

Efficient Tridiagonal Solvers on Multicomputers

Xian-He Sun, *Member, IEEE*, Hong Zhang, and Lionel M. Ni, *Senior Member, IEEE*

Abstract—Three parallel algorithms, namely the parallel partition LU (PPT) algorithm, the parallel partition hybrid (PPH) algorithm, and the parallel diagonal dominant (PDD) algorithm are proposed for solving tridiagonal linear systems on multicomputers. These algorithms are based on the divide-and-conquer parallel computation model. The PPT and PPH algorithms support both pivoting and nonpivoting. The PPT algorithm is good when the number of processors is small; otherwise, the PPH algorithm is better. When the system is diagonal dominant, the PDD algorithm is highly parallel and provides an approximate solution which equals to the exact solution within machine accuracy. Both computation and communication complexities of the three algorithms are presented. All three methods proposed in this paper have been implemented on a 64-node nCUBE-1 multicomputer. The analytic results match closely with the results measured from the nCUBE-1 machine.

Index Terms—Communication complexity, divide-and-conquer, LU decomposition, matrix partitioning, parallel numerical algorithms, multicomputers, tridiagonal systems.

I. INTRODUCTION

WE are interested in solving the following tridiagonal linear system of equations

$$Ax = d \quad (1.1)$$

where A is a tridiagonal matrix of order n

$$A = \begin{bmatrix} b_0 & c_0 & & & & \\ a_1 & b_1 & c_1 & & & \\ & \cdot & \cdot & \cdot & & \\ & & & & \cdot & \\ & & & & & \cdot \\ & & & & & & a_{n-2} & b_{n-2} & c_{n-2} \\ & & & & & & a_{n-1} & b_{n-1} & \end{bmatrix} \quad (1.2)$$

$x = (x_0, \dots, x_{n-1})^T$ and $d = (d_0, \dots, d_{n-1})^T$. We assume that A , x , and d have real coefficients. Extension to the complex case is straightforward.

Solving tridiagonal linear systems is one of the most important problems in scientific computing. It is involved in the solution of differential equations and in various other areas of science and engineering. A variety of methods have

been developed for solving (1.1) on parallel computers. A good survey of these methods can be found in [9], [13], [14], and [18]. Among them, the *recursive doubling method* (RCD) developed by Stone [22] and the *cyclic reduction or odd-even reduction method* (OER) developed by Hockney [10] are able to solve (1.1) in $O(\log n)$ time using n processors. Here and throughout, \log denotes \log_2 . However, in a realistic parallel processing environment, the number of processors, p , is usually less or much less than the dimension of the matrix, n . Lawrie and Sameh [16] proposed an algorithm in 1984 for a practical parallel environment. Recently, a *parallel prefix* (PP) method developed by Egecioglu, Koc, and Laub [7] modifies the RCD method to be applied to a limited number of processors, i.e., $p < n$. Wang's *partition* algorithm [23] is another relatively new method which is designed for a realistic parallel environment, and it has been efficiently implemented in [11]. While the OER algorithm is a popular choice for vector machines, the Lawrie-Sameh's algorithm, RCD algorithm, and Wang's algorithm are good candidates for multiprocessors.

A *multicomputer* is a distributed-memory multiprocessor in which each memory module is physically associated with each processor. A point-to-point interprocessor communication network provides a mechanism for communication between processors [12]. Multicomputers with hundreds or thousands of processors are commercially available. Multicomputers hold the promising potential for providing massive parallelism. It is for this reason the algorithms for multicomputers are studied.

Three parallel algorithms are proposed in this paper. All these algorithms are developed based on Sheman and Morrison matrix modification formula [20]. To start with, a tridiagonal matrix A is partitioned into p submatrices. Solving these p subsystems in parallel, it leads to a reduced tridiagonal system of order $2(p-1)$. If p is small, the *parallel partition LU* (PPT) algorithm solves this small tridiagonal system in a serial mode. If p is large, the reduced system can be solved in parallel by the PP method. This algorithm is called the *parallel partition hybrid* (PPH) algorithm. In many applications, the tridiagonal matrix A is evenly diagonal dominant. In this case half of the off-diagonal elements of the small $2(p-1)$ -dimensional matrix converge to zero exponentially. A fast and highly parallel algorithm, namely the *parallel diagonal dominant* (PDD) algorithm, is proposed. All the three algorithms have less communication requirement than Wang's algorithm [17], [23]. The PPT algorithm achieves a time complexity very close to the algorithm proposed by Ho and Johnsson [11]. Based on a matrix modification formula, our matrix partitioning scheme ends up very similar to that of Lawrie and Sameh [16]. However, our algorithms use different strategies in solving the reduced system, which reduce the communication overhead

Manuscript received March 3, 1989; revised October 15, 1990 and March 7, 1991. This work was supported in part by the NSF Grants CCR-87-16833, ECS-88-14027, and by the Applied Mathematical Sciences Program of the Ames Laboratory, U.S. Department of Energy under Contract W-7405-ENG-82.

X.-H. Sun was with the Department of Computer Science, Michigan State University, East Lansing, MI 48824. He is now with the Ames Laboratory, Iowa State University, Ames, IA 50011.

H. Zhang is with the Mathematical Sciences Department, Clemson University, Clemson, SC 29634-1907.

L. M. Ni is with the Department of Computer Science, Michigan State University, East Lansing, MI 48824.

IEEE Log Number 9104167.

significantly. Particularly, the PDD algorithm gives a fast solution that approximates the exact solution to the machine accuracy when $n/p \gg 1$. Although the matrix modification formula being used is not crucial in designing these algorithms, it provides a nice mathematical equation that an error analysis can be conducted formally [25].

