

A Survey of Parallel Direct Methods for Block Bidiagonal Linear Systems on Distributed Memory Computers

P. AMODIO*

Dipartimento di Matematica, Università di Bari
Via Orabona 4, 70125 Bari, Italy
00110570@vm.csata.it

M. PAPRZYCKI

Department of Mathematics and Computer Science
University of Texas of the Permian Basin, Odessa, TX 79762, U.S.A.
paprzycki_m@gusher.pb.utexas.edu

T. POLTI

Dipartimento di Matematica, Politecnico di Bari
Via Orabona 4, 70125 Bari, Italy
pptt@max.uniba.it

(Received February 1995; accepted May 1995)

Abstract—Four parallel algorithms for the solution of block bidiagonal linear systems on distributed memory computers are presented. All the algorithms belong to the class of direct methods. The first is a variant of the sequential algorithm and is suitable for a small number of processors. The remaining three algorithms are based on the parallel methods for banded systems and are much better suited for parallel computations on multiple processors. The arithmetical complexity functions of the proposed algorithms are derived. The results of experiments with the four algorithms implemented in Parallel Fortran on a linear array of 32 Transputers are presented and discussed.

Keywords—Block bidiagonal systems, Parallel algorithms.

1. INTRODUCTION

Matrices with exploitable sparse block structure arise in many applications. For example, several numerical methods for the solution of ODEs, PDEs and BVPs lead to block matrices with only few non-null diagonals [1–3].

We consider the parallel solution of a linear system

$$Ax = b \quad (1)$$

where the coefficient matrix A has a block bidiagonal structure

$$\begin{pmatrix} D_1 & & & & \\ C_2 & D_2 & & & \\ & \ddots & \ddots & & \\ & & & C_n & D_n \end{pmatrix} \quad (2)$$

We would like to thank the anonymous referees for their useful comments leading to a number of improvements in the paper. We would also like to express our thanks to K. Paprzycka for her help with the English.

*Work supported by the Ministero della Ricerca Scientifica, 60% project.

size m followed by a block of zeroes of size $(k-2)m \times m$, and e_{k-1} consists of a block of zeroes of size $(k-2)m \times m$ followed by an identity matrix of size m), and

$$A^{(i)} = \begin{pmatrix} D_{(i-1)k+1} & & & & \\ C_{(i-1)k+2} & D_{(i-1)k+2} & & & \\ & & \ddots & & \\ & & & \ddots & \\ & & & & C_{ik-1} & D_{ik-1} \end{pmatrix}.$$

Except for the last, each processor stores 2 block rows of the partitioning of A (e.g., processor i contains the block rows with $A^{(i)}$ and D_{ik} , and the last processor contains the block row with $A^{(p)}$). Section 2 presents a modification of the sequential algorithm that is suitable for a parallel computer with a small number of processors. Sections 3 and 4 introduce the two algorithms based on different parallel factorizations. Section 5 is devoted to the generalized cyclic reduction algorithm. Section 6 contains the study of the arithmetical complexity functions and the memory requirements of the proposed algorithms. In Section 7 the results of experiments on a network of 32 transputers are presented and discussed.

2. QUASI-SEQUENTIAL ALGORITHM

The first algorithm we consider is a simple modification of the sequential one. First, the LU decomposition is used in parallel to invert the main diagonal blocks D_i (this is the most expensive part of the sequential algorithm). Then the algorithm continues in the following way: the first processor solves its part of the system and sends the last block component of the solution to the second processor. The second processor, waiting for the data from the first, scales its first q_2 block equations (where q_2 is a positive integer number less than or equal to k). When it receives the data, it updates the first q_2 block equations, solves the remaining equations and sends its last block of the solution to the next processor. Each processor j performs the same operations as the second processor; waiting for a vector from processor $j-1$, it scales the first q_j block equations (obviously the value of q_j must be proportional to j , see Sections 6 and 7), then it solves the equations and sends the data to processor $j+1$. The algorithm for processor j can be thus summarized as follows:

```

for  $i = (j-1)k + 1, jk$ 
    determine  $P_i, L_i$  and  $U_i$  such that  $P_i L_i U_i = D_i$ 
end
for  $i = (j-1)k + 1, (j-1)k + q_j$ 
     $E_i = D_i^{-1} C_i$ 
     $g_i = D_i^{-1} b_i$ 
end
receive  $x_{(j-1)k}$  from processor  $j-1$ 
for  $i = (j-1)k + 1, (j-1)k + q_j$ 
     $x_i = g_i - E_i x_{i-1}$ 
end
for  $i = (j-1)k + q_j + 1, jk$ 
     $g_i = b_i - C_i x_{i-1}$ 
     $x_i = D_i^{-1} g_i$ 
end
send  $x_{jk}$  to processor  $j+1$ 

```

This algorithm is characterized by sequential communication between processors (each processor waits for data from another processor before sending its data to the next one). If the number of processors is relatively large in comparison to the block sizes, this does not make it very efficient (see Sections 6 and 7). At the same time, the algorithm has some advantages: it does

not produce fill-in vectors and requires only vector transmissions. If a very large system is to be solved on a machine with a large number of processors (and if each of these processors has a substantial local memory), it is possible to improve the algorithm by gradually increasing the number of blocks stored in later processors.

3. REDUCTION ALGORITHM

The second algorithm is similar to the approach used by Brugnano in [9] to solve tridiagonal linear systems and by Ascher and Chan in [6] to solve ABD systems. This algorithm leads to the solution of a reduced block system of size p in which the unknowns are located in the first block component of the solution in each of the processors. It may be considered as a variant of the sequential algorithm in the sense that we calculate the first block of the vector solution in each processor in order to solve the remaining system using the sequential solver. First we scale the matrix A in order to obtain identity matrices on the main block diagonal. Then supposing we know the first block component of the solution in each processor j , we can express, by using the sequential algorithm iteratively, each vector \mathbf{x}_{jk+i} , for $i = 2, \dots, k$, as a function of \mathbf{x}_{jk+1} , for $j = 0, \dots, p-1$. Hence, always using the same recursion, we relate \mathbf{x}_{jk+1} to $\mathbf{x}_{(j-1)k+1}$, for $j = 1, \dots, p-1$, obtaining a block bidiagonal linear system of size p with unitary block main diagonal

$$\begin{pmatrix} I_1 & & & & & & \\ T_2 & I_{k+1} & & & & & \\ & T_3 & I_{2k+1} & & & & \\ & & & \ddots & & & \\ & & & & T_p & I_{(p-1)k+1} & \end{pmatrix} \begin{pmatrix} \mathbf{x}_1 \\ \mathbf{x}_{k+1} \\ \mathbf{x}_{2k+1} \\ \vdots \\ \mathbf{x}_{(p-1)k+1} \end{pmatrix} = \begin{pmatrix} \mathbf{d}_1 \\ \mathbf{d}_2 \\ \mathbf{d}_3 \\ \vdots \\ \mathbf{d}_p \end{pmatrix}, \quad (3)$$

where I_j represents the identity matrix of order m_j ,

$$T_j = (-1)^{k+1} \prod_{i=1}^k E_{(j-2)k+2-i}, \quad \text{for } j = 2, \dots, p,$$

$$\mathbf{d}_1 = \mathbf{g}_1,$$

$$\mathbf{d}_j = \mathbf{g}_{(j-1)k+1} - \sum_{i=(j-2)k+2}^{(j-1)k} \left(\prod_{l=0}^{(j-1)k-i} (-E_{(j-1)k-l+1}) \right) \mathbf{g}_i, \quad \text{for } j = 2, \dots, p,$$

E_i are the subdiagonal blocks of the scaled matrix, and \mathbf{g}_i are the corresponding right-hand side vectors.

