

On some parallel banded system solvers

Jack J. DONGARRA *

Mathematics and Computer Science Division, Argonne National Laboratory, Argonne, IL 60439, U.S.A.

Ahmed H. SAMEH **

Department of Computer Science, University of Illinois, Urbana Champaign, Urbana, IL 61801, U.S.A.

Received March 1984

Abstract. This paper describes algorithms for solving narrow banded systems and the Helmholtz difference equations that are suitable for multiprocessing systems. The organization of the algorithms highlight the large grain parallelism inherent in the problems.

Keywords. Banded systems, parallelism, multiprocessor, partitioning.

1. Introduction

We consider algorithms for solving narrow-banded diagonally dominant linear systems which are suitable for multiprocessors. We describe a direct solver similar to that in [12] for tridiagonal systems, and in [9] for solving a banded system on a linearly connected set of processors. We will also provide and analyze a parallel implementation of the partitioning algorithm and the matrix decomposition which we refer to as a hybrid solver (direct and iterative) which is superior to the direct solver especially for strongly diagonally dominant systems. When the interconnection network is not sufficiently powerful, a bottleneck develops in one of the stages of the direct solver in which the cost of the computation is proportional to the number of processors. This inefficiency may be alleviated by using an iterative scheme in this stage that takes full advantage of the parallelism offered even by a linear array of p processors.

A similar approach is also used to handle the positive-definite system that arises from the standard five-point finite-difference discretization of the Helmholtz equation. This problem arises frequently in situations where fast solvers are of primary importance. In this paper we consider the matrix decomposition solver that has been described in several papers, e.g. [1,2,6,10,11].

2. A partitioning algorithm for banded systems

Let the linear system under consideration be denoted by

$$Ax = f \tag{1}$$

where A is a banded diagonally dominant matrix of order n . We assume that the number of superdiagonals $m \ll n$ is equal to the number of subdiagonals. On a sequential machine such a system would be solved via Gaussian elimination without pivoting at a cost of $O(m^2n)$ arithmetic operations. We describe here an algorithm for solving this system on a multiprocessor of p processing units. Each unit may be a sequential machine, a vector machine, or an

* Work supported in part by the Applied Mathematical Sciences Research Program (KC-04-02) of the Office of Energy Research of the U.S. Department of Energy under contract W-31-109-Eng-38.

** Work supported in part by the National Science Foundation under grant US NSF MCS 81-17010.

array of processors. In this paper, however, we consider only p sequential processing units.

Let the system (1) be partitioned into the block-tridiagonal form shown below

$$\begin{pmatrix} A_1 & B_1 & & & & & \\ C_2 & A_2 & B_2 & & & & \\ & \cdot & \cdot & \cdot & & & \\ & & & & C_{p-1} & A_{p-1} & B_{p-1} \\ & & & & & C_p & A_p \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \\ \vdots \\ x_{n-1} \\ x_n \end{pmatrix} = \begin{pmatrix} f_1 \\ f_2 \\ \vdots \\ f_{n-1} \\ f_n \end{pmatrix} \quad (2)$$

where A_i , $1 \leq i \leq p-1$, is a banded matrix of order $q = \lceil n/p \rceil$ and bandwidth $2m+1$ (same as A),

$$B_i = \begin{pmatrix} 0 & 0 \\ \hat{B}_i & 0 \end{pmatrix}, \quad 1 \leq i \leq p-1 \quad (3a)$$

and

$$C_{i+1} = \begin{pmatrix} 0 & \hat{C}_{i+1} \\ 0 & 0 \end{pmatrix}, \quad (3b)$$

in which \hat{B}_i and \hat{C}_{i+1} are lower and upper triangular matrices, respectively, each of order m . The algorithm consists of four stages.

2.1. Stage 1

Obtain the LU-factorization

$$A_i = L_i U_i, \quad 1 \leq i \leq p \quad (4)$$

using Gaussian elimination without pivoting, one processor per factorization. Here L_i is unit lower triangular and U_i is a non-singular upper triangular matrix. Note that each A_i is also diagonally dominant.

The cost of this stage is $O(m^2 n/p)$ arithmetic operations, no inter-processor communication is required.