Communication mechanism has a great impact on the performance of multicomputers. Thus, the communication pattern of parallel algorithms should be carefully designed to reduce the communication complexity. In this study, the algorithms are evaluated based on both computation and communication complexities.

This paper is organized as follows. Section II briefly reviews the parallel prefix method, which is used in the development of our parallel algorithms. The matrix partitioning method is presented in Section III. Based on the matrix partitioning scheme, three parallel algorithms, PPT, PPH, and PDD are discussed in Section IV, Section V, and Section VI, respectively. These parallel algorithms are implemented on a 64-node nCUBE multicomputer. Section VII presents analytic and experimental results and comparison discussion. Advantages of our proposed algorithms are summarized in Section VIII.

II. PRELIMINARY

Two methods will be used in our algorithms. They are LU decomposition method and parallel prefix method. For completeness, the latter is briefly presented in this section.

The RCD method uses $O(\log n)$ parallel computation time to solve (1.1) on a parallel computer with n processors [10], [22]. The PP method [7] modifies the RCD method for p processors with $p < n$.

Equation (1.1) can be written as

$$a_i x_{i-1} + b_i x_i + c_i x_{i+1} = d_i, \quad 0 \leq i \leq n-1 \quad (2.1)$$

where $x_{-1} = x_n = 0$. Solving for x_{i+1} , we have

$$x_{i+1} = \left(-\frac{b_i}{c_i}\right)x_i + \left(-\frac{a_i}{c_i}\right)x_{i-1} + \left(\frac{d_i}{c_i}\right) = \alpha_i x_i + \beta_i x_{i-1} + \gamma_i. \quad (2.2)$$

Here $c_i \neq 0$ is assumed. Equation (2.2) then can be written in the matrix form as

$$\begin{bmatrix} x_{i+1} \\ x_i \\ 1 \end{bmatrix} = \begin{bmatrix} \alpha_i & \beta_i & \gamma_i \\ 1 & 0 & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x_i \\ x_{i-1} \\ 1 \end{bmatrix}.$$

Define

$$\mathbf{X}_{i+1} = \begin{bmatrix} x_{i+1} \\ x_i \\ 1 \end{bmatrix} \quad \text{with } x_{-1} = x_n = 0.$$

Equation (2.2) becomes

$$\mathbf{X}_{i+1} = \mathbf{B}_i \mathbf{X}_i, \quad 0 \leq i \leq n-1, \quad (2.3)$$

and \mathbf{X}_i ($1 \leq i \leq n$) can be expressed in terms of \mathbf{X}_0 as

$$\mathbf{X}_i = \mathbf{B}_{i-1} \mathbf{B}_{i-2} \cdots \mathbf{B}_1 \mathbf{B}_0 \mathbf{X}_0, \quad 1 \leq i \leq n.$$

Now solving (1.1) becomes the finding of all the partial products of matrices \mathbf{B}_i for $0 \leq i \leq n-1$ and x_0 . If $p < n$, we

first evenly distribute matrices \mathbf{B}_i 's to p processors, perform sequential matrix multiplication on each processor, then use the prefix method on p processors. There are $(\log p)+1$ parallel communication steps in applying the prefix method with p processors.

Let $\mathbf{C}_i^j = \mathbf{B}_j \mathbf{B}_{j-1} \cdots \mathbf{B}_i$. Then \mathbf{C}_i^i is a 3×3 matrix. Since the last row of \mathbf{C}_i^j always equals $[0, 0, 1]$, only six entries of \mathbf{C}_i^j need to be transferred at each parallel communication. Fig. 1 shows the communication pattern of the parallel prefix method for obtaining all the partial matrix products with $n = 16$ and $p = 8$.

The actual communication complexity of the PP method depends on the mapping of computation units into processors, the interconnection topology of the multicomputer, and the underlying communication mechanism. In the case of hypercube topology, the communication pattern shown in Fig. 1 can be mapped such that the dilation cost is no greater than 2 [13].

One drawback of the PP method is that the method is unstable when $|c_i|$ is small relative to $|a_i| + |b_i|$ in (2.1). It has been shown that some stability problems arise in the use of the algorithm when the size of the system is large [6].

III. A MATRIX PARTITIONING TECHNIQUE

Our parallel algorithms are based on the divide and conquer model of parallel computation. In the divide part, the matrix \mathbf{A} is partitioned into submatrices. For convenience we assume that $n = pm$. The matrix \mathbf{A} in (1.2) can be written as

$$\mathbf{A} = \tilde{\mathbf{A}} + \Delta \mathbf{A}, \quad (3.1)$$

where

$$\tilde{\mathbf{A}} = \begin{bmatrix} \mathbf{A}_0 & & & \\ & \mathbf{A}_1 & & \\ & & \ddots & \\ & & & \mathbf{A}_{p-1} \end{bmatrix} \quad \Delta \mathbf{A} = \begin{bmatrix} & c_{p-1} & & \\ a_m & & c_{2m-1} & \\ & a_{2m} & & c_{3m-1} \\ & & \ddots & \\ & & & a_{m(p-1)} & c_{p(p-1)m} \end{bmatrix}$$