After the reduced system has been solved, the other components of the solution are easily obtained in parallel by using the sequential algorithm. The following summarizes the algorithm for a generic processor j :

```

for  $i = (j-1)k+1, jk$ 
  determine  $P_i, L_i$  and  $U_i$  such that  $P_i L_i U_i = D_i$ 
   $E_i = D_i^{-1} C_i$ 
   $\mathbf{g}_i = D_i^{-1} \mathbf{b}_i$ 
end
 $T_j = E_{(j-1)k+2}$ 
 $\mathbf{d}_j = \mathbf{g}_{(j-1)k+2}$ 
for  $i = (j-1)k+3, jk$ 
   $T_j = -E_i T_j$ 
   $\mathbf{d}_j = \mathbf{g}_i - E_i \mathbf{d}_j$ 
end
send  $T_j$  and  $\mathbf{d}_j$  to processor  $j+1$ 

```


The unknown elements of N and Q are obtained from (5) by direct identification:

$$\begin{aligned} \mathbf{v}^{(j)} &= (A^{(j)})^{-1} \mathbf{e}_1 C_{(j-1)k+1}, & \text{for } j = 2, \dots, p, \\ T_j &= -D_{jk}^{-1} C_{jk} \mathbf{e}_{k-1}^T \mathbf{v}^{(j)}, & \text{for } j = 2, \dots, p-1. \end{aligned} \quad (6)$$

The factorization (5), the solution of N , and the updating of the right-hand side by $\mathbf{v}^{(j)}$ are performed in parallel with no data transmission between processors. Data transmissions are only required when solving the reduced linear system

$$\begin{pmatrix} I_k & & & & & \\ T_2 & I_{2k} & & & & \\ & T_3 & I_{3k} & & & \\ & & \ddots & \ddots & & \\ & & & T_{p-1} & I_{(p-1)k} & \end{pmatrix} \begin{pmatrix} \mathbf{x}_k \\ \mathbf{x}_{2k} \\ \mathbf{x}_{3k} \\ \vdots \\ \mathbf{x}_{(p-1)k} \end{pmatrix} = \begin{pmatrix} \mathbf{g}_k \\ \mathbf{g}_{2k} \\ \mathbf{g}_{3k} \\ \vdots \\ \mathbf{g}_{(p-1)k} \end{pmatrix}. \quad (7)$$

The following is the algorithm for the generic processor j :

```

for  $i = (j-1)k+1, jk$ 
  determine  $P_i, L_i$  and  $U_i$  such that  $P_i L_i U_i = D_i$ 
end
 $V_{(j-1)k+1} = D_{(j-1)k+1}^{-1} C_{(j-1)k+1}$ 
for  $i = (j-1)k+2, jk-1$ 
   $V_i = D_i^{-1} C_i V_{i-1}$ 
end
 $T_j = -D_{jk}^{-1} C_{jk} V_{jk-1}$ 
 $\mathbf{g}_{(j-1)k+1} = D_{(j-1)k+1}^{-1} \mathbf{b}_{(j-1)k+1}$ 
for  $i = (j-1)k+2, jk$ 
   $\mathbf{b}_i = \mathbf{b}_i - C_i \mathbf{g}_{i-1}$ 
   $\mathbf{g}_i = D_i^{-1} \mathbf{b}_i$ 
end
obtain  $\mathbf{x}_{(j-1)k}$  and  $\mathbf{x}_{jk}$  by solving the reduced system
for  $i = (j-1)k+1, jk-1$ 
   $\mathbf{x}_i = \mathbf{g}_i - V_i \mathbf{x}_{(j-1)k}$ 
end

```

The algorithm for the solution of the reduced system created by the parallel factorization is quite similar to that seen in the previous section for the solution of the system created by the reduction algorithm. The main difference is that the parallel factorization algorithm gives a reduced system of smaller size (the last processor does not work to solve it). Moreover, since each processor j also needs to know vector $\mathbf{x}_{(j-1)k}$ at the end of the solution of the reduced system one additional vector transmission is required. The algorithm for the processor j is thus:

```

for  $i = 0, \lceil \log_2(p-1) \rceil - 1$ 
  send  $\mathbf{b}_{jk}$  and  $V_{jk}$  to processor  $j+2^i$ 
  receive  $\mathbf{b}_{(j-2^i)k}$  and  $V_{(j-2^i)k}$  from processor  $j-2^i$ 
   $\mathbf{b}_{jk} = \mathbf{b}_{jk} - V_{jk} \mathbf{b}_{(j-2^i)k}$ 
   $V_{jk} = -V_{jk} V_{(j-2^i)k}$ 
end
send  $\mathbf{x}_{jk}$  to processor  $j+1$ 
receive  $\mathbf{x}_{(j-1)k}$  from processor  $j-1$ 

```

Figure 2 shows the data communications required for the solution of the reduced system on $p = 10$ processors.

where I is a block identity matrix and

$$A_1 = \begin{pmatrix} I_2 & \\ \hat{C}_4 & I_4 & \\ & & \dots & & \dots & & & & & & & & & & & & & & & & & \\ & & & & \hat{C}_k & & I_k & & & & & & & & & & & & & & & \\ & & & & & & \hat{C}_{k+2} & I_{k+2} & & & & & & & & & & & & & & \\ & & & & & & & & & \dots & & & & & & & & & & & & \dots \end{pmatrix}.$$

In a similar way,

$$QPLP^T Q^T = \begin{pmatrix} I & \\ L_1 & I \end{pmatrix}, \quad \text{and} \quad QPUP^T Q^T = \begin{pmatrix} I & U_1 \\ & I \end{pmatrix},$$

where the block identity matrices on the first and the second row of M , L and U have, respectively, $p \lfloor k/2 \rfloor$ and $p \lceil k/2 \rceil$ elements.

Observe that if k is even, then this step corresponds to the first step of the classical cyclic reduction applied to the whole matrix. The reduction process is repeated by each processor on successive block matrices with block bidiagonal submatrices (see Figure 3) until (after $\lceil \log_2 k \rceil$ steps) a reduced system of size $p - 1$ is created.




Figure 3. Reduction of a 17×17 matrix on 3 processors. Black points represent block rows involved in the reduced system.

The algorithm to obtain the solution of the problem (1) by means of the cyclic reduction algorithm, for a processor j , is thus:

```

for  $i = (j - 1)k + 1, jk$ 
  determine  $P_i$ ,  $L_i$  and  $U_i$  such that  $P_i L_i U_i = D_i$ 
   $E_i = D_i^{-1} C_i$ 
   $g_i = D_i^{-1} b_i$ 
end
 $s = 1$ 
while  $s < k$ 
  for  $i = (j - 1)k + 2s, jk$ , step  $2s$ 
     $b_i = b_i - E_i b_{i-s}$ 
     $E_i = -E_i E_{i-s}$ 
  end
  if  $i - s < jk$ 
     $b_{jk} = b_{jk} - E_{jk} b_{(j-1)k+s}$ 
     $E_{jk} = -E_{jk} E_{(j-1)k+s}$ 
  end
   $s = 2s$ 
end
obtain  $x_{(j-1)k}$  and  $x_{jk}$  by solving the reduced system
while  $s \geq 1$ 
  for  $i = (j - 1)k + s, jk - 1$ , step  $2s$ 

```

```

       $\mathbf{x}_i = \mathbf{g}_i - E_i \mathbf{x}_{i-s}$ 
    end
     $s = s/2$ 
  end

```

The solution of the reduced system (7) should be obtained by using the same approach on $p/2$ processors, $p/4$, and so on. Because we are interested in solving large problems (1), the cost of any algorithm for the solution of (7) is negligible. Hence, we will use the same algorithm applied in the previous section with the parallel factorization algorithm.

6. ARITHMETICAL COMPLEXITY AND MEMORY REQUIREMENTS

In this section we introduce and compare the presented solvers from the point of view of computational cost and memory requirement. For simplicity, we assume that blocks in (2) have the same dimension. Hence, the parameters that need to be taken into account when deriving the computational cost and memory requirements of the proposed algorithms are:

- m size of each block;
- p number of processors involved;
- n number of block rows of the problem;
- q_i in the quasi-sequential algorithm the number of rows which are scaled in the i^{th} processor while previous processors complete their factorization step.