2.2. Stage 2

If we premultiply both sides of (2) by

$$\text{diag}(A_1^{-1}, A_2^{-1}, \dots, A_p^{-1})$$

we obtain a system of the form

$$\begin{pmatrix} I_q & E_1 & & & & & \\ F_2 & I_q & E_2 & & & & \\ & \cdot & \cdot & \cdot & & & \\ & & & & F_{p-1} & I_q & E_{p-1} \\ & & & & & F_p & I_q \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \\ \vdots \\ x_{n-1} \\ x_n \end{pmatrix} = \begin{pmatrix} g_1 \\ g_2 \\ \vdots \\ g_{n-1} \\ g_n \end{pmatrix} \quad (5)$$

where

$$E_i = (\hat{E}_i, 0), \quad F_i = (0, \hat{F}_i),$$

in which \hat{E}_i and \hat{F}_i are matrices of m columns given by

$$\hat{E}_i = A_i^{-1} \begin{pmatrix} 0 \\ \hat{B}_i \end{pmatrix} \quad \text{and} \quad \hat{F}_i = A_i^{-1} \begin{pmatrix} \hat{C}_i \\ 0 \end{pmatrix}$$

and will in general be full. In other words \hat{E}_i , \hat{F}_i , and g_i are obtained by solving the linear

(6)

The cost of the algorithm to be used for solving (6) depends on the interconnection network. Processor 1 contains T_1 and h_1 , processor j , $2 \leq j \leq p-1$, contains P_j , Q_j , S_j , T_j , and h_{2j-2} , h_{2j-1} , and processor p contains P_p , and h_{2p-2} . Hence, if the processors are linearly connected we can only solve (6) sequentially at the cost of $O(m^2p)$ steps, where a step is the cost of an arithmetic operation or the cost of transmitting a floating-point number from one processor to either of its immediate neighbors. We should add here that since A is diagonally dominant it can be shown that (6) is also diagonally dominant and hence can be solved via Gaussian elimination without pivoting.

2.4. Stage 4

Once the y_i 's are obtained, with y_1 in processor 1, y_{2j-2} and y_{2j-1} in processor j ($2 \leq j \leq p-1$), and y_{2p-2} in processor p , the rest of the components of the solution vector of (5) may be computed as follows. Processor k , $1 \leq k \leq p$, evaluates

$$z_k = w_k - M_k y_{2k-3} - N_k y_{2k} \quad (7)$$

with processors 1 and p performing the additional tasks

$$y_0 = h_0 - S_1 y_2 \quad \text{and} \quad y_{2p-1} = h_{2p-1} - Q_p y_{2p-3}, \quad (8)$$

respectively (M_1 and N_p are non-existence and are taken to be zero in this equation). The cost of this stage is $O(mn/p)$ steps, with no inter-processor communication.

It can be shown that for a linear array of processors, the speedup of this algorithm over the classical sequential algorithm behaves as shown in Fig. 1 where p_0 and σ_0 are $O(\sqrt{nm})$. Stage 2 dominates the computation until p_0 , then the communication costs impact the performance and Stage 3 has a greater influence.

For a linear array of processors, the bottleneck in this parallel algorithm is the process of solving the reduced system (Stage 3). It is the only stage whose cost increases as p becomes large. This inefficiency may be alleviated by solving the reduced system (6) using an iterative scheme that takes full advantage of the parallelism offered even by a linear array of p processors. Since the reduced system is diagonally dominant, the simplest iterative scheme that is suited for such a linear array of processors is the block-Jacobi algorithm which we outline below.

Let

$$G_i = \begin{pmatrix} I_m & T_i \\ P_{i+1} & I_m \end{pmatrix}, \quad i = 1, 2, \dots, p-1, \quad G = \text{diag}(G_1, G_2, \dots, G_{p-1}), \quad (9)$$

and let H be the block-tridiagonal matrix

$$H = \left[\begin{pmatrix} -Q_j & 0 \\ 0 & 0 \end{pmatrix}, 0, \begin{pmatrix} 0 & 0 \\ 0 & -S_{j+1} \end{pmatrix} \right],$$

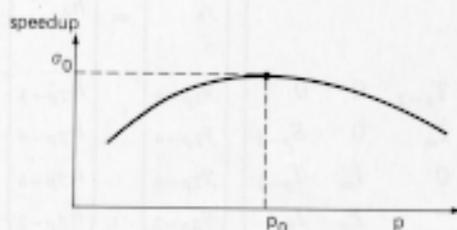


Fig. 1.

$$\begin{pmatrix} N & E & & & & \\ E^T & N & E & & & \\ & & & \ddots & & \\ & & & & E^T & N & E \\ & & & & & E^T & N \end{pmatrix} \begin{pmatrix} \hat{u}_1 \\ \hat{u}_2 \\ \vdots \\ \hat{u}_{p-1} \\ u_p \end{pmatrix} = \begin{pmatrix} \hat{f}_1 \\ \hat{f}_2 \\ \vdots \\ \hat{f}_{p-1} \\ \hat{f}_p \end{pmatrix}, \quad (13)$$

where

$$E = \begin{pmatrix} 0 & 0 \\ -I_n & 0 \end{pmatrix}$$

and $N = [-I_n, T, -I_n]$ are of order qn ,

$$\hat{u}_i = (u_{(i-1)q+1}^T, \dots, u_{iq}^T)^T \quad \text{and} \quad \hat{f}_i = (f_{(i-1)q+1}^T, \dots, f_{iq}^T)^T, \quad 1 \leq i \leq p.$$

If the multiprocessor consist of p processing units, the matrix decomposition algorithm (MD-algorithm) may be organized as follows.