The submatrices \mathbf{A}_i ($i = 0, \dots, p-1$) are $m \times m$ tridiagonal matrices. Let \mathbf{e}_i be a column vector with its i th ($0 \leq i \leq n-1$) element being 1 and all the other entries being zero. We have

$$\Delta \mathbf{A} = [a_m \mathbf{e}_m, c_{m-1} \mathbf{e}_{m-1}, a_{2m} \mathbf{e}_{2m}, c_{2m-1} \mathbf{e}_{2m-1}, \dots, c_{(p-1)m-1} \mathbf{e}_{(p-1)m-1}] \begin{bmatrix} \mathbf{e}_{m-1}^T \\ \mathbf{e}_m^T \\ \vdots \\ \mathbf{e}_{(p-1)m-1}^T \\ \mathbf{e}_{(p-1)m}^T \end{bmatrix} = \mathbf{V} \mathbf{E}^T$$

where both \mathbf{V} and \mathbf{E} are $n \times 2(p-1)$ matrices. Thus, we have

$$\mathbf{A} = \tilde{\mathbf{A}} + \mathbf{V} \mathbf{E}^T. \quad (3.2)$$

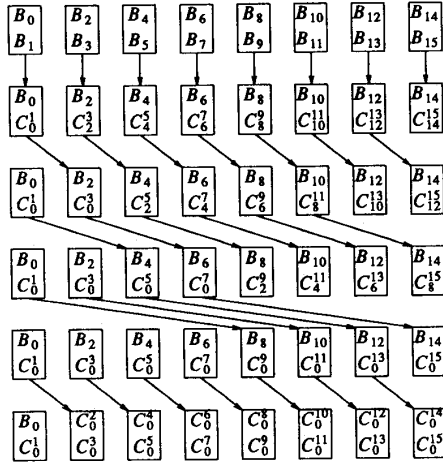


Fig. 1. The communication pattern of the parallel prefix method for obtaining partial products.

Based on the matrix modification formula originally defined by Sherman and Morrison [20] for rank-one changes and generalized by Woodbury [5], [24], and assuming that all A_i 's are invertible, (1.1) can be solved by

$$\begin{aligned} \mathbf{x} &= \mathbf{A}^{-1}\mathbf{d} = (\tilde{\mathbf{A}} + \mathbf{V}\mathbf{E}^T)^{-1}\mathbf{d} \\ &= \tilde{\mathbf{A}}^{-1}\mathbf{d} - \tilde{\mathbf{A}}^{-1}\mathbf{V}(\mathbf{I} + \mathbf{E}^T\tilde{\mathbf{A}}^{-1}\mathbf{V})^{-1}\mathbf{E}^T\tilde{\mathbf{A}}^{-1}\mathbf{d}. \end{aligned} \quad (3.3)$$

Note that \mathbf{I} is an identity matrix and $\mathbf{I} + \mathbf{E}^T\tilde{\mathbf{A}}^{-1}\mathbf{V}$ is a pentadiagonal matrix of order $2(p-1)$. We introduce a permutation matrix \mathbf{P} such that

$$\mathbf{P}\mathbf{z} = (z_1, z_0, z_3, z_2, \dots, z_{2p-3}, z_{2(p-2)})^T$$

for all $\mathbf{z} \in \mathbf{R}^{2(p-1)}$.

Equation (3.3) then becomes

$$\mathbf{x} = \tilde{\mathbf{A}}^{-1}\mathbf{d} - \tilde{\mathbf{A}}^{-1}\mathbf{V}\mathbf{P}(\mathbf{P} + \mathbf{E}^T\tilde{\mathbf{A}}^{-1}\mathbf{V}\mathbf{P})^{-1}\mathbf{E}^T\tilde{\mathbf{A}}^{-1}\mathbf{d}. \quad (3.4)$$

Note that $\mathbf{Z} = (\mathbf{P} + \mathbf{E}^T\tilde{\mathbf{A}}^{-1}\mathbf{V}\mathbf{P})$ is a tridiagonal matrix of order $2(p-1)$. Let

$$\tilde{\mathbf{A}}\tilde{\mathbf{x}} = \mathbf{d}, \quad (3.5)$$

$$\tilde{\mathbf{A}}\mathbf{Y} = \mathbf{V}\mathbf{P}, \quad (3.6)$$

$$\mathbf{h} = \mathbf{E}^T\tilde{\mathbf{x}}, \quad (3.7)$$

$$\mathbf{Z} = \mathbf{P} + \mathbf{E}^T\mathbf{Y}, \quad (3.8)$$

$$\mathbf{Z}\mathbf{y} = \mathbf{h}, \quad (3.9)$$

$$\Delta\mathbf{x} = \mathbf{Y}\mathbf{y}. \quad (3.10)$$

From (3.5)–(3.10), (3.4) becomes

$$\mathbf{x} = \tilde{\mathbf{x}} - \Delta\mathbf{x}. \quad (3.11)$$

In (3.5) and (3.6), $\tilde{\mathbf{x}}$ and \mathbf{Y} are solved by the LU decomposition method. By the structure of $\tilde{\mathbf{A}}$ and $\mathbf{V}\mathbf{P}$, this is equivalent to solve

$$\mathbf{A}_i[\tilde{\mathbf{x}}^{(i)}, \mathbf{v}^{(i)}, \mathbf{w}^{(i)}] = [\mathbf{d}^{(i)}, a_{im}\mathbf{e}_{0, c(i+1)m-1}\mathbf{e}_{m-1}],$$

$i = 0, \dots, p-1. \quad (3.12)$

Here we have $a_0 = c_{n-1} = 0$, $\mathbf{e}_0, \mathbf{e}_{m-1} \in \mathbf{R}^m$, $\tilde{\mathbf{x}}^{(i)}$ and $\mathbf{d}^{(i)}$ are the i th block of $\tilde{\mathbf{x}}$ and \mathbf{d} , respectively, and $\mathbf{v}^{(i)}, \mathbf{w}^{(i)}$ are possible nonzero column vectors of the i th row block of \mathbf{Y} . Equation (3.12) implies that we only need to solve three linear systems of order m with the same LU decomposition for each i ($i = 0, \dots, p-1$). In addition, we can skip the first $m-1$ forward substitutions for the third system since the first $m-1$ components of the vector at the right-hand side are all zeros.