We will assume that the coefficient matrix may be divided in an optimal way among the processors ($n = kp - 1$), and that $q_j = \min((j - 1)q, k)$, for $j = 1, \dots, p$.

6.1. Arithmetical Complexity

The basic operations that are performed by the algorithms may be identified as (we use BLAS based notation):

- GETRF* LU factorization with partial pivoting of a block of size m , $((2/3)m^3 - (1/2)m^2 - (1/6)m)$ operations);
- TRSM* solution of an already decomposed linear system of size m , $(2m^2 - m)$ operations);
- GEMM* matrix-matrix (block-block) multiplication, $(2m^3 - m^2)$ operations);
- GEMV* rank 2 update block-vector multiplication followed by a vector addition, $(2m^2)$ operations);
- TRAN(s)* transmission of s data elements.

Based on the above notation, the sequential solver performs the following operations:

$$[n] \text{ GETRF} + [n] \text{ TRSM} + [n - 1] \text{ GEMV}$$

and its computational cost is

$$C_s = \left(\frac{2}{3}m^3 + \frac{7}{2}m^2 - \frac{7}{6}m \right) n - 2m^2.$$

For all four presented algorithms, the only part which is completely parallelizable is the factorization of the blocks on the main diagonal. We must then follow two different paths to obtain the computational cost of the quasi-sequential algorithm and of the remaining three parallel algorithms.

In the quasi-sequential algorithm, the update of the right-hand side blocks remains sequential. The computational cost of this part is thus the sum of the computational costs on all the processors. The solution of the block linear systems is partially parallelized. If we suppose that each

processor j performs the scaling of the main diagonal blocks of the first q_j rows when the previous processors are solving their part of the system, then the operations for this phase overlap and are not counted. This means that the operations performed are

$$[k] \text{ GETRF} + \left[\sum_{j=0}^{p-1} (k - jq)_+ \right] \text{ TRSM} + [n] \text{ GEMV} + [p-1] \text{ TRAN}(m),$$

where $(x)_+ = x$ if $x \geq 0$ and $(x)_+ = 0$ otherwise, and $q \leq \lfloor (4m-1)/(2m^2+m-1)k \rfloor$.

The total cost is thus

$$C_{qs} = \left(\frac{2}{3}m^3 - \frac{1}{2}m^2 + \frac{1}{6}m \right) \frac{(n+1)}{p} + \sum_{j=0}^{p-1} \left(2m - 1 - \frac{4m-1}{m+1}j \right)_+ \frac{m(n+1)}{p} + 2m^2n \quad (8)$$

arithmetical operations and

$$T_{qs} = (p-1)t(m)$$

transmissions, where $t(m)$ is the cost of transmission of a vector of length m .

If $p > m/2$, the summation does not depend on p , but essentially on $m^2/2$. Then formula (8) simplifies to

$$C_{qs} = \left(\frac{7}{6}m^3 - \frac{1}{2}m^2 + \frac{1}{6}m \right) \frac{(n+1)}{p} + 2m^2n.$$

The number of arithmetical operations consists of two factors: the first depends on m and n but does not depend on p ; the second one depends on $1/p$ and decreases when the number of processors increases.

For the reduction algorithm, the parallel factorization and the cyclic reduction algorithms, the operations leading to the creation of the reduced system as well as the appropriate data transmissions are completely parallelizable. The total number of operations performed by each processor will thus essentially depend on two terms that are proportional to $1/p$ and to $\lfloor \log_2 p \rfloor$.

The reduction algorithm requires the following operations in its initial phase:

$$[k] \text{ GETRF} + [k(m+1)] \text{ TRSM} + [k-1] \text{ GEMM} + [2k-2] \text{ GEMV} + [1] \text{ TRAN}(m^2+m)$$

followed by the solution of the block reduced system of size p that requires

$$\lfloor \log_2 p \rfloor [1] \text{ TRAN}(m^2+m) + [1] \text{ TRAN}(m) + \lfloor \log_2 p \rfloor [1] \text{ GEMM} + \lfloor \log_2 p \rfloor \text{ GEMV}.$$

The total number of arithmetical operations is

$$C_{ra} = \left(\frac{14}{3}m^3 + \frac{7}{2}m^2 + \frac{7}{6}m \right) \frac{(n+1)}{p} + (2m^3 + m^2) \lfloor \log_2 p \rfloor - 4m^3,$$

and the data transmission cost is

$$T_{ra} = \lfloor \log_2 p \rfloor t(m^2+m) + t(m).$$

The parallel factorization and the cyclic reduction algorithms require in their initial parallel phase the following operations:

$$[k] \text{ GETRF} + [k(m+1)] \text{ TRSM} + [k-1] \text{ GEMM} + [2k-2] \text{ GEMV}$$

followed by the solution of the reduced system of size $p-1$ which is characterized by similar arithmetical complexity function as above. The total cost of these algorithms is almost the same as that of the reduction algorithm. The arithmetical complexity is

$$C_{pf} = C_{cr} = \left(\frac{14}{3}m^3 + \frac{7}{2}m^2 + \frac{7}{6}m \right) \frac{(n+1)}{p} + (2m^3 + m^2) \lfloor \log_2(p-1) \rfloor - 4m^3,$$

while the cost of transmissions is

$$T_{pf} = T_{cr} = (\lceil \log_2(p-1) \rceil - 1)t(m^2 + m) + 2t(m).$$

Assuming (as in most common cases of computational practice) that n is much larger than m , that m is large enough so that the m^3 terms dominate lower order terms, and that the costs of data transmission are negligible, we can simplify the above formulas. In Table 1, we present the simplified formulas (the reduction, parallel factorization and cyclic reduction algorithms have the same cost and are labeled “parallel algorithms”).

Table 1. Computational costs for the proposed algorithms.

Algorithm	Computational Cost
sequential	$\frac{2}{3}m^3n$
quasi-sequential	$\frac{7}{6}m^3n/p + 2m^2n$
parallel algorithms	$\frac{14}{3}m^3n/p + 2m^3\lceil \log_2(p) \rceil$

A number of observations can be made when considering the arithmetic complexity functions of Table 1. When comparing the sequential and the quasi-sequential algorithms, for a large number of processors, the term $(7/6)m^3n/p$ disappears. Therefore, the limit on the speedup of the quasi-sequential algorithm is $m/3$ (see also Figures 4 and 5).

