

Engineering Physics and Mathematics Division

Mathematical Sciences Section

**PARALLEL MATRIX TRANSPOSE ALGORITHMS  
ON DISTRIBUTED MEMORY CONCURRENT COMPUTERS**

Jaeyoung Choi §  
Jack J. Dongarra §†  
David W. Walker §

§ Mathematical Sciences Section  
Oak Ridge National Laboratory  
P.O. Box 2008, Bldg. 6012  
Oak Ridge, TN 37831-6367

† Department of Computer Science  
University of Tennessee at Knoxville  
107 Ayres Hall  
Knoxville, TN 37996-1301

Date Published: October 1993

Research was supported by the Applied Mathematical Sciences Research Program of the Office of Energy Research, U.S. Department of Energy, by the Defense Advanced Research Projects Agency under contract DAAL03-91-C-0047, administered by the Army Research Office, and by the Center for Research on Parallel Computing

Prepared by the  
Oak Ridge National Laboratory  
Oak Ridge, Tennessee 37831  
managed by  
Martin Marietta Energy Systems, Inc.  
for the  
U.S. DEPARTMENT OF ENERGY  
under Contract No. DE-AC05-84OR21400

**MASTER**

DISTRIBUTION OF THIS DOCUMENT IS UNLIMITED *p-2*

## Contents

1	Introduction . . . . .	1
2	Design Issues . . . . .	2
3	Matrix Transpose Algorithms . . . . .	5
3.1	$P$ and $Q$ : relatively prime . . . . .	7
3.2	$P$ and $Q$ : not relatively prime . . . . .	9
4	Results . . . . .	13
5	Conclusions and Remarks . . . . .	18
6	References . . . . .	18

# PARALLEL MATRIX TRANSPOSE ALGORITHMS ON DISTRIBUTED MEMORY CONCURRENT COMPUTERS

Jaeyoung Choi

Jack J. Dongarra

David W. Walker

## Abstract

This paper describes parallel matrix transpose algorithms on distributed memory concurrent processors. We assume that the matrix is distributed over a  $P \times Q$  processor template with a block scattered data distribution.  $P$ ,  $Q$ , and the block size can be arbitrary, so the algorithms have wide applicability.

The communication schemes of the algorithms are determined by the greatest common divisor ( $GCD$ ) of  $P$  and  $Q$ . If  $P$  and  $Q$  are relatively prime, the matrix transpose algorithm involves *complete exchange* communication. If  $P$  and  $Q$  are not relatively prime, processors are divided into  $GCD$  groups and the communication operations are overlapped for different groups of processors. Processors transpose  $GCD$  wrapped diagonal blocks simultaneously, and the matrix can be transposed with  $LCM/GCD$  steps, where  $LCM$  is the least common multiple of  $P$  and  $Q$ .

The algorithms make use of non-blocking, point-to-point communication between processors. The use of nonblocking communication allows a processor to overlap the messages that it sends to different processors, thereby avoiding unnecessary synchronization.

Combined with the matrix multiplication routine,  $C = A \cdot B$ , the algorithms are used to compute parallel multiplications of transposed matrices,  $C = A^T \cdot B^T$ , in the PUMMA package [5]. Details of the parallel implementation of the algorithms are given, and results are presented for runs on the Intel Touchstone Delta computer.

## 1. Introduction

Matrix transposition is a fundamental matrix operation of linear algebra [8,14] and arises in many scientific and engineering applications. On a uniprocessor, an algorithm involving a transposed matrix may not actually require the matrix data to be transposed in physical memory. Instead, it may be accessed simply by exchanging the row and column indices. However, in a distributed-memory multiprocessor environment, we cannot simply interchange the global row and column indices. Instead, the data must be physically moved from one processor to another.

Transposition of a matrix is a redistribution of its elements. Many researchers have considered the problem for different architectures. In 1972, Eklundh [7] considered the problem of directly accessing rows or columns of a matrix when its size is larger than the available high-speed storage. O'Leary [12] implemented an algorithm for transposing an  $N \times N$  matrix on a one-dimensional systolic array. Azari, Bojanczyk and Lee [1] developed an algorithm for transposing an  $M \times N$  matrix on an  $N \times N$  mesh-connected array processor, and Johnsson and Ho [10] presented an algorithm for a Boolean n-cube, or hypercube.

Current advanced architecture computers possess hierarchical memories in which accesses to data in the upper levels of the memory hierarchy (registers, cache, and/or local memory) are faster than those in lower levels (shared or off-processor memory). To exploit the power of such machines, block-partitioned algorithms are preferred for dense linear algebra computations, in which operations are performed on submatrices, rather than individual matrix elements. In distributing matrix data over processors we therefore assume a block scattered decomposition [4,6]. The block scattered decomposition can reproduce the most common data distributions used in dense linear algebra, as described briefly in the next section.

In this paper, the parallel matrix transpose algorithms are presented based on the block scattered decomposition. The algorithms are implemented on the Intel Touchstone Delta computer. The communication schemes of the algorithms are determined by the greatest common divisor ( $GCD$ ) of the number of rows and columns ( $P$  and  $Q$ ) of the processor template. If  $P$  and  $Q$  are relatively prime, the matrix transpose algorithm involves *complete exchange* communication. This is called all-to-all personalized communication, in which each of  $N_p = P \cdot Q$  processors is required to send distinct subblocks to each of the remaining  $N_p - 1$  processors, and receive distinct subblocks from each of them. Bokhari and Berryman [2] have developed binary exchange and quadrant exchange algorithms on a circuit switched mesh, where  $P$  and  $Q$  are powers of 2. The complete exchange communication in our transpose algorithms arises for any processor configuration, and is not limited to the case where  $P$  and  $Q$  are powers of 2. We implemented the complicated two-dimensional complete exchange communication problem by generalizing the one-dimensional complete exchange communication based on direct point-to-point communication. Details are discussed in Section 3.1.

We have presented the Parallel Universal Matrix Multiplication Algorithms (PUMMA) in [5] for performing  $C \leftarrow \alpha op(A) \cdot op(B) + \beta C$ , where  $op(X) = X$  or  $X^T$ , based on the block scattered decomposition. One of the cases in the PUMMA package,  $C \leftarrow \alpha A^T \cdot B^T + \beta C$ , is implemented in two steps ( $T \leftarrow \alpha B \cdot A$ ;  $C \leftarrow T^T + \beta C$ ). The second step involves parallel matrix transposition. The performance of this algorithm for evaluating  $C = A^T \cdot B^T$  is compared with the algorithm for evaluating  $C = A \cdot B$  on the Intel Delta machine in Section 4.