Equation (3.7) only picks $2(p-1)$ specified components from the vector $\tilde{\mathbf{x}}$. The evaluation of (3.8) uses those possible nonzero entries of specified $2(p-1)$ rows of \mathbf{Y} together with \mathbf{P} to form matrix \mathbf{Z} . There is no computation or communication involved in computing \mathbf{h} and \mathbf{Z} .

Solving (3.9) is the major computation involved in the conquer part of our algorithms. Since \mathbf{Y} has at most two nonzero entries at each row, the evaluation of (3.10) takes four arithmetic operations per row.

Based on the above observations and together with a careful scaling process, the computational complexity required to solve (1.1) in a sequential processor is stated in the following theorem.

Theorem 3.1: Equation (1.1) can be solved using (3.5)–(3.11) with $17n - 6p - 23$ arithmetic operations without pivoting and $24n - 13p - 34$ arithmetic operations with pivoting.

IV. PARALLEL PARTITION LU (PPT) ALGORITHM

Based on the matrix partitioning technique described previously, in following sections we show how to implement parallel algorithms on multicomputers. The *Parallel Partition LU* (PPT) algorithm is described first in this section.

Using p processors, the PPT algorithm to solve (1.1) consists of the following steps:

- Step 1. Allocate $\mathbf{A}_i, \mathbf{d}^{(i)}$ and elements $a_{im}, c_{(i+1)m-1}$ to the i th node, where $0 \leq i \leq p-1$.
- Step 2. Use the LU decomposition method to solve (3.12). All computations can be executed in parallel and independently on p processors.
- Step 3. Send $\tilde{x}_0^{(i)}, \tilde{x}_{m-1}^{(i)}, v_0^{(i)}, v_{m-1}^{(i)}, w_0^{(i)}, w_{m-1}^{(i)}$ ($0 \leq i \leq n-1$) to all other nodes from the i th node to form matrix \mathbf{Z} and vector \mathbf{h} [(3.7) and (3.8)] on each node. Here and throughout the subindex indicates the component of the vector.
- Step 4. Use the LU decomposition method to solve $\mathbf{Z}\mathbf{y} = \mathbf{h}$ [(3.9)] on all nodes simultaneously. Note that \mathbf{Z} is a $2(p-1)$ dimensional tridiagonal matrix.
- Step 5. Compute (3.10) and (3.11). We have

$$\begin{aligned} \Delta\mathbf{x}^{(i)} &= [\mathbf{v}^{(i)}, \mathbf{w}^{(i)}] \begin{bmatrix} y_{2i-1} \\ y_{2i} \end{bmatrix}, \\ \mathbf{x}^{(i)} &= \tilde{\mathbf{x}}^{(i)} - \Delta\mathbf{x}^{(i)}. \end{aligned} \quad (4.1)$$

This step is executed in parallel on p processors.

As mentioned in Section I, the underlying communication mechanism of multicomputers has a great impact on the performance of parallel algorithms. In the *store-and-forward* mechanism, which is used in all first generation hypercube

multicomputers, the message transfer time between two adjacent processors can be expressed as $\alpha + \beta S$, where α is the communication latency, β is the transmission time per byte, and S is the number of bytes in the message [8]. If a message is delivered to h hops away, the message transfer time can be roughly estimated as $h(\alpha + \beta S)$. In second generation multicomputers, advanced communication mechanisms are adopted, such as the *circuit switching* used in iPSC-2 [2] and the *wormhole routing* used in Ametek 2010 [1]. In these new communication paradigms, the message transfer time is almost independent of the distance (number of hops) between processors [2]. In this case, if network contention is not considered, the transfer time of a message with S bytes can always be expressed as $\alpha + \beta S$ regardless of the distance that a message has to traverse.

The PPT algorithm has an all-to-all communication which is also called *total data-exchange* or *all-to-all broadcasting* [15] communication. Fig. 2 shows the communication pattern of the PPT algorithm for the case of $p = 8$.

One of the best ways to handle the all-to-all communication is using the *butterfly* communication fashion as shown in Fig. 3 for the case of $p = 8$.

The butterfly communication pattern can be perfectly embedded in the hypercube topology [19]. Note that in the data transfer, each processor, upon receiving a message, has to add in its own data and forward the enlarged message to the following processors. As a result, the communication time required in total data-exchange can be estimated to be

$$(\log p)\alpha + (p - 1)S\beta. \quad (4.2)$$

Equation (4.2) shows that the total data-exchange with butterfly fashion has the same time complexity as using the *spanning binomial tree* technique [15]. By Johnsson and Ho [15], it is optimum within a factor of two. Now we are ready to state the computation and communication complexities of the PPT algorithm.

Theorem 4.1: The PPT algorithm solves (1.1) in $17(n/p) + 16p - 45$ and $24(n/p) + 22p - 69$ parallel arithmetic operations for nonpivoting and pivoting, respectively. With 4-byte per data element, it requires $\alpha \log p + 16(p - 1)\beta$ communication times, where p is the number of processors and $n = mp$.