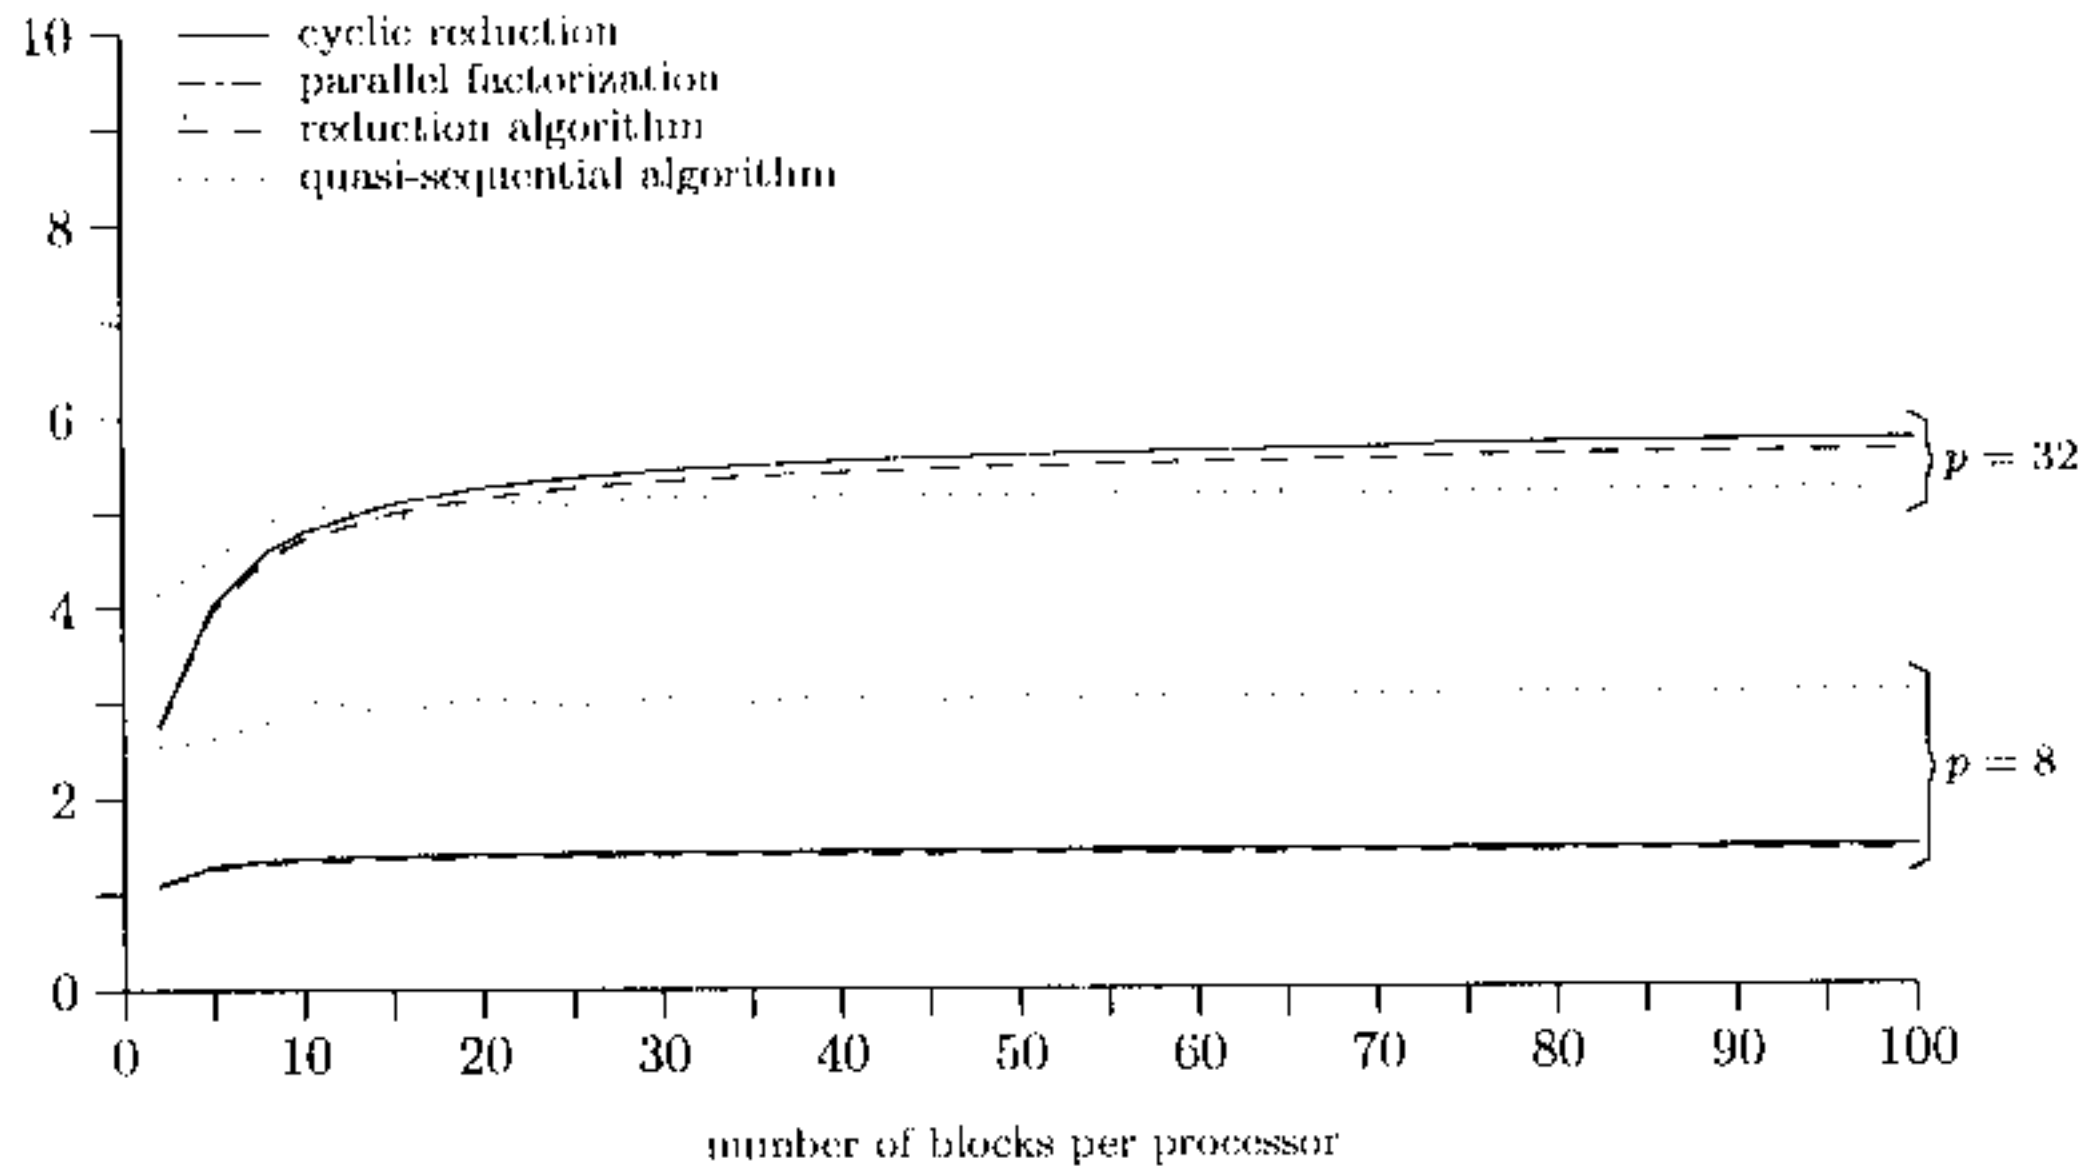


Figure 4. Speedups of the four algorithms for $m = 15$ and $p = 8, 32$.

The arithmetical complexity functions for the parallel algorithms have a structure similar to other *divide-and-conquer* algorithms [8,14] and have similar characteristics. For fixed m and p and for large n , further increase in the value of n does not lead to additional speedup. Similarly (excluding the communication cost), for fixed n and p , an increase in m will not produce further speedup increase.

Finally, given the values of n , m and p , we can estimate (see Figure 6) when the quasi sequential algorithm should be used instead of either of the parallel algorithms. It can be observed that for a given number of processors p , there exists a value of m starting from which the quasi-sequential algorithm outperforms the parallel algorithms. It can also be observed that if the data transmission costs are taken into consideration (and the cost of one transmission of