## 2. Design Issues

The way in which an algorithm's data are distributed over the processors of a concurrent computer has a major impact on the load balance and communication characteristics of the concurrent algorithm, and hence largely determines its performance and scalability. The block scattered decomposition provides a simple, yet general-purpose way of distributing a block-partitioned matrix on distributed memory concurrent computers. In the block scattered decomposition, described in detail in [4,6], an  $M \times N$  matrix is partitioned into blocks of size  $r \times s$ , and blocks separated by a fixed stride in the column and row directions are assigned to the same processor. If the stride in the column and row directions is  $P$  and  $Q$  blocks respectively, then we require that  $P \cdot Q$  equal the number of processors,  $N_p$ . Thus, it is useful to imagine the processors arranged as a  $P \times Q$  mesh, or template. The processor at position  $(p, q)$  ( $0 \leq p < P$ ,  $0 \leq q < Q$ ) in the template is assigned the blocks indexed by,

$$(p + i \cdot P, q + j \cdot Q), \quad (1)$$

where  $i = 0, \dots, [(M_b - p - 1)/P]$ ,  $j = 0, \dots, [(N_b - q - 1)/Q]$ , and  $M_b \times N_b$  is the size in blocks of the matrix ( $M_b = \lceil M/r \rceil$ ,  $N_b = \lceil N/s \rceil$ ).

Blocks are scattered in this way so that good load balance can be maintained in parallel algorithms, such as LU factorization [3,6]. The nonscattered decomposition (or pure block distribution) is just a special case of the scattered decomposition in which the block size is given by  $r = \lceil M/P \rceil$  and  $s = \lceil N/Q \rceil$ . A purely scattered decomposition (or two-dimensional wrapped distribution) is another special case in which the block size is given by  $r = s = 1$ .

If  $P$  and  $Q$  are relatively prime, the matrix transpose algorithm involves a two-dimensional complete exchange communication, where each of  $N_p$  processors is required to send distinct subblocks to each of the remaining  $N_p - 1$  processors, and receive distinct subblocks from each of them. We implemented the complicated two-dimensional complete exchange algorithm by generalizing the one-dimensional complete exchange algorithm. Three one-dimensional complete exchange communication schemes are shown in Figure 1, where each processor needs one subblock from each other processor, and the number in parentheses denotes the number of




Figure 1: Three complete exchange communication schemes on 8 processors. The number in parentheses denotes the amount of data to transmit.

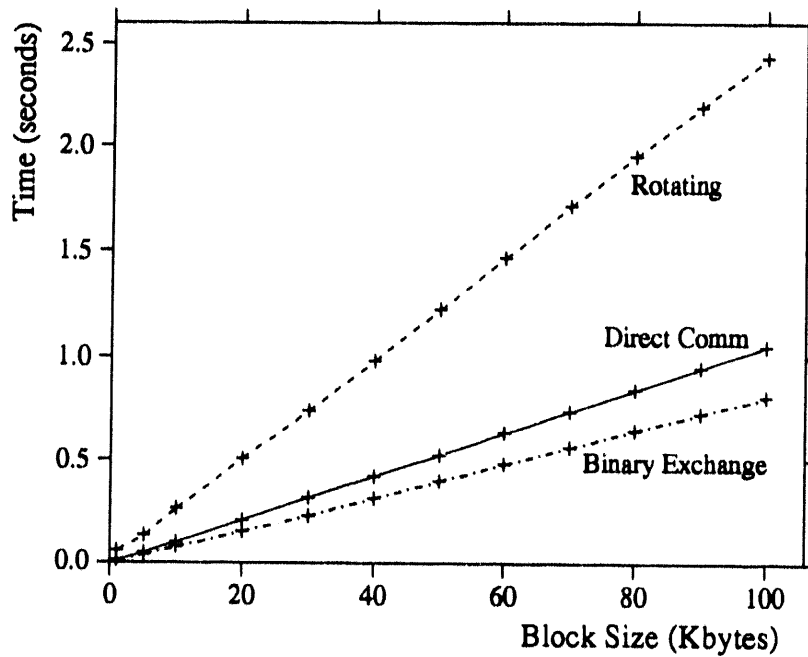


Figure 2: Comparison of three exchange communication schemes on 16 processors.

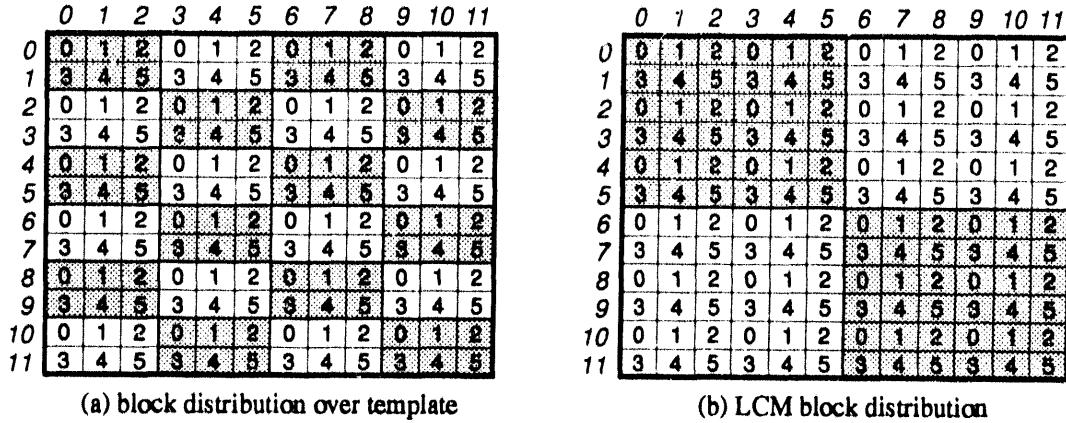
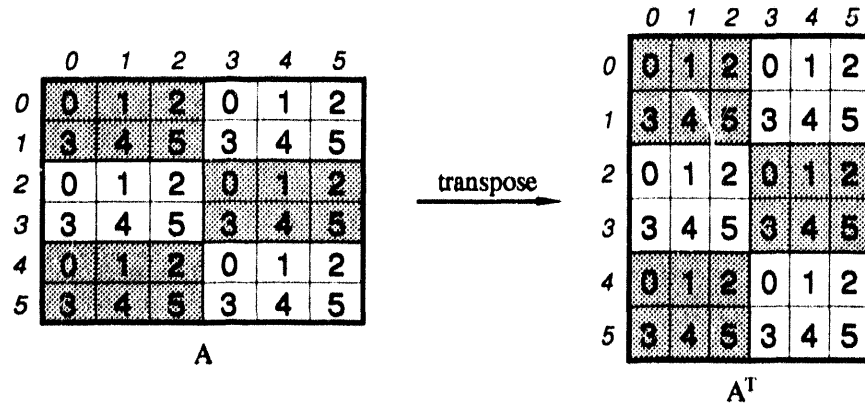