Let T_{LU} , T_{SPT} , and T_{PPT} be the time required to solve (1.1) using the sequential LU decomposition algorithm, the sequential partitioning algorithm [(3.5)–(3.11)], and the PPT algorithm, respectively. Let τ_{comp} represent the unit of a computation operation normalized to the communication time. From those theorems shown previously, we have

$$T_{LU} = (8n - 7)\tau_{comp} \quad \text{without pivoting} \quad (4.3)$$

$$T_{LUP} = (11n - 12)\tau_{comp} \quad \text{with pivoting} \quad (4.4)$$

$$T_{SPT} = (17n - 6p - 23)\tau_{comp} \quad \text{without pivoting} \quad (4.5)$$

$$T_{SPT} = (24n - 13p - 34)\tau_{comp} \quad \text{with pivoting} \quad (4.6)$$

$$T_{PPT} = \left(17\frac{n}{p} + 16p - 45\right)\tau_{comp} + (\log p)\alpha + 16(p - 1)\beta \quad \text{without pivoting} \quad (4.7)$$

$$T_{PPT} = \left(24\frac{n}{p} + 22p - 69\right)\tau_{comp} + (\log p)\alpha$$




Fig. 2. The communication pattern of the PPT algorithm.




Fig. 3. The butterfly communication pattern.

$$+ 16(p - 1)\beta \quad \text{with pivoting} \quad (4.8)$$

In Section VII, the theoretical results obtained here will be compared with the measured results obtained from experiments on a 64-node nCUBE-1.

V. THE PARALLEL PARTITION HYBRID (PPH) ALGORITHM

The PPT algorithm competes well with any existing parallel algorithms for solving (1.1) when $p \ll n$. However, the efficiency of the PPT algorithm decreases as the number of processors, p , increases. This is due to the fact that the major computation in the conquer part (Step 4 in Section IV) of the PPT algorithm is redundant down on each node. For this reason, we use the PP method (see Section II), the limited processor version of the RCD method, to solve (3.9). In order to apply the PP method, all the superdiagonal elements of the coefficient matrix must be nonzero. The following theorem is needed in applying the PP method to (3.9).

Theorem 5.1: If all the superdiagonal elements of the matrix \mathbf{A} are nonzero, the tridiagonal matrix $\mathbf{Z} = \mathbf{P} + \mathbf{E}^T \mathbf{A}^{-1} \mathbf{V} \mathbf{P}$ has nonzero superdiagonal elements.

Proof: The superdiagonal elements of \mathbf{Z} are either equal to 1 or the first components of the solutions

$$\mathbf{A}_i \mathbf{w}^{(i)} = c_{(i+1)m-1} \mathbf{e}_{m-1}, \quad i = 1, \dots, p - 2 \quad (5.1)$$

where $\mathbf{w}^{(i)}, \mathbf{e}_{m-1} \in \mathbf{R}^m$. Suppose $w_0^{(i)} = 0$, then $w_j^{(i)} = 0$ for $j = 1, \dots, m - 1$, since the superdiagonal elements of \mathbf{A}_i are nonzero. Therefore, we have $\mathbf{A}_i \mathbf{w}^{(i)} = 0$, which is a contradiction to (5.1). \square

The new algorithm, namely the *parallel partition hybrid* (PPH) algorithm, is similar to the PPT algorithm except the following changes. After Step 2 of the PPT algorithm, the $2i-1$ th and $2i$ th, for $i = 0, \dots, p-1$, rows of the tridiagonal matrix Z are stored in the i th node (here the -1 th and $2(p-1)$ th rows of Z are assumed to be zero). Step 3 of the PPT algorithm, which performs total data-exchange, is eliminated. Step 4 of the PPT algorithm is then replaced by the PP method as described in Section II.

The PP method has two communication patterns in Step 4. The first communication pattern shown in Fig. 1 computes all partial matrix products. The second communication pattern is a *broadcast* which broadcasts the computed x_0 to all other nodes. Broadcast can be achieved by using the *spanning binomial tree* technique [15]. Considering second generation multicomputers, based on the communication model discussed in Section IV, the communication time required to obtain all partial products of B_i 's (Fig. 1) is $(\log p + 1)\alpha + 24(\log p + 1)\beta$. Here we have $S = 24$ because each message transfer has 6 data elements and we assume that each floating point number has 4 bytes. Since the broadcast communication takes $\log p$ steps and the message has one data element, it takes $(\log p)\alpha + (4 \log p)\beta$ time. The communication pattern of the PPH algorithm is shown in Fig. 4 for the case of $p = 8$. Counting the arithmetic operations and communication time, we have the following theorem.

Theorem 5.2: With $n = pm$, the PPH algorithm solves (1.1) in $17(n/p) + 20 \log p + 17$ and $24(n/p) + 20 \log p + 4$ parallel arithmetic operations for nonpivoting and pivoting, respectively. It requires $(2 \log p + 1)\alpha + (28 \log p + 24)\beta$ communication steps.

Let T_{SPH} be the time required to execute the PPH algorithm in a sequential processor and T_{PPH} be the time required to execute the PPH algorithm on p processors, we have

$$\begin{aligned} T_{SPH} &= (17n + 8p - 41)\tau_{comp} && \text{without pivoting} \\ T_{SPH} &= (24n - 5p - 41)\tau_{comp} && \text{with pivoting} \\ T_{PPH} &= \left(17\frac{n}{p} + 20 \log p + 17\right)\tau_{comp} + (2 \log p + 1)\alpha \\ &\quad + (28 \log p + 24)\beta && \text{without pivoting} \\ T_{PPH} &= \left(24\frac{n}{p} + 20 \log p + 4\right)\tau_{comp} + (2 \log p + 1)\alpha \\ &\quad + (28 \log p + 24)\beta. && \text{with pivoting} \end{aligned}$$

It is interesting to notice that the PPH algorithm can reach a speedup of 2 over the PP algorithm for $1 < p \ll n$. When $p = n$, no matrix partitioning is needed and the PPH algorithm is virtually the RCD algorithm. When $p = 1$, there is no conquer part and as the LU decomposition method is used in the dividing part, the algorithm becomes the LU decomposition algorithm.