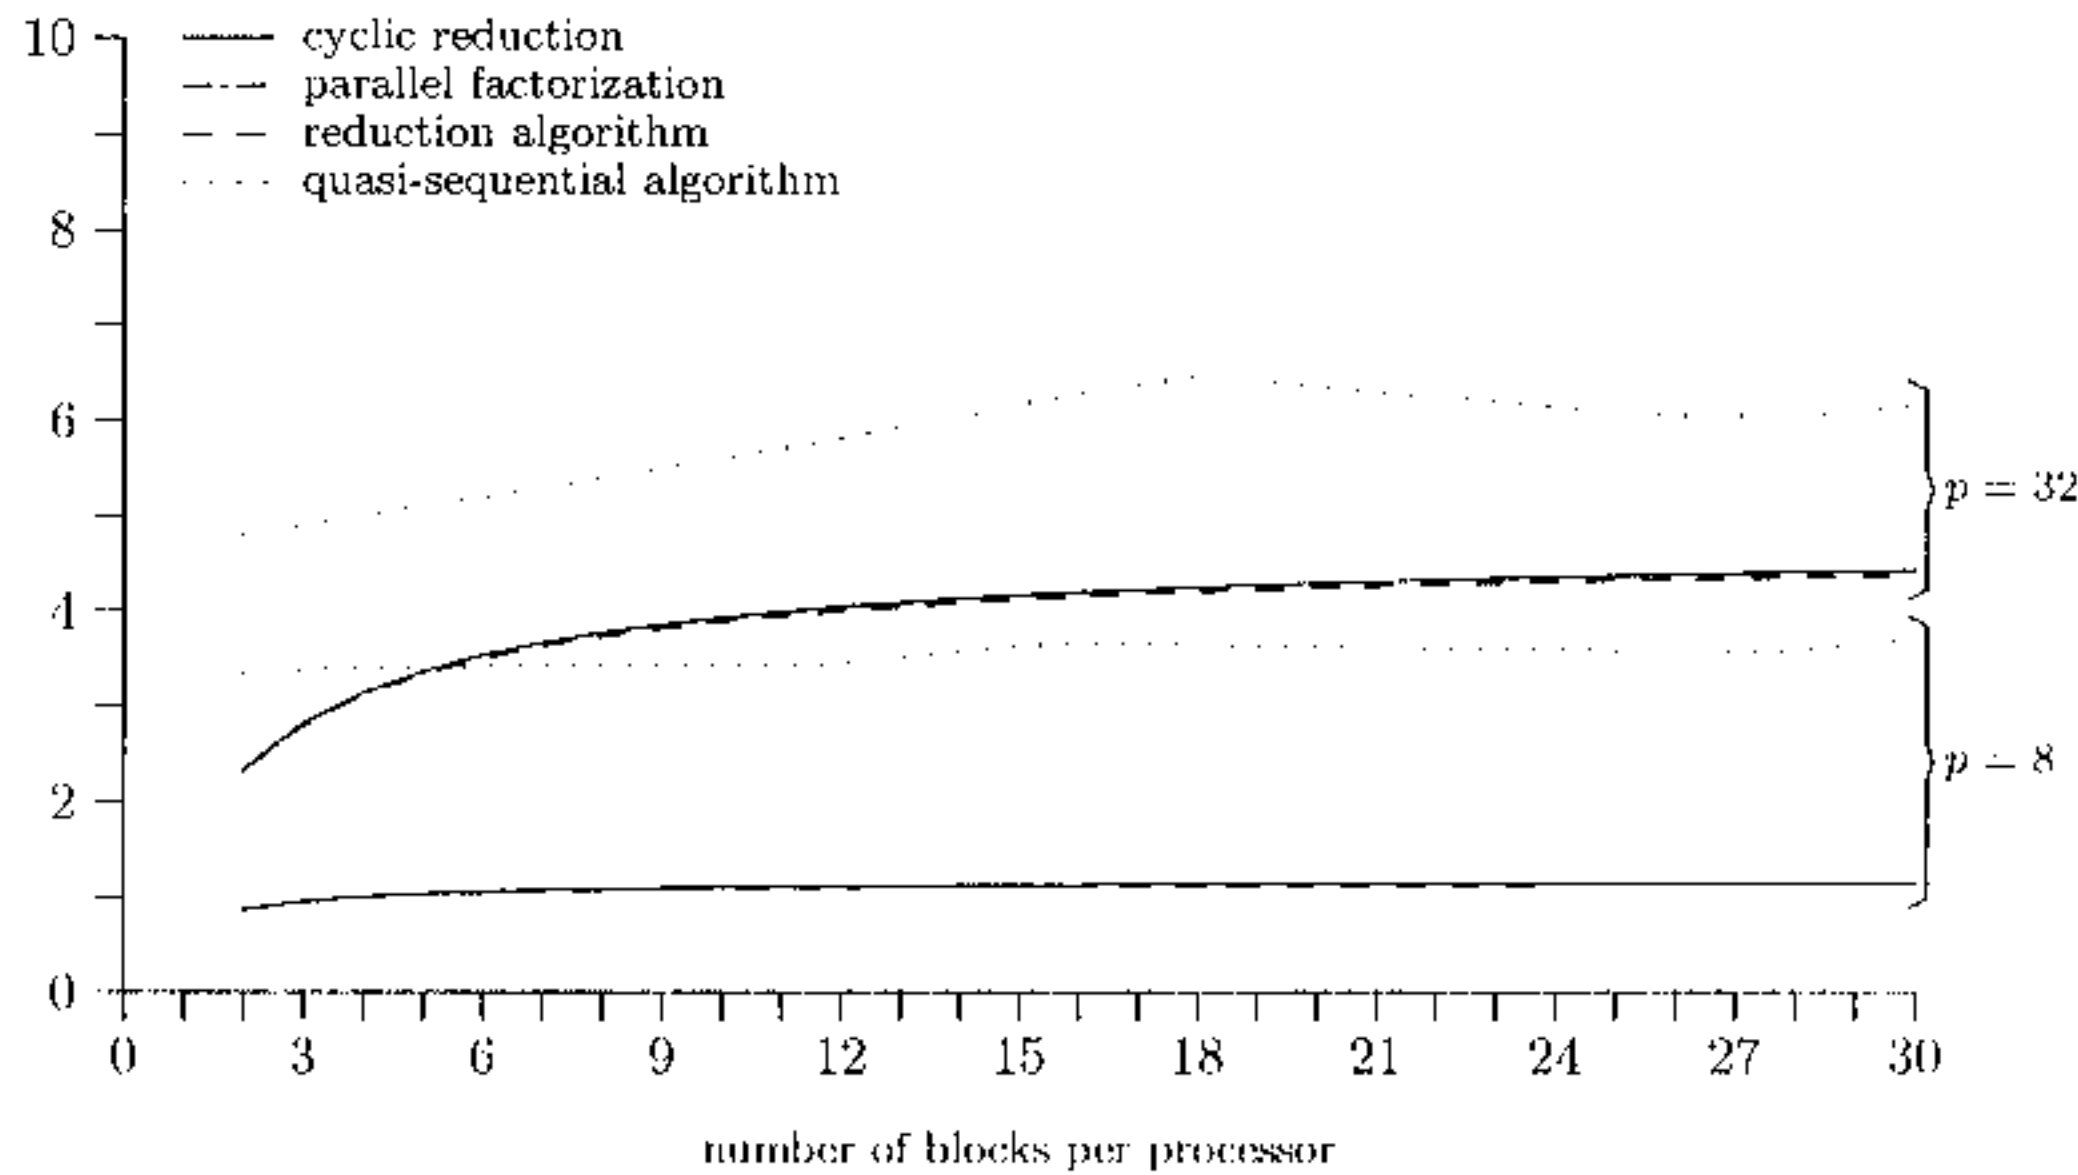


Figure 5. Speedups of the four algorithms for $m = 30$ and $p = 8, 32$.

vector of s elements is almost the same as that of s transmissions of scalars), then the quasi-sequential algorithm requires $O(pn)$ data transmissions whereas the parallel algorithms require $O(\log_2(p)m^2)$ data transmissions. Therefore, for moderate p and large m , the m^2 term will dominate the data transmission cost, giving the quasi-sequential algorithm additional advantage over the parallel algorithms. As the number of processors increases, the situation reverses. Overall the cost of data transmission of the parallel algorithms is smaller than that of the quasi-sequential algorithm for $m < p/\log_2(p)$. These results are also confirmed by our experiments (see Figures 4, 5, and 7-9).