3.1. Stage 1

Each matrix T has the spectral decomposition

$$T = Q\Lambda Q$$

where Q is an orthogonal matrix whose elements are given by

$$e_i^T Q e_j = \sqrt{2/(n+1)} \sin(ij\pi/(n+1)), \quad 1 \leq i, j \leq n$$

and $\Lambda = \text{diag}(\lambda_1, \dots, \lambda_n)$, with

$$\lambda_j = (2 + \alpha^2) + 4 \sin^2(j\pi/(n+1)), \quad 1 \leq j \leq n.$$

Now, (13) is reduced to the form

$$\begin{pmatrix} M & E & & & & \\ E^T & M & E & & & \\ & & & \ddots & & \\ & & & & E^T & M & E \\ & & & & & E^T & M \end{pmatrix} \begin{pmatrix} \hat{v}_1 \\ \hat{v}_2 \\ \vdots \\ \hat{v}_{p-1} \\ \hat{v}_p \end{pmatrix} = \begin{pmatrix} \hat{g}_1 \\ \hat{g}_2 \\ \vdots \\ \hat{g}_{p-1} \\ \hat{g}_p \end{pmatrix}, \quad (14)$$

where $M = [-I_n, \Lambda, -I_n]$ is of order qn .

$$\hat{v}_i = \hat{Q}\hat{u}_i, \quad \text{and} \quad \hat{g}_i = \hat{Q}\hat{f}_i$$

in which $\hat{Q} = \text{diag}(Q, \dots, Q)$ is of order qn . Here, the i th processing unit performs the q sine transforms

$$g_{(i-1)q+k} = Qf_{(i-1)q+k}, \quad 1 \leq k \leq q, \quad (15)$$

where $\hat{g}_i = (g_{(i-1)q+1}^T, \dots, g_{iq}^T)^T, 1 \leq i \leq p$.

3.2. Stage 2

Premultiplying both sides of (14) by $\text{diag}(M^{-1}, \dots, M^{-1})$, we obtain a system of the form

$$\begin{pmatrix} I & F & & & & \\ G & I & F & & & \\ & & & \ddots & & \\ & & & & G & I & F \\ & & & & & G & I \end{pmatrix} \begin{pmatrix} \hat{v}_1 \\ \hat{v}_2 \\ \vdots \\ \hat{v}_{p-1} \\ \hat{v}_p \end{pmatrix} = \begin{pmatrix} \bar{h}_1 \\ \bar{h}_2 \\ \vdots \\ \bar{h}_{p-1} \\ \bar{h}_p \end{pmatrix}, \quad (16)$$

The cost of the arithmetic (alone) of this algorithm is $O(n \log_2 n)$ and is dominated by the sine transforms in Stages 1 and 5. If the arithmetic is overlapped with the inter-processing communication (through the global memory), the total cost remains $O(n \log_2 n)$.

4. Numerical results

The algorithms described above require more arithmetic operations than the well-known sequential algorithms currently in use for solving the problems described above. Hence, they should not be competitive on sequential machines. We have used all of our parallel schemes as sequential solvers and have obtained results of equal numerical accuracy with LINPACK [4], for solving narrow banded systems, and with FISHPACK for solving the Helmholtz equation.

PBAND consumed twice the time as the routines SGBFA and SGBSL from LINPACK on a VAX-780 for solving the system of linear equations (1) with $n = 512$, $m = 5$, and the number of processors $p = 16$. This ratio of time consumed held for three examples: $A_k = [a_{ij}^{(k)}]$, $1 \leq k \leq 3$, with $a_{ij}^{(k)}$, $i \neq j$, obtained by a random number generator, and the diagonal elements $a_{ii}^{(k)} = a_k$ given by $a_1 = 32$, $a_2 = 5$, and $a_3 = 3$. All A_k are non-singular with A_1 being the only diagonally dominant matrix. In Table 2 we compare the number of iterations required by PBAND1, 2, and 3 to achieve a residual of 2-norm less than 10^{-5} . We see, then, that if the multiprocessor possesses a reasonably powerful interconnection network, the PBAND 3 version of Orthomin (1) can be an effective scheme for solving the system when A is not diagonally dominant but the reduced system has a positive definite symmetric part.

We ran our FORTRAN program on the CRAY X-MP-4 using the multitasking features available on the machine [3]. The CRAY X-MP-4 has four processors which can be used by a single FORTRAN program. In running the band solver we using one, two, three, and four processors to solve the problem. Table 3 shows the results for PBAND3 for a positive-definite random matrix of various orders. For large problems using two processors the speedups are almost perfect. Results with three and four processors show some degradation in performance, but considering the level of granularity the results are quite impressive.