Figure 3: A matrix with  $12 \times 12$  blocks is distributed over a  $2 \times 3$  processor template. (a) Each shaded and unshaded area represents different templates. The numbered squares represent blocks of elements, and the number indicates at which location in the processor template the block is stored - all blocks labeled with the same number are stored in the same processor. The *slanted* numbers, on the left and on the top of the matrix, represent indices of row block and column block, respectively. (b) The matrix has  $2 \times 2$  LCM blocks. Blocks belong to the same processor if the relative locations of blocks are the same in each square LCM block. The definition of the LCM block is defined in the text.


subblocks to transmit.

The binary exchange scheme completes in  $\lceil \log_2 P \rceil$  steps and the amount of data transmitted in each step is fixed at  $2^{\lceil \log_2 P \rceil - 1}$  subblocks, where  $P$  is the number of processors. The rotating scheme can avoid network congestion, but requires  $P - 1$  steps and the amount of data transmitted in the initial steps is large. In the direct point-to-point communication scheme, the number of steps is the same in the rotating scheme, but the amount of data transmitted in each step is only one subblock.

The three schemes have been implemented on 16 nodes of the Delta and their performances are compared in Figure 2. The binary exchange and the rotating schemes are implemented with an odd-even communication scheme, which is preferable to a simultaneous communication scheme on the Delta [5,11]. In this algorithm, odd-numbered processors send their own blocks and even-numbered processors receive them in the first step, and even-numbered processors send and odd-numbered processors receive in the next step. On  $P = 2^d$  processors, as shown in Figure 2, the binary exchange scheme is the fastest. However, if  $P$  is not a power of 2, then this scheme becomes very complicated and may be slower than the direct communication scheme. The direct communication scheme is about 20% slower than the binary exchange scheme for the worst case ( $P = 2^d$ ), but it is simple to implement on an arbitrary number of processors. We adopted the simple direct communication scheme for the implementation of the matrix transpose algorithms, which are explained in detail in the next section.



(a) matrix transpose from matrix point-of-view




(b) matrix transpose from processor point-of-view

Figure 4: An example of matrix transpose for a block scattered decomposition, when  $P = 2$ ,  $Q = 3$ , and  $M_b = N_b = 6$ .


### 3. Matrix Transpose Algorithms

We assume that a matrix is distributed over a two-dimensional processor mesh, or template, so that in general each processor has several blocks of the matrix as shown in Figure 3 (a), where a matrix with  $12 \times 12$  blocks is distributed over a  $2 \times 3$  template. Denoting the least common multiple of  $P$  and  $Q$  by  $LCM$ , we refer to a square of  $LCM \times LCM$  blocks as an  $LCM$  block. Thus, the matrix may be viewed as a  $2 \times 2$  array of  $LCM$  blocks, as shown in Figure 3 (b). The concept of the  $LCM$  block was introduced in [5], and is very useful for implementing algorithms that use a block scattered data distribution. Blocks belong to the same processor if their relative locations are the same in each square  $LCM$  block. An algorithm may be developed for the first  $LCM$  block, and then it can be directly applied to the other  $LCM$  blocks, which all have the same structure and the same data distribution as the first  $LCM$  block. That is, when an operation is executed on a block of the first  $LCM$  block, the same operation can be done simultaneously on other blocks, which have the same relative location in each  $LCM$  block.





(a) matrix transpose from matrix point-of-view



(b) matrix transpose from processor point-of-view

Figure 5: An example of matrix transpose for a block scattered decomposition, when  $P = 3$ ,  $Q = 3$ , and  $M_b = N_b = 6$ .

We now describe parallel matrix transpose algorithms. A matrix  $A$ , distributed over a  $P \times Q$  processor template, has  $M_b \times N_b$  blocks and each block consists of  $r \times s$  elements, where  $r$  and  $s$  are arbitrary. Figure 4 (a) shows an example of a matrix transpose on a  $2 \times 3$  template. If  $A$  is transposed, the transposed matrix  $A^T$  is distributed over the same  $P \times Q$  template, and it has  $N_b \times M_b$  blocks and each block has  $s \times r$  elements. The elements of each block remain in the same block, but may be in a different processor, and each block is itself transposed. Figure 4 (b) shows the same example from the processor point-of-view. If  $P$  and  $Q$  are relatively prime, as shown in the figure, blocks in the first processor  $P_0$  are scattered to all processors. As shown in Figure 5, which is the same example on a  $3 \times 3$  square template, the blocks in each processor are not dispersed, but they are moved as one entity to a different processor. Parallel matrix transpose algorithms for the block scattered data distribution have several communication patterns determined by the greatest common divisor ( $GCD$ ) of  $P$  and  $Q$ .

```

DO J = 0, Q - 1
  DO I = 0, P - 1
    [ Copy all blocks of A required by P(p + I, q - J) to T1
      (in condensed and transposed form) ]
    [ Send T1 to P(p + I, q - J) ]
    [ Receive T2 from P(p - I, q + J) ]
    [ Copy T2 to C ]
  END DO
END DO

```

Figure 6: A parallel matrix transpose algorithm from the processor point-of-view, when  $P$  and  $Q$  are relatively prime.  $P(p, q)$  represents  $P_{\text{MOD}(p, P), \text{MOD}(q, Q)}$ . Processor  $P_{p, q}$  ( $0 \leq p < P$  and  $0 \leq q < Q$ ) communicates with  $P(p + I, q - J)$  to send, and  $P(p - I, q + J)$  to receive based on direct point-to-point communication. ' $p + I$ ' and ' $q - J$ ' can be replaced with a different combination of signs.