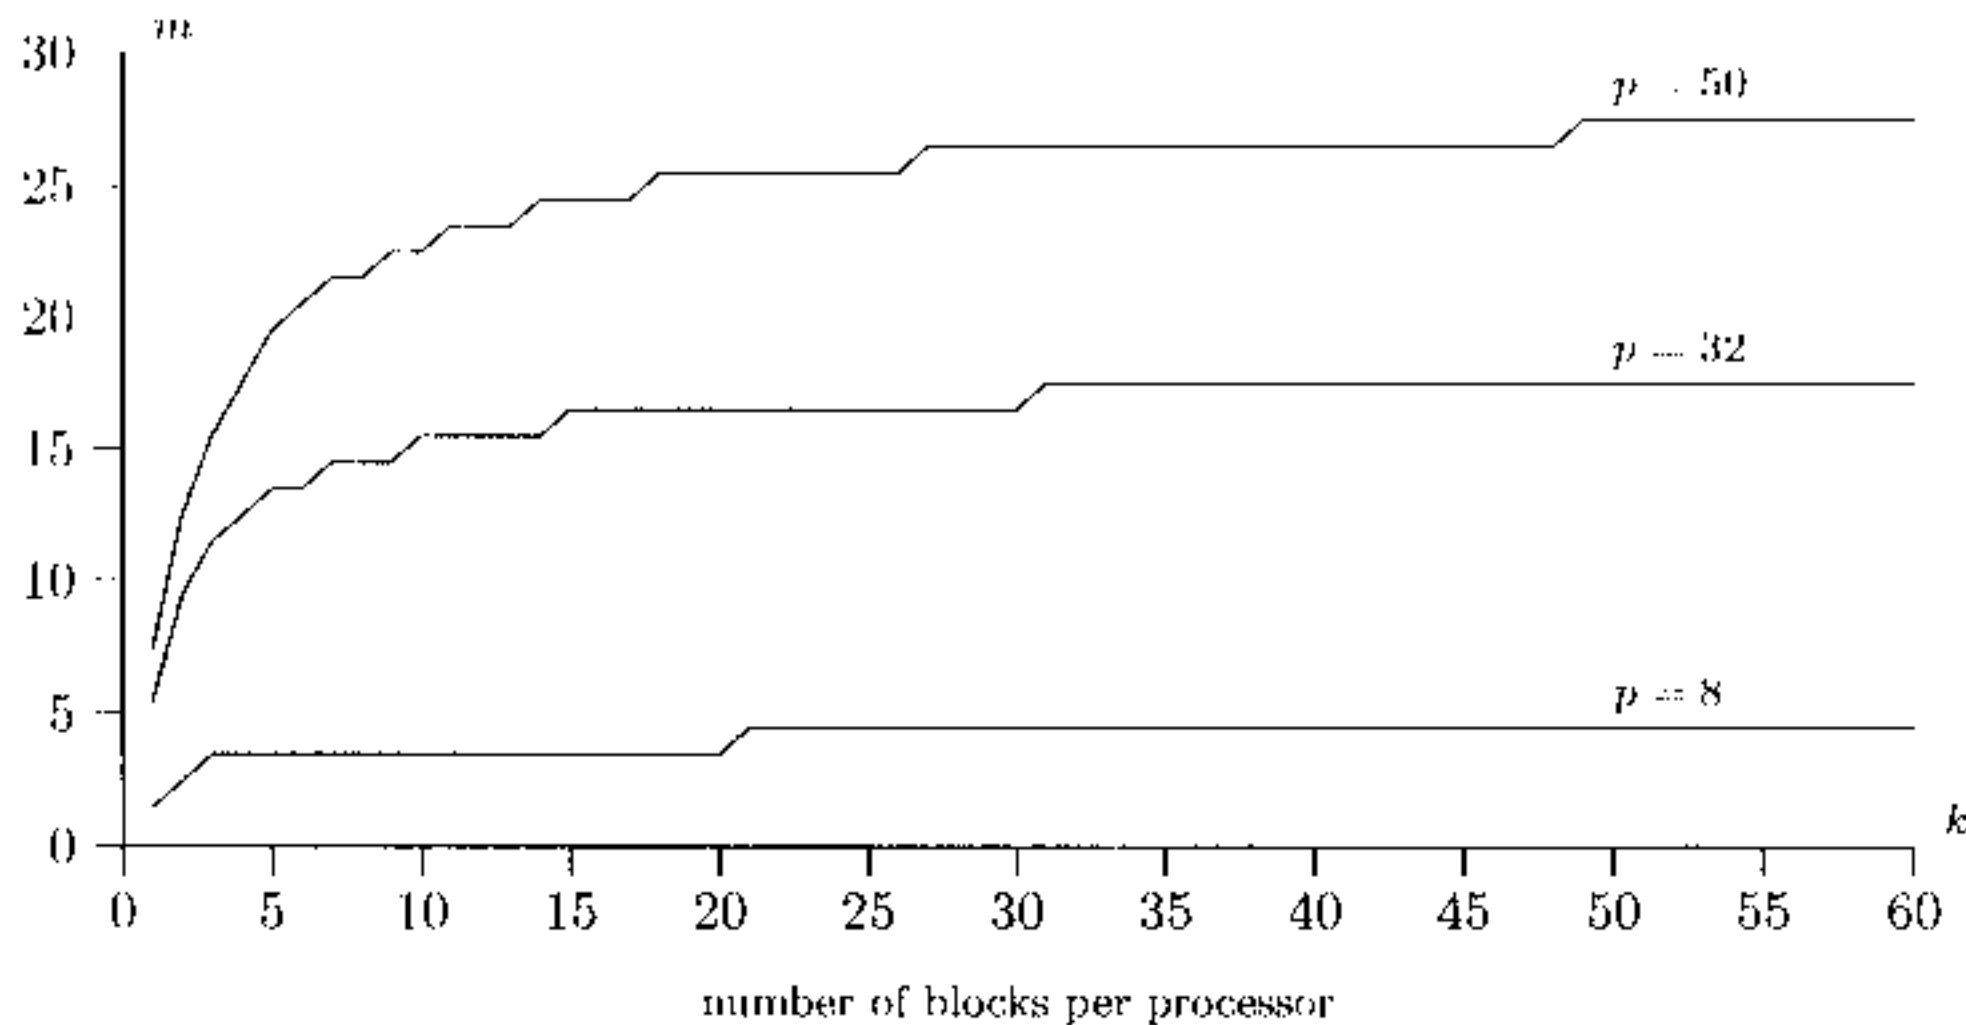


Figure 6. Theoretical comparison between the parallel factorization and the quasi-sequential algorithm. For each line (representing a given number of processors p) points above it are values of k and m for which the quasi-sequential algorithm outperforms the parallel factorization.

6.2. Memory Requirements

We will consider two cases of memory requirements of the proposed algorithms depending on when the solution of the linear system is performed. If the factorization step and the linear system solution step are performed separately, the parallel factorization algorithm requires to store one fill-in vector, and the cyclic reduction algorithm requires to store the off-diagonal elements of the




Figure 7. Speedups of the four algorithms for $m = 5$ and $p = 8, 32$.




Figure 8. Speedups of the algorithms for $p = 32$ and $k = \lfloor 800/m \rfloor$ variable. The dashed lines represent the theoretical speedups of the two algorithms.

Table 2. Per-processor memory requirements for the proposed algorithms when the coefficient matrix is factored before the solution of the associated linear system.

Algorithm	Workspace
sequential	$2nm^2 + nm$
quasi-sequential	$2kn^2 + (k + 1)m$
reduction	$(2k + \lceil \log_2 p \rceil)m^2 + (k + 1)m$
parallel factorization	$(3k + \lceil \log_2(p - 1) \rceil - 1)m^2 + (k + 1)m$
cyclic reduction	$(3k + \lceil \log_2(p - 1) \rceil - 1)m^2 + (k + 1)m$