As we mentioned earlier, the PP method is unstable if superdiagonal elements are small. This may be true for the PPH method. However, in this method, the PP algorithm is used only for solving a smaller linear system $Zy = h$, and half of superdiagonal elements of the matrix Z equal to 1 as we proved in Theorem 5.1. This indicates that the PPH method

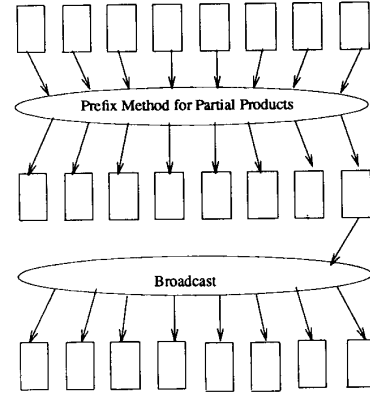


Fig. 4. The communication pattern of the PPH algorithm.

might be more stable than the PP method, but still inherits the instability of the PP method.

VI. THE PARALLEL DIAGONAL DOMINANT (PDD) ALGORITHM

Theorem 5.1 guarantees that the superdiagonal elements of matrix Z are nonzero if A has the same property. In practice, especially for a diagonal dominant tridiagonal matrix A , the magnitude of the last component of $v^{(i)}$ and the first component of $w^{(i)}$, which turn out to be the off-diagonal elements of the matrix Z , may be smaller than the machine accuracy when $p \ll n$. In this section we show that the magnitudes of $v_{m-1}^{(i)}$ and $w_0^{(i)}$ converge to zero when $m = n/p \rightarrow \infty$. An estimated rate of convergence will also be provided. Thus, when the tridiagonal system satisfies a strong diagonal dominance condition, a fast and highly efficient parallel algorithm gives an approximate solution which equals the exact solution within machine accuracy.

A. Theoretical Results for Diagonal Dominant Systems

For notational convenience, we write a tridiagonal matrix A as a triple of vectors as

$$A = \begin{bmatrix} b_0 & c_0 & & & & & \\ a_1 & b_1 & c_1 & & & & \\ & \cdot & \cdot & \cdot & & & \\ & & \cdot & a_{n-2} & b_{n-2} & c_{n-2} & \\ & & & & a_{n-1} & b_{n-1} & \end{bmatrix} = (a_i, b_i, c_i) \quad (i = 0, \dots, n-1)$$

where $a_0 = c_{n-1} = 0$. The matrix Z and the j th block of the matrix A can be similarly written as

$$\begin{aligned} Z &= (a_i^Z, b_i^Z, c_i^Z) \quad (i = 0, \dots, 2(p-1) - 1), \\ A_j &= (a_i^{(j)}, b_i^{(j)}, c_i^{(j)}) \quad (i = 0, \dots, m-1). \end{aligned}$$

The proof of Theorem 6.2 is a direct result of (6.3) and Lemma 6.1. It is not difficult to verify that

$$\frac{1}{\prod_{s=1}^m (\lambda_s + \epsilon)} \leq \frac{1}{(m+1)(1 + \epsilon/4)^m}. \quad (6.7)$$

Thus, the convergent rate of those elements is at least exponential if $\epsilon > 0$. Note that the bound in (6.7) is exact for $\epsilon = 0$ and would be more accurate for small ϵ . A sharper bound for larger ϵ is given by [3].

The fast convergence of these elements has a simple intuitive explanation. The solution of a tridiagonal system requires that every equation influences every other equation. For diagonal dominant systems, the influence diminishes with the distance between equations.

As $v_{m-1}^{(i)}, w_0^{(i)}$ ($0 \leq i \leq 2(p-1)$) become zero within machine accuracy, the matrix Z becomes a diagonal block matrix with each block of size 2×2 . In order to carry out the PDD algorithm, we have to guarantee that the 2×2 diagonal blocks of matrix Z

$$\begin{bmatrix} w_{m-1}^{(i)} & 1 \\ 1 & v_0^{(i+1)} \end{bmatrix}, \quad i = 0, \dots, p-2.$$

are nonsingular. This is proved by the following theorem and its proof is stated in the Appendix.

Theorem 6.3: If a real tridiagonal matrix A of order n satisfies condition (6.1), then

$$|w_{m-1}^{(i)} \cdot v_0^{(i+1)}| < 1, \quad i = 0, \dots, p-2. \quad (6.8)$$

B. The Parallel Diagonal Dominant (PDD) Algorithm

Theorem 6.2 implies that the matrix Z is a diagonal block matrix with block size 2×2 within machine accuracy if m is large enough. Then the procedure for solving (3.9), the conquer part of our algorithm, becomes very simple. If (1.1) satisfies condition (6.1), the parallel diagonal dominant (PDD) algorithm has the following steps modified from the PPT algorithm.

- 3) Send $\hat{x}_0^{(i)}, v_0^{(i)}$ from the i th node to the $(i-1)$ th node for $i = 1, \dots, p-1$.
- 4) Solve

$$\begin{bmatrix} w_{m-1}^{(i)} & 1 \\ 1 & v_0^{(i+1)} \end{bmatrix} \begin{bmatrix} y_{2i} \\ y_{2i+1} \end{bmatrix} = \begin{bmatrix} \hat{x}_{m-1}^{(i)} \\ \hat{x}_0^{(i+1)} \end{bmatrix}$$

in parallel on all i th nodes for ($0 \leq i \leq p-2$). Then send y_{2i+1} from the i th node to the $(i+1)$ th node for $i = 0, \dots, p-2$.