### 3.1. $P$ and $Q$ : relatively prime

We start with the simple case where  $P$  and  $Q$  are relatively prime, i. e.  $GCD = 1$ . In this case blocks in  $P_0$  are scattered to all processors after being locally transposed as shown in Figure 4 (b). This case involves the two-dimensional complete exchange communication. That is, every processor needs to communicate with every other processor. The complete exchange problem is implemented by direct communication between sender and receiver.


Figure 6 shows the pseudocode from the processor point-of-view, where  $P(p, q)$  represents  $P_{\text{MOD}(p, P), \text{MOD}(q, Q)}$  in the processor template. Processor  $P(p, q)$  ( $0 \leq p < P$  and  $0 \leq q < Q$ ) starts to transpose blocks whose transposed blocks belong to itself. Then it deals with blocks whose transposition are in processors in the same column of the template ( $P(p - i, q)$ ,  $0 \leq i < P$ ). The processor sends blocks to its top neighbor,  $P(p - 1, q)$ , and receives blocks from its bottom neighbor,  $P(p + 1, q)$ . Before sending the blocks, it is necessary to copy the blocks to be sent into a contiguous message buffer. Next it sends blocks to the next top processor,  $P(p - 2, q)$ , and receives blocks from the next bottom processor,  $P(p + 2, q)$ .

After it completes its operations with the processors in the same column, it sends blocks to the processors to the left in the template ( $P(p - i, q - 1)$ ,  $0 \leq i < P$ ), and receives blocks from the processors to the right ( $P(p + i, q + 1)$ ). All operations are completed in  $P \times Q = LCM$  steps.


We interpret the algorithm from the matrix point-of-view. In the first  $LCM$  block, the above algorithm performs the operation by transposing one (wrapped) diagonal blocks at one step. (Actually the algorithm transposes every  $LCM$  diagonal blocks of the matrix at each step.) The first step of the algorithm in Figure 6 requires no explicit communication. It corresponds




(a) zeroth diagonal ( $A(i,j)$ ,  $\text{MOD}(j-i, \text{LCM})=0$ )




(b) third diagonal ( $A(i,j)$ ,  $\text{MOD}(j-i, \text{LCM})=3$ )




(c) fourth diagonal ( $A(i,j)$ ,  $\text{MOD}(j-i, \text{LCM})=4$ )



(d) first diagonal ( $A(i,j)$ ,  $\text{MOD}(j-i, \text{LCM})=1$ )



(e) second diagonal ( $A(i,j)$ ,  $\text{MOD}(j-i, \text{LCM})=2$ )



(f) fifth diagonal ( $A(i,j)$ ,  $\text{MOD}(j-i, \text{LCM})=5$ )

Figure 7: Snapshots of matrix transposition when  $P = 2$ ,  $Q = 3$  and  $M_b = N_b = 6$ . The small *slanted* number in the upper left corner in each block represents processor identification number. One wrapped block diagonal is transposed in each step. The darkly shaded area represents blocks to be shifted, and lightly shaded area stands for their transpositions.

```

DO J = 0, Q - 1
  K = J      \* Determine K-th diagonal block to transpose *\
  WHILE (MOD(K, P) ≠ 0) DO K = MOD(K + Q, LCM) END DO
  DO I = 0, P - 1
    [ Copy every (K : Nb : LCM)-th wrapped diagonal blocks in Pp,q to T1 ]
    [ Move T1 from Pp,q to P(p + I, q - J) ]
    [ Copy the received T1 to C ]
    K = MOD(K + Q, LCM)
  END DO
END DO

```

Figure 8: A parallel matrix transpose algorithm from the matrix point-of-view, when  $P$  and  $Q$  are relatively prime. One diagonal block is transposed at one step. The 'While' statement should be executed until  $\text{MOD}(K, P)$  becomes 0. ( $start : limit : intv$ ) represents values of  $x$ , where  $x = start + intv \cdot y$ ,  $y = 0, 1, \dots$ , and  $x$  can't exceed  $limit$ .

to an in-place transpose of the diagonal blocks of  $\mathbf{A}$  ( $\mathbf{A}(i, i)$ ) (See Fig. 7(a)). Then every  $P$  diagonal blocks of  $\mathbf{A}$  ( $\mathbf{A}(i, j)$ ,  $\text{MOD}(j - i, P) = 0$ ) (See Fig. 7(b)) are transposed in the first outer loop ( $J = 0$ ) of Figure 6. In the next outer loop ( $J = 1$ ), the next  $P$  diagonal blocks ( $\mathbf{A}(i, j)$ ,  $\text{MOD}(j - i, P) = 1$ ) are transposed. In Figures 7 (c) and (d),  $P_0$  ( $P(0, 0)$ ) sends blocks to  $P_2$  ( $P(0, 2)$ ), and receives from  $P_1$  ( $P(0, 1)$ ), where  $P_0$ ,  $P_1$  and  $P_2$  are in the same row. Then  $P_0$  sends blocks to  $P_5$  ( $P(1, 2)$ ), and receives from  $P_4$  ( $P(1, 1)$ ), and so on. The pseudocode for the algorithm from the matrix point-of-view is shown in Figure 8. Processors need to determine a diagonal block of  $\mathbf{A}$  ( $\mathbf{A}(i, j)$ ,  $\text{MOD}(j - i, LCM) = K$ ) which they start to transpose in the outer  $J$  loop in order to communicate with other processors in the same row of the template. The two lines before the inner DO-loop compute the value of  $K$ .

### 3.2. $P$ and $Q$ : not relatively prime

In the previous section, we have investigated the case when  $P$  and  $Q$  are relatively prime, which involves complete exchange communication. In this section we consider the case of  $GCD > 1$ . The former algorithm is a special case ( $GCD = 1$ ) of this algorithm.

