helsinki, I decided to give pointers to recent work by myself and others rather than writing a review article. Hopefully my experience at the meeting will help me in choosing an appropriate selection of references.

A wide range of iterative, direct, and preconditioning techniques with an emphasis on the exploitation of parallelism is considered at length in the recent book by Dongarra, Duff, Sorensen, and van der Vorst [5]. A more recent bibliographic tour is presented by Duff and van der Vorst [7].

2. Iterative methods

There have been several books published recently on the iterative solution of large sparse linear systems although most are aimed at practising mathematicians rather than applications scientists. Perhaps one of the most accessible is the book by Saad [13].

Several projects have developed software that implements iterative methods for parallel computers. These include:

- AZTEC. For more information see: `http://www.cs.sandia.gov/HPCCIT/aztec.html`
- ITPACK. For the parallel implementation of this package see: [10].
- PARASOL. For more information see Section 5.
- PETSc. For more information see: `http://www.mcs.anl.gov/petsc/`
- PIM. Information for this parallel software package can be obtained from: `ftp://unix.hensa.ac.uk/pub/misc/netlib/pim/`
- PINEAPL. For more information see: `http://www.nag.co.uk/projects/PINEAPL/`
- SPARSKIT and SPARSELIB. For more information see: `http://www/cs.umn.edu/~saad/`

[*] Technical reports available through Web site: http://www.numerical.rl.ac.uk/reports/reports.html

3. Direct methods

The two main books on the direct solution of sparse linear equations are the book by George and Liu [8] and the one by Duff, Erisman, and Reid [6]. The former restricts its discussion to symmetric positive definite systems and emphasizes graph theoretic aspects, while the latter considers both symmetric and unsymmetric systems and includes a discussion of some of the algorithms used in the Harwell Subroutine Library (HSL) [9]. The HSL has the largest number of direct sparse codes in a single library and has a few codes for iterative solution also. A subset of HSL is marketed by NAG as the Harwell Sparse Matrix Library (HSML). Other sparse direct software can be found through the netlib repository http://www.netlib.org.

4. Solution of least-squares problems

As I mentioned in my presentation at the meeting, I feel something of a heretic to even dare suggest that the breeding equations could be solved by other than using the mixed model equations. However, the original problem is represented by a set of overdetermined equations and some variance and covariance matrices and the MME is only one way of solving them corresponding to the normal equations weighted by the variances and covariances. These normal equations are, of course symmetric (although not always definite) and are commonly solved in the cattle breeding community by combinations of block Jacobi and Gauss-Seidel, although conjugate gradients preconditioned by block Jacobi is also now used.

There are, however, other methods for solving the least-squares problem. The most robust uses a QR factorization of the coefficient matrix and has been used by Amestoy, Duff, and Puglisi [1] in the solution of problems in pig breeding. However, to use such techniques in the solution of problems of the dimension discussed at the meeting, some partitioning of the matrix would be required.

The iterative method of Paige and Saunders, LSQR, [11] is, like most iterative techniques for square matrices, based on Krylov subspaces and, like these methods, will work far better with a good preconditioner.

5. Methods coupling direct and iterative techniques

In my opinion, the most promising set of techniques for the solution of problems of the size encountered by the cattle breeding community are those that combine both direct and iterative methods. This is in effect a sophisticated preconditioning for the iterative method. Indeed, one of the approaches discussed by the HPBREEDING project at the meeting does precisely this, namely solving subblocks using a Cholesky factorization and using this as a block Jacobi preconditioning for the conjugate gradient algorithm. In fact, this technique is very similar to methods used in the solution of problems from discretizations of PDEs using domain decomposition which can be viewed as permuting the matrix to bordered block diagonal form. In other applications, the use of overlapping blocks has been found to significantly improve the convergence [12]. In addition, further preconditioners are used both for the Schur complement and also a coarse preconditioner for the overall problem. A good discussion of these preconditioners can be found in the thesis of Luiz Carvalho [4]. I think that it could be beneficial for the cattle breeding community to investigate the effect of such extra preconditioning on their current solution techniques.

Block projection methods fall into this category of solution techniques and can be used either to solve the MME or the original overdetermined systems. One of the simplest of these (and the most parallelizable) is the block Cimmino algorithm that can be accelerated using conjugate gradients [2, 3].