Fig. 5 depicts the communication pattern of the PDD algorithm. In Step 3, all nodes transfer a message to their left neighbor. In Step 4, all nodes transfer a message to their right neighbor. Since it is unnecessary for node 0 to communicate with node $p-1$, the algorithm only needs a one dimensional array connection, which can be perfectly embedded in first generation hypercube multicomputers. The optimal and simple communication property makes the PDD algorithm an ideal




Fig. 5. The communication pattern of the PDD algorithm.

algorithm for multicomputers. The computation and communication complexities of the PDD algorithm is stated in the following theorem.

Theorem 6.4: If $n/p \gg 1$ and matrix A satisfies condition (6.1), the PDD algorithm solves (1.1) in $17(n/p) - 14$ parallel arithmetic operations and $2\alpha + 12\beta$ communications on multicomputers with p processors. When the left-hand side of (6.7) is within machine accuracy, the approximate solution equals the exact solution within machine accuracy.

Let T_{SDD} and T_{PDD} be the time required to execute the PDD algorithm in a sequential processor and p parallel processors, respectively. We have

$$T_{SDD} = (17n - 14p - 8)\tau_{comp},$$

$$T_{PDD} = \left(17\frac{n}{p} - 14\right)\tau_{comp} + 2\alpha + 12\beta.$$

Many tridiagonal systems arising in science and engineering problems satisfy condition (6.1). However, condition (6.1) is sufficient but not necessary. The algorithm can be used for any positive definite or diagonal dominant matrix whose resulting $|v_{m-1}^{(j)}|, |w_0^{(j)}|$ [see (6.3)] become underflow as $n \gg p$. One of the sufficient conditions is that the determinants of scaled matrices A_j' ($j = 0, \dots, p-1$) become overflow as $n \gg p$. For example, if $A = (1, i, 1)$, $\det(A_j')$ are approximately larger than $m!$ since its i th eigenvalue $\lambda_i \approx i$ except few smallest and largest ones. The magnitude of those entries of Z can be as small as 10^{-64} for $n/p \geq 50$.

VII. EXPERIMENTAL RESULTS

The arithmetic operation counts and communication steps of each individual algorithm are summarized in Table I.

Fig. 6 shows the estimated and measured speedup of the PP, PPT, PPH, and PDD algorithms with respect to the subroutine SGTSL of LINPACK [4]. These algorithms are implemented on a 64-node nCUBE 1 multicomputer. nCUBE 1 is a first generation multicomputer adopting the store-and-forward communication mechanism. In our nCUBE 1 machine, the following system parameters are measured: $\alpha = 5.0$, $\beta = 0.013$, and $\tau_{comp} = 0.08$ (without normalization). The dimension of matrix A is chosen as $n = 6400$. This value is limited by the memory constraint of individual processors. As p increases, the PPH algorithm will outperform the PPT algorithm because the dimension of the matrix Z [(3.9)] increases as p increases, which favors the parallel approach used in the PPH algorithm. The performance of the PPH algorithm seems to be underestimated compared with the

TABLE I
COMPUTATION AND COMMUNICATION COUNTS OF TRIDIAGONAL SOLVERS

Algorithm	Computation	Communication
PPT (nonpivoting)	$17 \frac{n}{p} + 16p - 45$	$(\log p)\alpha + 16(p-1)\beta$
PPT (pivoting)	$24 \frac{n}{p} + 22p - 69$	$(\log p)\alpha + 16(p-1)\beta$
PPH (nonpivoting)	$17 \frac{n}{p} + 20 \log p + 17$	$(2 \log p + 1)\alpha + (28 \log p + 24)\beta$
PPH (pivoting)	$24 \frac{n}{p} + 20 \log p + 4$	$(2 \log p + 1)\alpha + (28 \log p + 24)\beta$
PDD	$17 \frac{n}{p} - 14$	$2\alpha + 12\beta$




Fig. 6. The speedup over the LU decomposition method, where $n = 6400$. (a) Estimated speedup. (b) Measured speedup.




Fig. 7. Efficiency over the LU decomposition method, $n = 6400$. (a) Estimated efficiency. (b) Measured efficiency.

measured results. This is mainly caused by assuming a dilation of 2 for all communications occurred in Fig. 1. However, in actual implementation, some communications have dilation 1.

As n goes to infinity, the asymptotic speedups, compared with the LU decomposition method, of all our methods are $0.471p$. And the asymptotic speedup for the PP method is $0.229p$. Dividing the speedup by the number of processors, p , Fig. 7 shows the *efficiencies* of our methods. For the case of $n = 6400$, only the PDD algorithm has an efficiency closing to the asymptotic efficiency. For all methods, the efficiency decreases as the number of processors p increases. This is mainly caused by the increasing ratio of the communication and the computation complexities.

There are two commonly accepted measures for the speedup and the efficiency of parallel algorithms [18]. One focuses on how much faster a problem can be solved by p processors. Thus, it compares the best serial algorithm with the parallel algorithm under consideration. It is defined as

$$S_p = \frac{\text{execution time using the fastest sequential algorithm on one processor}}{\text{execution time using the parallel algorithm on } p \text{ processors}}$$

Both Figs. 6 and 7 are based on the above measure, and the best sequential algorithm chosen is LINPACK subroutine, SGTSL. Another measure interests in the parallel inheritance of the algorithm and is defined as

$$S'_p = \frac{\text{execution time using one processor}}{\text{execution time using } p \text{ processors}}$$

We call the latter as the *relative speedup*. The relative speedups




Fig. 8. The performance of relative speedup and relative efficiency, where $n = 6400$. (a) Estimated relative speedup. (b) Measured relative speedup.

and efficiencies of our methods are shown in Fig. 8. As it shows, the relative speedup of the PDD algorithm is very close to the ideal speedup. This is due to the high efficiency of the PDD algorithm as we discussed in Section VI-A.