Figure 9 shows an example of transposing a  $12 \times 12$  matrix on a  $4 \times 6$  template from the processor point-of-view. Each processor has its own  $3 \times 2$  ( $= LCM/P \times LCM/Q$ ) array of blocks. The processors can transpose the matrix with 6 ( $= LCM/P \cdot LCM/Q = LCM/GCD$ ) communications steps. As shown in Figure 10, a processor  $P(p, q)$  starts to communicate with  $P(\tilde{p}, \tilde{q})$ , where  $\tilde{p}$  and  $\tilde{q}$  are computed from  $p$  and  $q$  (details are explained later of this section). Once  $P(\tilde{p}, \tilde{q})$  is determined, it communicates with other processors, whose vertical and horizontal intervals are  $GCD$  from  $P(\tilde{p}, \tilde{q})$ . The two loops of the algorithm in Figure 6 are




Figure 9: A matrix transpose example on a  $4 \times 6$  template.

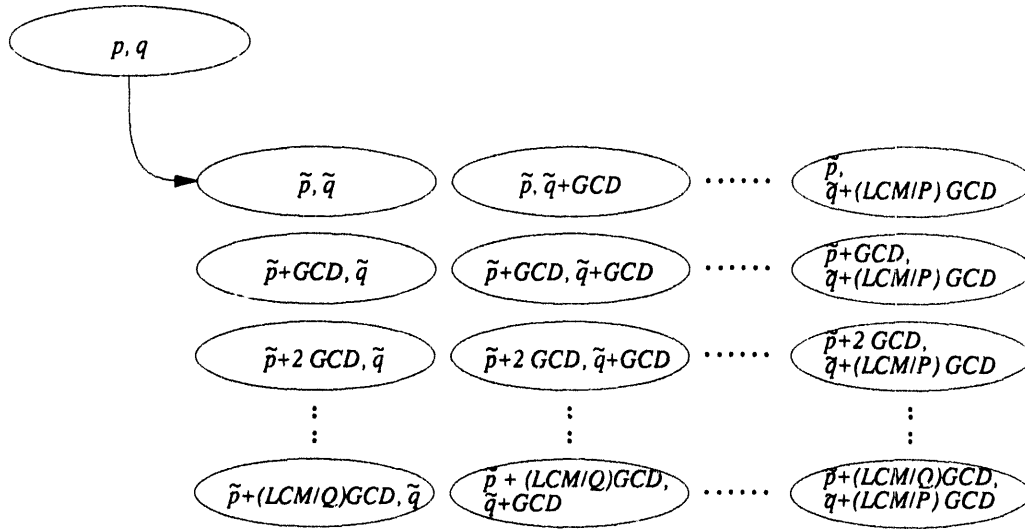


Figure 10: Processor map for communication. A processor  $P\langle p, q \rangle$  starts to communicate with  $P\langle \tilde{p}, \tilde{q} \rangle$ , then it communicates with other processors, whose vertical and horizontal intervals are  $GCD$  from  $P\langle \tilde{p}, \tilde{q} \rangle$ .

```

PARDO  $K = 1, GCD$ 
   $g = \text{MOD}(q - p, GCD)$ 
   $\tilde{p} = \text{MOD}(p + g, P)$ ;  $\tilde{q} = \text{MOD}(q - g, Q)$ 
  DO  $J = 0, LCM/P - 1$ 
    DO  $I = 0, LCM/Q - 1$ 
      [ Copy to T1 (in condensed and transposed form) all blocks of A
        required by  $P(\tilde{p} + I \times GCD, \tilde{q} - J \times GCD)$  ]
      [ Send T1 to  $P(\tilde{p} + I \times GCD, \tilde{q} - J \times GCD)$  ]
      [ Receive T2 from  $P(\tilde{p} - I \times GCD, \tilde{q} + J \times GCD)$  ]
      [ Copy T2 to C ]
    END DO
  END DO
END PARDO

```

Figure 11: A modified matrix transpose algorithm from the processor point-of-view. Operations of  $GCD$  groups of processors are overlapped.




Figure 12: Two snapshots of matrix transposition for transposing the zeroth and first wrapped block diagonals, when  $P = 4$ ,  $Q = 6$  and  $M_b = N_b = 12$ . In this example, transposing of even numbered wrapped block diagonals can be overlapped with that of odd numbered.




Figure 13: Matrix transposition when  $P = Q = GCD = 3$ . Processors transpose 3 ( $= GCD$ ) diagonal blocks at one step, so that the transposition is done in one step.

```

PARDO  $K = 1, GCD$ 
   $g = MOD(q - p, GCD)$ 
   $\tilde{p} = MOD(p + g, P); \tilde{q} = MOD(q - g, Q)$ 
  DO  $J = 0, LCM/P - 1$ 
     $K = J$       \* Determine  $K$ -th diagonal block to transpose *
    WHILE ( $MOD(K - g, P) \neq 0$ ) DO  $K = MOD(K + Q, LCM)$  END DO
    DO  $I = 0, LCM/Q - 1$ 
      [ Copy every  $(K : N_b : LCM)$ -th diagonal blocks in  $P(p, q)$  to T1 ]
      [ Move T1 from  $P(p, q)$  to  $P(\tilde{p} + I \times GCD, \tilde{q} - J \times GCD)$  ]
      [ Copy the received T1 to C ]
       $K = MOD(K + Q, LCM)$ 
    END DO
  END DO
END PARDO

```

Figure 14: A modified matrix transpose algorithm from matrix point-of-view.  $GCD$  diagonal blocks are transposed simultaneously.

changed from  $Q$  and  $P$  to  $LCM/P$  and  $LCM/Q$ . The pseudocode of the algorithm is shown in Figure 11.

Figure 12 shows two snapshots of the same example, from the matrix point-of-view, to transpose the zeroth and the first diagonal blocks of  $\mathbf{A}$  ( $\mathbf{A}(i, j)$ ,  $\text{MOD}(j - i, LCM) = 0$  and  $1$ , respectively.) The processors which have the blocks to send out are shaded at the bottom. In the example, only  $P \cdot Q / GCD$  processors are involved in block communication in each step. Processors are split into  $GCD$  groups of processors, and a processor  $P(p, q)$  belongs to a group  $g$  if it has the same value of  $g = \text{MOD}(q - p, GCD)$ . Processors in a group  $g$  send and receive their blocks to other processors in another group  $g' = \text{MOD}(GCD - g, GCD)$ . The operations of each group can be overlapped.