The PARASOL Project is an ESPRIT IV Long Term Research Project (No 20160) for ``An Integrated Environment for Parallel Sparse Matrix Solvers''. In this Project new techniques in direct and iterative techniques have been developed and are included as separate solvers or as combinations within the framework of a single Library. The main goal of this Project, which started on January 1 1996, is to build and test a portable library for solving large sparse systems of equations on distributed memory systems. There are twelve partners in five countries, five of whom are code developers, five end users, and two software houses. The software is written in Fortran 90 and uses MPI for message passing. The solvers being developed in this consortium are: two domain decomposition codes by Bergen and ONERA, a multigrid code by GMD, and a parallel multifrontal method (called MUMPS) by CERFACS and RAL. The final library will be in the public domain. For more information on the PARASOL project, see the web site at http://www.genias.de/parasol.

6. Final remarks

Although the problems described at the Cattle Breeding Conference are of very high dimension, they are very sparse and are comparable with the size of some of the larger problems solved when discretizing partial differential equations. For example, some problems from the ASCI program in the United States have dimensions of several million and are often significantly denser than the problems encountered in cattle breeding. I feel that it is certainly worthwhile to contemplate other formulations than the mixed model equations but perhaps in the short term, to avoid too much of a cultural change, some of the current techniques with additional preconditioning might be worth pursuing.

References

[1] Patrick R. Amestoy, Iain S. Duff, and Chiara Puglisi. Multifrontal QR factorization in a multiprocessor environment. Numerical Linear Algebra with Applications, 3(4):275-300, 1996.

[2] Mario Arioli, Iain S. Duff, Joseph Noailles, and Daniel Ruiz. A block projection method for sparse matrices. SIAM J. Scientific and Statistical Computing, 13:47--70, 1992.

[3] Mario Arioli, Iain S. Duff, Daniel Ruiz, and Miloud Sadkane. Block Lanczos techniques for accelerating the Block Cimmino method. SIAM J. Scientific Computing, 16(6):1478--1511, 1995.

[4] Luiz Mariano Carvalho. Preconditioned Schur complement methods in distributed memory environments. INPT PhD Thesis TH/PA/97/41, CERFACS, Toulouse, France, 1997.

[5] Jack J. Dongarra, Iain S. Duff, Danny C. Sorensen, and Henk A. van der Vorst. Numerical Linear Algebra for High-Performance Computers. SIAM Press, Philadelphia, 1998.

[6] Iain S. Duff, A. M. Erisman, and John K. Reid. Direct Methods for Sparse Matrices. Oxford University Press, Oxford, England, 1986.

[7] Iain S. Duff and Henk A. van der Vorst. Parallel hoojy. Technical Report RAL-TR-1999-027, Rutherford Appleton Laboratory, 1999.

[8] A. George and J. W. H. Liu. Computer Solution of Large Sparse Positive Definite Systems. Englewood Cliffs, New Jersey: Prentice-Hall, 1981.

[9] HSL. Harwell Subroutine Library. A Catalogue of Subroutines (Release 12). AEA Technology, Harwell Laboratory, Oxfordshire, England, 1996. For information concerning HSL contact: Dr Scott Roberts, AEA Technology, 477 Harwell, Didcot, Oxon OX11 0RA, England (tel: +44-1235-432682, fax: +44-1235-432023, email: Scott.Roberts@aeat.co.uk).

[10] D. R. Kincaid and T. C. Oppe. Recent vectorization and parallelization of ITPACKV. In O. Axelsson and L. Yu. Kolotilina, editors, Preconditioned Conjugate Gradient Methods, pages 58--74, Berlin, 1990. Nijmegen 1989, Springer Verlag. Lecture Notes in Mathematics 1457.

[11] C. C. Paige and M. A. Saunders. Lsqr: an algorithm for sparse linear equations and sparse least squares. ACM Trans. Math. Softw., 8:43--71, 1982.

[12] G. Radicati di Brozolo and Y. Robert. Parallel conjugate gradient-like algorithms for solving sparse non-symmetric systems on a vector multiprocessor. Parallel Computing, 11:223--239, 1989.

[13] Y. Saad. Iterative methods for sparse linear systems. PWS Publishing, New York, NY, 1996.