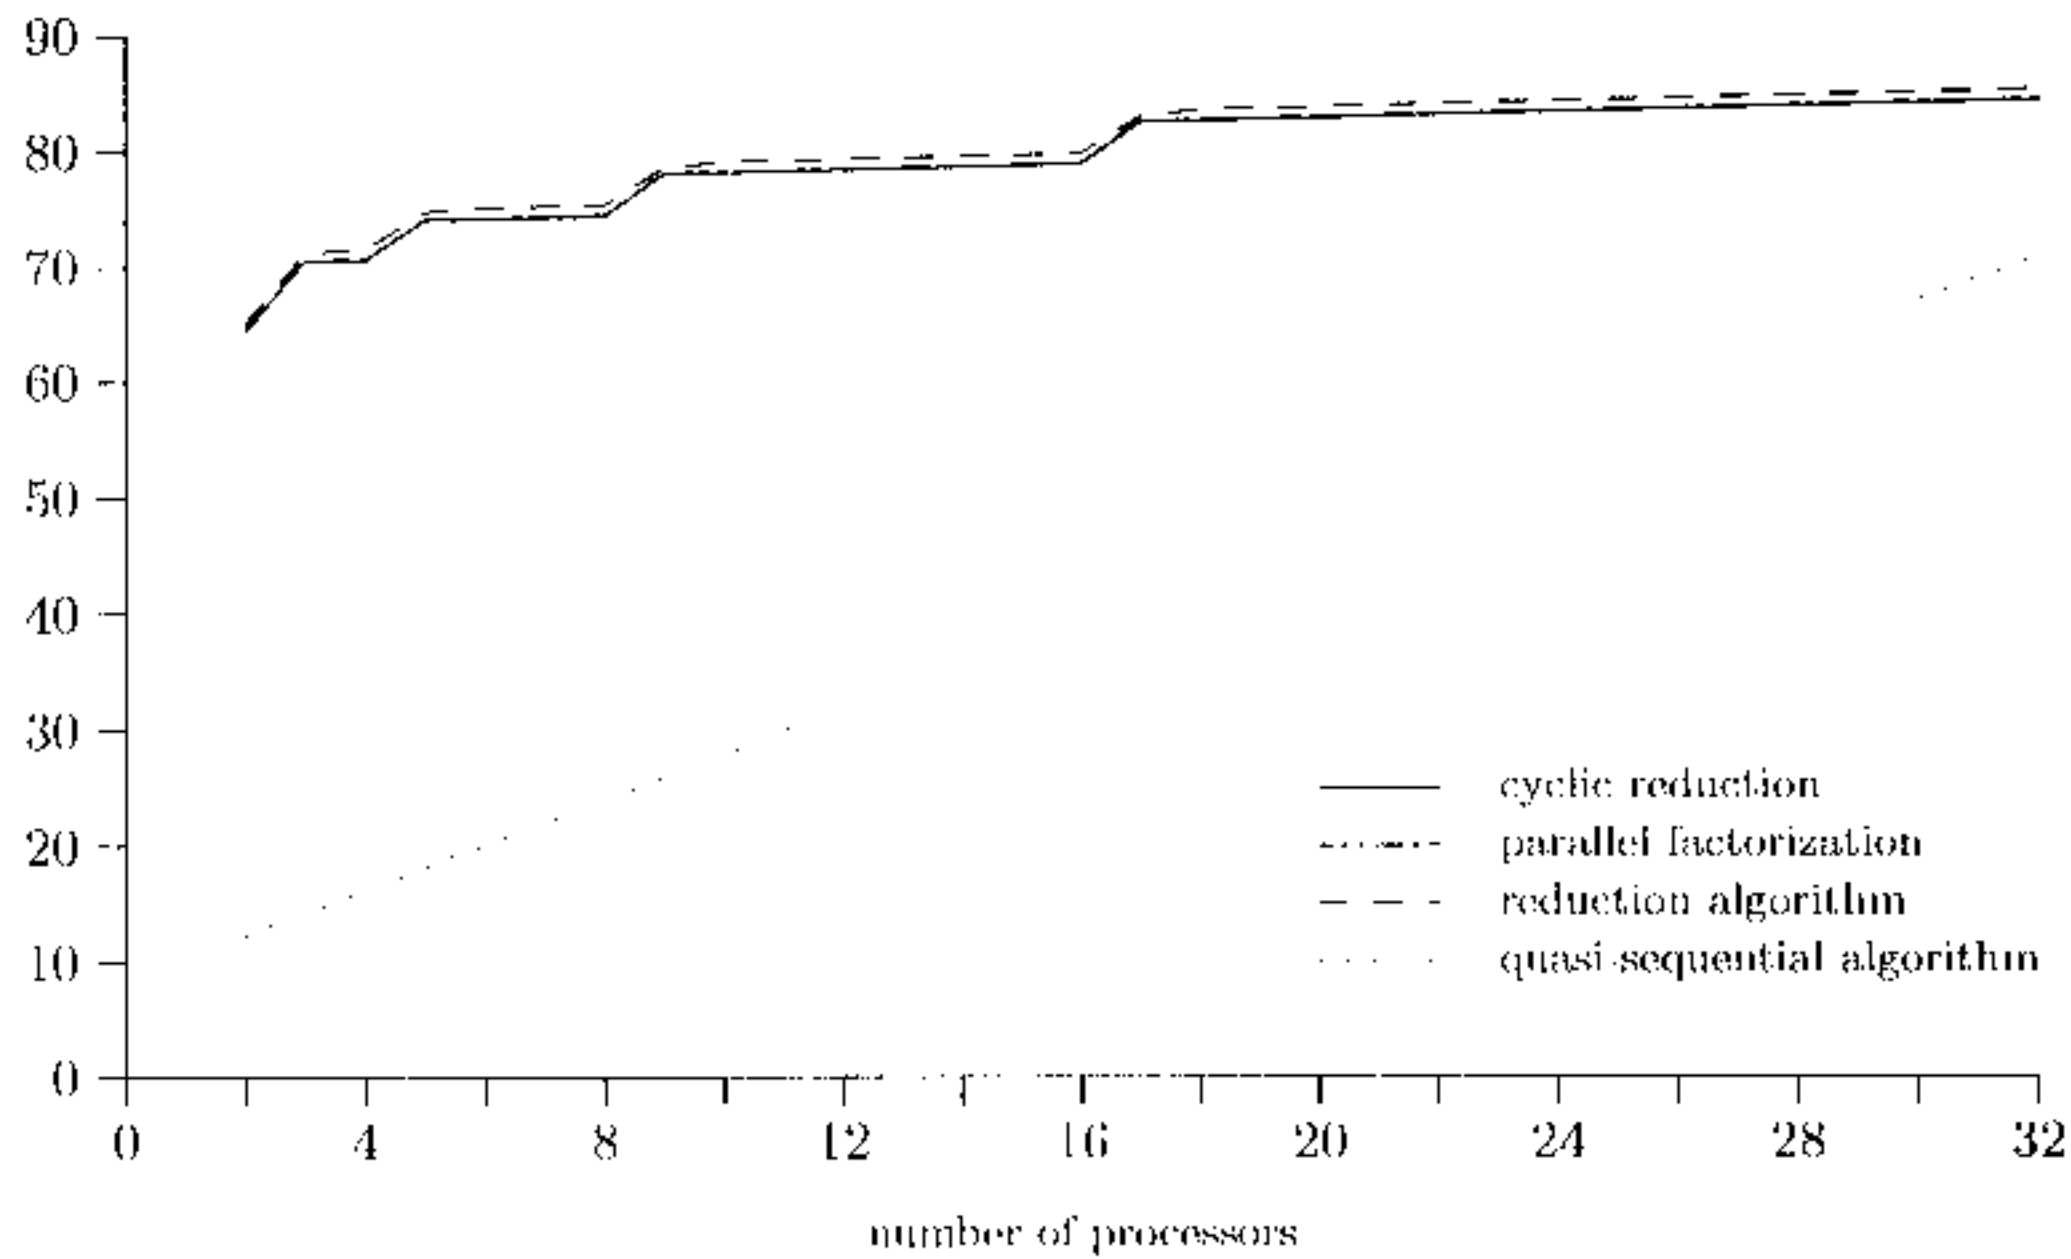


Figure 9. Time (in seconds) of execution for $k = 10$ and $m = 30$ on varying numbers of processors.

submatrices obtained at each step of the reduction. Table 2 presents the per-processor memory requirements of the proposed algorithms. It can be observed that the sequential and quasi-sequential algorithms have precisely the same overall memory requirements. The per-processor memory requirements of the reduction algorithm are slightly higher than those of the quasi-sequential algorithm for a moderate p , and substantially lower than those of the remaining two parallel algorithms for large k . The overall memory requirement of the parallel factorization and cyclic reduction algorithms is larger by km^2 elements than that of the reduction algorithm and by $(k + \lceil \log_2(p - 1) \rceil - 1)m^2$ elements than that of the quasi-sequential algorithm.