The problem is interpreted from the matrix point-of-view. In general, for transposing the  $k$ -th diagonal block of  $\mathbf{A}$  ( $\mathbf{A}(i, j)$ ,  $\text{MOD}(j - i, LCM) = k$ ), a group of processors  $g_k = \text{MOD}(k, GCD)$  send the blocks to another group  $g'_k = \text{MOD}(GCD - g_k, GCD)$ . Since the operations are overlapped over different groups of processors, processors transpose  $GCD$  diagonal blocks simultaneously. So, the matrix can be transposed with  $LCM/GCD$  steps. For the extreme case of  $P = Q = GCD = 3$  as shown in Figure 13, processors transpose 3 ( $= GCD$ ) diagonal blocks at one step. That is, the transposition is done in one step. A processor  $P(p, q)$  exchanges data with processor  $P(q, p)$ . The pseudocode of the algorithm from the matrix point-of-view is shown in Figure 14. The code includes the case of  $GCD = 1$ .

96 processors		64 processors		48 processors	
$P \times Q$	Time (second)	$P \times Q$	Time (second)	$P \times Q$	Time (second)
$6 \times 16$	0.404	$4 \times 16$	0.596	$4 \times 12$	0.652
$8 \times 12$	0.330	$8 \times 8$	0.572	$6 \times 8$	0.546
$12 \times 8$	0.307	$16 \times 4$	0.475	$8 \times 6$	0.527
$16 \times 6$	0.381			$12 \times 4$	0.547

Table 1: Dependence of performance on template configuration for fixed number of processors ( $M = N = 2400$ ).

## 4. Results

In this section we present performance results of the parallel matrix transpose algorithms on the Intel Touchstone Delta computer. The performance of the transpose algorithms cannot be represented in floating point operations per second (flops), since there is no multiplications or additions in the transpose algorithms. The algorithms are combined with a matrix multiplication routine in the PUMMA to compute  $\mathbf{C} = \alpha \mathbf{A}^T \cdot \mathbf{B}^T + \beta \mathbf{C}$  in two steps ( $\mathbf{T} \leftarrow \alpha \mathbf{B} \cdot \mathbf{A}$ ;  $\mathbf{C} \leftarrow \mathbf{T}^T + \beta \mathbf{C}$ ). We assume that  $\alpha = 1$  and  $\beta = 0$  in our test. The performance of  $\mathbf{A}^T \cdot \mathbf{B}^T$  is compared with that of  $\mathbf{A} \cdot \mathbf{B}$ .






Figure 15: Performance comparison of  $A \cdot B$  and  $A^T \cdot B^T$  on  $15 \times 16$  template. ( $P = 15, Q = 16$ , and  $GCD = 1$ ).  $C \Leftarrow A^T \cdot B^T$  is implemented in two steps,  $T \Leftarrow B \cdot A$ , and then  $C \Leftarrow T^T$ .

Matrix elements are generated uniformly on the interval  $[-1, 1]$  in double precision. Conversions between measured runtimes and performance in gigaflops (Gflops) are made assuming an operation count of  $2MNL$  for the multiplication of a  $M \times L$  by a  $L \times N$  matrix. In our test examples, all processors have the same number of blocks so there is no load imbalance. The algorithms were implemented with force type communication [9].

First, we considered how, for a fixed number of processors  $N_p = P \times Q$ , performance depends on the configuration of the processor template. Some typical results are presented in Table 1 for a fixed number of processors. In the test, the block size is fixed at  $5 \times 5$  elements. It may be seen that the template configuration does have some effect on performance. The performance difference is between 19 and 24 %. For rectangular templates with different aspect ratios, the algorithm prefers those with small  $Q$  to those with small  $P$ . On the Delta, communication speed along vertical links seems faster than along horizontal links.

Figures 15 ~ 19 show the performance of the routines on  $15 \times 16$  ( $GCD = 1$ , i.e.,  $P$  and  $Q$  are relatively prime),  $14 \times 16$  ( $GCD = 2$ ),  $12 \times 16$  ( $GCD = 4$ ),  $8 \times 16$  ( $GCD = 8$ ), and  $16 \times 16$  ( $P = Q = GCD = 16$ ) templates, respectively. In all cases the block size is fixed at  $5 \times 5$  elements. The solid and the dashed lines show the performance of  $A^T \cdot B^T$  and  $A \cdot B$ , respectively. The difference of the two lines shows the loss of performance due to matrix transposition.

The transposed multiplication routine shows good performance relative to matrix multiplication. The loss of performance due to the matrix transpose routine is about 2 or 3 %. The




Figure 16: Performance comparison of  $A \cdot B$  and  $A^T \cdot B^T$  on  $14 \times 16$  template. ( $P = 14$ ,  $Q = 16$ , and  $GCD = 2$ )




Figure 17: Performance comparison of  $A \cdot B$  and  $A^T \cdot B^T$  on  $12 \times 16$  template. ( $P = 12$ ,  $Q = 16$ , and  $GCD = 4$ )




Figure 18: Performance comparison of  $A \cdot B$  and  $A^T \cdot B^T$  on  $8 \times 16$  template. ( $P = 8, Q = 16$ , and  $GCD = 8$ )




Figure 19: Performance comparison of  $A \cdot B$  and  $A^T \cdot B^T$  on  $16 \times 16$  template. ( $P = Q = GCD = 16$ ).

$P \times Q$	Matrix Size	Block Size	Time (second)
$8 \times 16$	$4800 \times 4800$	$1 \times 1$	1.557
		$5 \times 5$	1.612
		$300 \times 300$	1.564
$12 \times 16$	$4800 \times 4800$	$1 \times 1$	1.280
		$5 \times 5$	0.893
		$100 \times 100$	0.882
$14 \times 16$	$5600 \times 5600$	$1 \times 1$	1.484
		$5 \times 5$	1.193
		$50 \times 50$	1.161
$15 \times 16$	$6000 \times 6000$	$1 \times 1$	1.740
		$5 \times 5$	1.437
		$25 \times 25$	1.426
$16 \times 16$	$6400 \times 6400$	$1 \times 1$	1.967
		$5 \times 5$	1.967
		$400 \times 400$	1.967

Table 2: Dependence of performance on block size.

transpose routine has excellent performance if  $P$  and  $Q$  are relatively prime. In other cases ( $GCD \geq 2$ ), network congestion may degrade the performance of the routine.