We also implemented the same program on the Denelcor HEP. The Denelcor HEP is a MIMD computer which supports tightly coupled parallel processing. The fully configured computing system offered by Denelcor consists of up to 16 processing elements (PEMs) sharing a large global memory through a crossbar switch. Within a single PEM, parallelism is achieved through pipelining independent serial instructions streams called processes. The principal pipeline that handles the numerical and logical operations consists of synchronous functional

Table 2

Algorithm	$a_1 = 32$	$a_1 = 5$	$a_1 = 3$
PBAND1	2	4	29
PBAND2	1	3	23
PBAND3	1	3	10

Table 3

n	Bandwidth	Time (seconds)	Speedup	Speedup	Speedup
		1 processor	2 processors	3 processors	4 processors
512	11	0.024	1.8	2.40	2.86
4096	15	0.256	1.98	2.86	3.75
16384	31	2.12	1.995	2.91	3.87

Table 4

n	Bandwidth	Speedup over running sequentially
512	11	6.
4096	15	8.
16384	31	Too large to run

Table 5

p	1	8	16	32
FISHPACK	12			
Helm		36	37	37

units that have been segmented into an eight-stage pipe. The HEP we used had only a single PEM, and the maximum speedup over the same process running sequentially is between eight and ten. For further details on the HEP architecture see the article [7] by H. Jordan. See also Table 4.

Similarly, we compared HELM with the cyclic reduction scheme of FISHPACK [13] for solving the Poisson equation ($\alpha = 0$) on the unit square with mesh size $1/129$, i.e. $n = 128$, for $p = 8, 16$, and 32 processors. Both schemes succeeded in producing residual of 2-norms less than 10^{-5} . Table 5 shows the time in seconds consumed by HELM and FISHPACK running on a VAX11/780 for different values of p .

The results in Table 5 reflect the fact that HELM performs redundant arithmetic operations compared to their sequential scheme in FISHPACK (cyclic reduction). One may also conclude that with an "appropriate" inter-connection of the p processors, HELM may achieve a speedup of $p/3$ over FISHPACK's sequential algorithm. True speedups achieved by these parallel algorithms over the sequential counterparts can only be measured by actual runs on a specific multiprocessor.

Acknowledgements

We would like to thank the people at CRAY Research, especially Chris Hsiung and Tom Hewitt, for helping with the runs on the CRAY X-MP-4 and Ken Hillstrom at Argonne for help with the HEP runs.

References

- [1] B. Buzbee, A fast Poisson solver amenable to parallel computation, *IEEE Trans. Comput.* **C-22** (1973) 793-796.
- [2] B. Buzbee, G. Golub and C. Nielson, On direct methods for solving Poisson's equation, *SIAM J. Numer. Anal.* **7** (1970) 627-656.
- [3] S.S. Chen, J.J. Dongarra and C.C. Hsiung, Multiprocessing linear algebra algorithms on CRAY X-MP-2: Experiences with small granularity, *J. Parallel and Distributed Comput.* **1** (1984) 22-31.
- [4] J. Dongarra, J. Bunch, C. Moler and G. Stewart, *LINPACK Users' Guide* (SIAM, Philadelphia, 1979).
- [5] H. Elman, Iterative methods for non-self-adjoint elliptic problems, *Elliptic Problem Solvers II* (Academic Press, New York, 1984) (to appear).
- [6] R. Hockney, A fast direct solution of Poisson's equation using Fourier analysis, *J. ACM* **12** (1965) 95-113.
- [7] H.F. Jordan, Experience with pipelined multiple instruction streams, *IEEE Proc.* (1984) 113-123.
- [8] T. Kailath, A. Vieira and M. Morf, Inverses of Toeplitz operators, innovations, and orthogonal polynomials, *SIAM Rev.* **20** (1978) 106-119.

- [9] D. Lawrie and A. Sameh, The computation and communication complexity of a parallel banded system solver, to appear in *ACM Trans. Math. Software* (1985).
- [10] A. Sameh, A fast Poisson solver for multiprocessors, *Elliptic Problems Solvers II* (Academic Press, New York, 1984) (to appear).
- [11] A. Sameh, S. Chen and D. Kuck, Parallel Poisson and biharmonic solvers, *Computing* **17** (1976) 219-230.
- [12] P. Swarztrauber and R. Sweet, Algorithm 541: Efficient Fortran subprograms for the solution of separable elliptic partial differential equations, *ACM Trans. Math. Software* **5** (1979) 352-364.
- [13] P. Vinsome, ORTHOMIN, an iterative method for solving sparse sets of simultaneous linear equations, *Proc. 4th Symposium on Reservoir Simulation*, Society of Petroleum Engineers of the American Institute of Mechanical Engineering (1976) 149-159.