If the factorization is performed together with the linear system solution, the per-processor memory requirements of the proposed algorithms are summarized in Table 3.

Table 3. Per-processor memory requirements for the proposed algorithms when the solution of the associated linear system and the factorization of the coefficient matrix are performed at the same time.

Algorithm	Workspace
sequential	$2mm^2 + nm$
quasi-sequential	$2km^2 + (k + 1)m$
reduction	$(2k + 1)m^2 + (k + 1)m$
parallel factorization	$2km^2 + (k + 1)m$
cyclic reduction	$2km^2 + (k + 1)m$

It can be observed that in all cases the per-processor memory requirements of all parallel algorithms are the same and that the overall memory requirement of all the algorithms are almost the same.

7. NUMERICAL TESTS

The algorithms presented in previous sections have been coded in Parallel Fortran [15] with the Express communication library [16] on a network of 32 Microway transputers T800-20, each with a local memory of 1 Mb. The sequential algorithm has been implemented in Fortran on a single transputer T800-20 with 16 Mb of memory. The codes use the level 3-BLAS routines

DGEMM to perform matrix-matrix products, DGEMV to perform matrix-vector products, and two subroutines of LINPACK package: DGEFA to factorize the main diagonal blocks with only local row pivoting, and DGESL to solve triangular linear systems factored by DGEFA.

As a test problem we have chosen a linear system with $m_i = m$ for all i , and

$$D_i = I - h_i E, \quad C_i = -I - h_i E$$

arising from the numerical solution of the initial value problem

$$\begin{aligned} \mathbf{y}' &= E\mathbf{y} + \mathbf{b}(t), & t \in [t_0, T], \\ \mathbf{y}(t_0) &= \mathbf{y}_0 \end{aligned} \tag{9}$$

by means of the trapezoidal rule. The obtained block bidiagonal system is in general used as preconditioning when problem (9) is solved by means of Boundary Value Methods (see [4,5]). We have selected the $m \times m$ matrix E with eigenvalues in the negative part of the complex half-plane in order to have a stable solution (for the purpose of our experiments, E was obtained using a random number generator). Moreover, $h_i = \gamma h_{i-1}$, for $\gamma > 1$ fixed.

In Figures 4, 5, and 7, the measured speedups of the cyclic reduction, parallel factorization, reduction and quasi-sequential algorithms are reported for $m = 5, m = 15, m = 30$, different values of $k = (n+1)/p$, and for different numbers of processors. For the quasi-sequential algorithm, we have chosen $q_j = (j+1)q$, where q is a positive integer number fixed for all processors, j is the processor number, $j = 1, \dots, p$, in such a way that the speedup was the highest (the optimal value of q_j will be architecture dependent). For example, $q = k/4.5$ if $m = 5$, $q = k/10$ if $m = 15$ and $q = 1$ for $m = 30$. It should be noted that in this last case the values of k are relatively small (less than 30).

The results of all the experiments accord with the theoretical results of the previous section. Figures 4, 5, and 7 show that if a small number of processors is used, the quasi-sequential algorithm is always the fastest among the presented algorithms. For increasing numbers of processors, parallel algorithms become preferable (but only when the size of each block (m) is not too large). For $m = 15$ and $p = 32$, the quasi-sequential algorithm behaves almost exactly like the parallel algorithms. The increase in m not only reduces the utility of the parallel algorithms but also reduces the obtained speedup; for $p = 32$, the reduction is from the speedup of about 9 for $m = 5$ to speedup of less than 7 for $m = 30$. It should be also observed that the cyclic reduction and parallel factorization algorithms behave exactly the same and always outperform the reduction algorithm.

Figure 8 compares the performance of cyclic reduction and quasi-sequential algorithms for $p = 32$ processors, for a matrix of size $n = 25600$ and for approximately fixed $k = \lceil 800/m \rceil$ while varying the value of m . It also compares the practical and the theoretical speedup for both algorithms.

Figure 8 shows that for suitable large dimensions of the coefficient matrix A , theoretical speedups are very good approximations to those obtained on the parallel computer. Moreover, the cyclic reduction algorithm outperforms the quasi-sequential algorithm only if $m \leq 16$. Finally, if the optimal algorithm is selected properly for given n , m and p , it is possible to obtain speedup greater than 5.5 on 32 processors.

The last Figure 9 represents the “scalability” of the presented algorithms, that is, the time of execution for a different number of processors with k and m constant. This graph shows that the parallel algorithms scale nicely as the number of processors increases (the jumps in the time are related to the fact that the workload in the solution of the reduced system depends on $\log_2 p$). As predicted, for fixed k and m , the quasi-sequential algorithm does not scale well. It can be observed again that cyclic reduction and factorization algorithms behave similarly, slightly outperforming the reduction algorithm. Moreover, it can be predicted that the crossover point will occur for approximately 40 processors.

8. CONCLUSIONS

We have presented four parallel algorithms for the solution of block bidiagonal linear systems, discussed their implementation details, derived and discussed their arithmetical complexity functions and the memory requirement functions. It was shown that there is no simple answer which of the presented algorithms should be applied. The answer to this question depends on the parameters of the linear system itself (block size and number of blocks) as well as characteristics of the distributed memory computer that the systems are to be solved on (available memory per-processor and the communication overhead). Some guidelines for the selection process have been presented.

REFERENCES

1. I. Gladwell and M. Paprzycki, Parallel solution of almost block diagonal systems on the CRAY Y-MP using level 3 BLAS, *J. of Comput. and Appl. Mathem.* **45**, 181–189 (1993).
2. K. Wright, Parallel treatment of block bidiagonal matrices in the solution of ordinary differential boundary value problems, *J. of Comput. and Appl. Mathem.* **45**, 191–200 (1993).
3. S.J. Wright, Stable parallel algorithms for two-point boundary value problems, *SIAM J. Sci. Stat. Comput.* **13**, 742–764 (1992).
4. P. Amodio and F. Mazzia, Parallel block preconditioning for the solution of boundary value methods, Report n. 12/93 of the Dipartimento di Matematica, Università di Bari, Bari, Italy.
5. P. Amodio and F. Mazzia, Parallel iterative solvers for boundary value methods, Report n. 10/94 of the Dipartimento di Matematica, Università di Bari, Bari, Italy.
6. U.M. Ascher and P.S. Chan, On parallel methods for boundary value ODEs, *Computing* **46**, 1–17 (1991).
7. G.H. Golub and C.F. Van Loan, *Matrix Computations*, The Johns Hopkins University Press, Baltimore, (1989).
8. M. Paprzycki and I. Gladwell, Solving almost block diagonal systems on parallel computers, *Parallel Computing* **17**, 133–153 (1991).
9. L. Brugnano, A parallel solver for tridiagonal linear systems for distributed memory parallel computers, *Parallel Computing* **17**, 1017–1023 (1991).
10. P. Amodio and L. Brugnano, Parallel factorizations and parallel solvers for tridiagonal linear systems, *Lin. Alg. Appl.* **172**, 813–823 (1992).
11. P. Amodio, L. Brugnano and T. Politi, Parallel factorizations for tridiagonal matrices, *SIAM J. Numer. Anal.* **30**, 813–823 (1993).
12. D. Heller, Some aspects of the cyclic reduction algorithm for block tridiagonal linear systems, *SIAM J. Numer. Anal.* **13**, 484–496 (1976).
13. J.M. Ortega, *Introduction to Parallel and Vector Solution of Linear Systems*, Plenum Press, New York, (1988).
14. M. Paprzycki and P. Stpiczyński, Solving linear recurrence systems on parallel computers, In *Proceedings of the Conference Mardi Gras '94*, Baton Rouge, February 10–12, 1994, Nova Science Publishers, New York (to appear).
15. *Parallel Fortran User Guide*, 3I Ltd, Livingstone, (1988).
16. *Express Fortran User's Guide*, ParaSoft Corp., Pasadena, (1990).