$P \times Q$	Matrix Size ( $\mathbf{A}, \mathbf{B}$ )	$\mathbf{A} \cdot \mathbf{B}$ (%)	$\mathbf{A}^T \cdot \mathbf{B}^T$ (%)
$1 \times 1$	$500 \times 500$	36.70 (100.0)	35.04 (100.0)
$8 \times 16$	$5600 \times 5600$	32.05 (87.3)	30.57 (87.3)
$12 \times 16$	$6720 \times 6720$	32.09 (87.4)	31.64 (90.3)
$14 \times 16$	$6720 \times 6720$	32.52 (88.6)	32.11 (91.6)
$15 \times 16$	$7200 \times 7200$	32.78 (89.3)	32.43 (92.6)
$16 \times 16$	$8000 \times 8000$	31.22 (85.1)	30.38 (86.7)

Table 3: Performance per node in Mflops. Block size is fixed to  $5 \times 5$  elements.  $1 \times 1$  template gives performance of assembly-coded matrix multiplication. Numbers in parentheses represent efficiency compared with the performance on 1 processor.

Table 2 shows how the block size affects the performance of the algorithms. It includes three cases of the block size, two extreme cases - the smallest and largest possible block sizes - and  $5 \times 5$  block of elements. If  $P = Q$ , processors directly copy all blocks at once, so block size does not affect the performance. For the case of the smallest block size ( $1 \times 1$  element) when  $P \neq Q$ , processors make a copy element by element, so it takes a little more time to make a copy. The routines with the smallest block sizes are slower than those with the largest possible block sizes by between 15% and 31%. This difference is negligible, compared with the total elapsed time of the matrix multiplication.

Performance per node is shown in Table 3. The  $1 \times 1$  template gives the performance of the assembly-coded level 3 BLAS matrix multiplication routine for the two cases  $\mathbf{A} \cdot \mathbf{B}$  and

$\mathbf{A}^T \cdot \mathbf{B}^T$ . Processors have about 85% efficiency for  $\mathbf{A} \cdot \mathbf{B}$ , and 87% for  $\mathbf{A}^T \cdot \mathbf{B}^T$  if  $P = Q = 16$ . The routines perform better on templates for which  $P \neq Q$ . Processors achieve about 89%, and 93% of efficiency for each case if  $P$  and  $Q$  are relatively prime.

## 5. Conclusions and Remarks

We have presented parallel matrix transpose algorithms based on the block scattered decomposition. The algorithms have good performance for arbitrary processor configurations on the Intel Delta computer.

If  $P$  and  $Q$  are relatively prime, the transpose routine involves complete exchange communication on a two-dimensional template. We have approached this complicated problem with a direct point-to-point communication scheme (see Section 2). When  $P$  and  $Q$  are not relatively prime ( $GCD > 1$ ), the processors' operations are overlapped over different groups, so that only  $LCM/GCD$  communications are required.

In our Fortran implementation, we assume that the first dimension of the matrix may be different from the number of rows of the matrix in a processor. Even when  $P = Q$ , the processor needs to copy blocks of  $\mathbf{A}$  to a communication buffer before sending, and copy the received buffer to blocks of  $\mathbf{C}$  after receiving.

The parallel matrix transpose algorithms have been combined with matrix multiplication routines. The integrated routines comprise a general-purpose matrix multiplication package, called PUMMA [5], for MIMD message-passing computers. The package has good performance for a wide range of decomposition parameters, that is, its performance depends weakly on processor configuration and block size.

The PUMMA package is currently available only for double precision real data, but will be implemented in the near future for other data types, i.e., single precision real and complex, and double precision complex. To obtain a copy of the software and a description of how to use it, send the message "send pumma from misc" to [netlib@ornl.gov](mailto:netlib@ornl.gov).

## Acknowledgments

The authors would like to thank Eduardo D'Azevedo at ORNL for his helpful suggestions to improve the quality of the paper. This research was performed in part using the Intel Touchstone Delta System operated by the California Institute of Technology on behalf of the Concurrent Supercomputing Consortium. Access to this facility was provided through the Center for Research on Parallel Computing.

## 6. References

- [1] N. G. Azari, A. W. Bojanczyk, and S.-Y. Lee. Synchronous and asynchronous algorithms

- for matrix transposition on MCAP. In *SPIE Vol. 975, Advanced Algorithms and Architecture for Signal Processing III*, pages 277-288, 1988.
- [2] S. H. Bokhari and H. Berryman. Complete exchange on a circuit switched mesh. In *Proceedings of the 1992 Scalable High Performance Computing Conference*, pages 300-306. IEEE Press, 1992.
- [3] J. Choi, J. J. Dongarra, R. Pozo, and D. W. Walker. ScaLAPACK: A scalable linear algebra library for distributed memory concurrent computers. In *Proceedings of Fourth Symposium on the Frontiers of Massively Parallel Computation (McLean, Virginia)*. IEEE Computer Society Press, Los Alamitos, California, October 19-21 1992.
- [4] J. Choi, J. J. Dongarra, and D. W. Walker. The design of scalable software libraries for distributed memory concurrent computers. In *Proceedings of Environment and Tools for Parallel Scientific Computing Workshop, (Saint Hilaire du Touvet, France)*. Elsevier Science Publishers, September 7-8 1992.
- [5] J. Choi, J. J. Dongarra, and D. W. Walker. PUMMA : Parallel universal matrix multiplication algorithms on distributed memory concurrent computers. Technical Report TM-12252, Oak Ridge National Laboratory, Mathematical Sciences Section, April 1993.
- [6] J. J. Dongarra, R. van de Geijn, and D. Walker. A look at scalable linear algebra libraries. In *Proceedings of the 1992 Scalable High Performance Computing Conference*, pages 372-379. IEEE Press, 1992.
- [7] J. O. Eklundh. A fast computer method for matrix transposing. *IEEE Transactions on Computers*, 21:801-803, 1972.
- [8] G. H. Golub and C. V. Van Loan. *Matrix Computations*. The Johns Hopkins University Press, Baltimore, MD, 1989. Second Edition.
- [9] Intel Corporation. *Touchstone Delta Fortran Calls Reference Manual*, April 1991.
- [10] S. L. Johnsson and C.-T. Ho. Algorithms for matrix transposition on boolean n-cube configured ensemble architecture. *SIAM J. Matrix Anal. Appl.* 9:419-454, July 1988.
- [11] R. Littlefield. Characterizing and tuning communications performance for real applications. In *Proceedings, First Intel Delta Application Workshop, CCSF-14-92, Pasadena, California*, pages 179-190, February 1992. presentation overheads.
- [12] D. P. O'Leary. Systolic arrays for matrix transpose and other reorderings. *IEEE Transactions on Computers*, 36:117-122, January 1987.