One of the advantages of our methods is its diversity. In the dividing part, the pivoting may be used when it is necessary. There has been a tacit assumption for most available parallel algorithms that no pivoting is required. Thus, our algorithms may be more stable than others in the cases pivoting is required. In the conquer part, basically the major computation required is to solve (3.9). The methods used in this part are irrelevant to the methods used in the dividing part. Other strategies may also be used to solve (3.9).

The nonsingularity of matrix \tilde{A} is required by the PPT and PPH algorithms. Unless the matrix A is positive definite or diagonal dominant, there is no general rule to guarantee

- [10] R. W. Hockney, "A fast direct solution of Poisson's equation using Fourier analysis," *J. ACM*, vol. 12, pp. 95-113, 1965.
- [11] C. T. Ho and S. L. Johnsson, "Optimizing tridiagonal solvers for alternating direction methods on Boolean cube multiprocessors," *SIAM J. Sci. Statist. Comput.*, vol. 11, no. 3, pp. 563-592, May 1990.
- [12] K. Hwang, "Advanced parallel processing with supercomputer architectures," *Proc. IEEE*, pp. 33-47, Oct. 1987.
- [13] S. L. Johnsson, "Solving tridiagonal systems on ensemble architectures," *SIAM J. Sci. Statist. Comput.*, vol. 8, no. 3, pp. 354-392, May 1987.
- [14] ———, "Communication efficient basic linear computations on hypercube multiprocessors," *J. Parallel Distributed Comput.*, no. 4, pp. 133-172, 1987.
- [15] S. L. Johnsson and C. T. Ho, "Spanning graphs for optimum broadcasting and personalized communication in hypercubes," *IEEE Trans. Comput.*, vol. 38, Sept. 1989.
- [16] D. H. Lawrie and A. H. Sameh, "The computation and communication complexity of a parallel banded system solver," *ACM Trans. Math. Software*, vol. 10, no. 2, pp. 185-195, June 1984.
- [17] P. H. Michielse and H. A. Vorst, "Data transport in Wang's partition method," in *Parallel Computing*, 7. New York: North-Holland, 1988, pp. 87-95.
- [18] J. M. Ortega and R. G. Voigt, "Solution of partial differential equations on vector and parallel computers," *SIAM Rev.*, pp. 149-240, June 1985.
- [19] Y. Saad and M. H. Schultz, "Data communication in hypercube," Res. Rep., YALEU/DCS/RR-428, Oct. 1985.
- [20] J. Sherman and W. J. Morrison, "Adjustment of an inverse matrix corresponding to changes in the elements of a given column or a given row of the original matrix," *Ann. Math. Statist.*, vol. 20, 621, 1949.
- [21] G. D. Smith, *Numerical Solution of Partial Differential Equations*. Oxford, England: Oxford University Press, 1985.
- [22] H. S. Stone, "An efficient parallel algorithm for the solution of a tridiagonal linear system of equations," *J. ACM*, vol. 20, no. 1, pp. 27-38, Jan. 1973.
- [23] H. H. Wang, "A parallel method for tridiagonal equations," *ACM Trans. Math. Software*, vol. 7, pp. 170-183, June 1981.
- [24] M. Woodbury, "Inverting modified matrices," Memo. 42, Statistics Research Group, Princeton, NJ, 1950.
- [25] H. Zhang, "On the accuracy of the parallel diagonal dominant algorithm," *Parallel Comput.*, pp. 265-272, 1991.



Xian-He Sun (S'88-M'90) received the B.S. degree in mathematics from Beijing Normal University, Beijing, China, in 1982 and the M.S. and Ph.D. degrees in computer science from Michigan State University, East Lansing, MI, in 1987 and 1990, respectively.

He is currently a postdoctoral fellow at Ames Laboratory, Ames, IA. His research interests include parallel processing, performance evaluation, parallel numerical algorithms, and database systems.

Dr. Sun is a member of the IEEE Computer Society and the Association for Computing Machinery.



Hong Zhang received the B.S. degree in mathematics from Beijing Normal University, Beijing, China, in 1982 and the M.S. and Ph.D. degrees in mathematics from Michigan State University in 1985 and 1989, respectively.

She is currently an Assistant Professor at Clemson University, Clemson, SC. Her research interests include numerical and scientific computation. Her current work is on the designing of parallel algorithms for solution of linear algebraic systems and eigenvalue problems.



Lionel M. Ni (S'78-M'81-SM'87) received the B.S. degree in electrical engineering from National Taiwan University in 1973, the M.S. degree in electrical and computer engineering from Wayne State University, Detroit, MI, in 1977, and the Ph.D. degree in electrical engineering from Purdue University, West Lafayette, IN, in 1980.

In 1981 he joined the faculty of the Department of Computer Science, Michigan State University, East Lansing, where he is currently a Professor and Director of the Advanced Computer Systems Laboratory. During the summers of 1979 and 1981, he was a Researcher at the IBM San Jose Research Laboratories. During the year 1987-1988, he was a Visiting Scientist in the Division of Mathematics and Computer Science at Argonne National Laboratory. He is an editor of the *Journal of Parallel and Distributed Computing*. His research interests include parallel processing and distributed computing.

Dr. Ni served as a Distinguished Visitor of the IEEE Computer Society from 1985 to 1988. He is a member of the Association for Computing Machinery and SIAM, and served as the program chairman of the Fifteenth Annual International Computer Software and Applications Conference.