- [13] S. R. Seidel. Broadcasting on linear arrays and meshes. Technical Report TM-12356, Oak Ridge National Laboratory, Mathematical Sciences Section, March 1993.
- [14] G. Strang. *Linear Algebra and Its Applications*. Harcourt Brace Jovanovich, Inc., San Diego, CA, 1988. Third Edition.

ORNL/TM-12309

### INTERNAL DISTRIBUTION

- |                       |  |
|-----------------------|--|
| 1. B. R. Appleton     | 23-27. S. A. Raby                      |
| 2-6. J. Choi          | 28. C. H. Romine                       |
| 7-8. T. S. Darland    | 29. T. H. Rowan                        |
| 9. E. F. D'Azevedo    | 30-34. R. F. Sincovec                  |
| 10-14. J. J. Dongarra | 35-39. D. W. Walker                    |
| 15. J. B. Drake       | 40-44. R. C. Ward                      |
| 16. T. H. Dunigan     | 45. P. H. Worley                       |
| 17. G. A. Geist       | 46. Central Research Library           |
| 18. L. J. Gray        | 47. ORNL Patent Office                 |
| 19. M. R. Leuze       | 48. K-25 Applied Technology Library    |
| 20. E. G. Ng          | 49. Y-12 Technical Library             |
| 21. C. E. Oliver      | 50. Laboratory Records Department - RC |
| 22. B. W. Peyton      | 51-52. Laboratory Records Department   |

### EXTERNAL DISTRIBUTION

- 53. Christopher R. Anderson, Department of Mathematics, University of California, Los Angeles, CA 90024
- 54. David C. Bader, Atmospheric and Climate Research Division, Office of Health and Environmental Research, Office of Energy Research, ER-76, U.S. Department of Energy, Washington, DC 20585
- 55. David H. Bailey, NASA Ames, Mail Stop 258-5, NASA Ames Research Center, Moffet Field, CA 94035
- 56. Edward H. Barsis, Computer Science and Mathematics, P. O. Box 5800, Sandia National Laboratory, Albuquerque, NM 87185
- 57. Colin Bennett, Department of Mathematics, University of South Carolina, Columbia, SC 29208
- 58. Dominique Bennett, CERFACS, 42 Avenue Gustave Coriolis, 31057 Toulouse Cedex, FRANCE
- 59. Marsha J. Berger, Courant Institute of Mathematical Sciences, 251 Mercer Street, New York, NY 10012



60. Mike Berry, Department of Computer Science, University of Tennessee, 107 Ayres Hall, Knoxville, TN 37996-1301
61. Ake Bjorck, Department of Mathematics, Linkoping University, S-581 83 Linkoping, Sweden
62. A. W. Bojanczyk, School of Electrical Engineering, Cornell University, ETC Building, Rm 337, Ithaca, NY 14853-6367
63. John H. Bolstad, Lawrence Livermore National Laboratory, L-16, P. O. Box 808, Livermore, CA 94550
64. George Bourianoff, Superconducting Super Collider Laboratory, 2550 Beckleymeade Avenue, Suite 260, Dallas, TX 75237-3946
65. Roger W. Brockett, Pierce Hall, 29 Oxford Street, Harvard University, Cambridge, MA 02138
66. Bill L. Buzbee, National Center for Atmospheric Research, P. O. Box 3000, Boulder, CO 80307
67. Thomas A. Callcott, Director, The Science Alliance Program, 53 Turner House, University of Tennessee, Knoxville, TN 37996
68. Captain Edward A. Carmona, Parallel Computing Research Group, U.S. Air Force Weapons Laboratory, Kirtland AFB, NM 87117
69. John Cavallini, Acting Director, Scientific Computing Staff, Applied Mathematical Sciences, Office of Energy Research, U.S. Department of Energy, Washington, DC 20585
70. I-liang Chern, Mathematics and Computer Science Division, Argonne National Laboratory, 9700 South Cass Avenue, Argonne, IL 60439
71. Alexandre Chorin, Mathematics Department, Lawrence Berkeley Laboratory, Berkeley, CA 94720
72. Ray Cline, Sandia National Laboratories, Livermore, CA 94550
73. James Corones, Ames Laboratory, Iowa State University, Ames, IA 50011
74. Jean Côté, RPN, 2121 Transcanada Highway, Suite 508, Dorval, Quebec H9P 1J3, CANADA
75. John J. Dorning, Department of Nuclear Engineering Physics, Thornton Hall, McCormick Road, University of Virginia, Charlottesville, VA 22901

76. Larry Dowdy, Computer Science Department, Vanderbilt University, Nashville, TN 37235
77. Donald J. Dudziak, Department of Nuclear Engineering, 110B Burlington Engineering Labs, North Carolina State University, Raleigh, NC 27695-7909
78. Iain S. Duff, Atlas Centre, Rutherford Appleton Laboratory, Didcot, Oxon OX11 0QX, England
79. John Dukowicz, Los Alamos National Laboratory, Group T-3, Los Alamos, NM 87545
80. Richard E. Ewing, Department of Mathematics, University of Wyoming, Laramie, WY 82071
81. Ian Foster, Mathematics and Computer Science Division, Argonne National Laboratory, 9700 South Cass Avenue, Argonne, IL 60439
82. Geoffrey C. Fox, Northeast Parallel Architectures Center, Syracuse University, Syracuse, NY 13244-4100
83. Chris Fraley, Statistical Sciences, Inc., 1700 Westlake Ave. N, Suite 500, Seattle, WA 98119
84. Paul O. Frederickson, RIACS, MS 230-5, NASA Ames Research Center, Moffet Field, CA 94035
85. Dennis B. Gannon, Computer Science Department, Indiana University, Bloomington, IN 47401
86. J. Alan George, Vice President, Academic and Provost, Needles Hall, University of Waterloo, Waterloo, Ontario, CANADA N2L 3G1
87. James Glimm, Department of Mathematics, State University of New York, Stony Brook, NY 11794
88. Gene Golub, Computer Science Department, Stanford University, Stanford, CA 94305
89. Phil Gresho, Lawrence Livermore National Laboratory, L-262, P. O. Box 808, Livermore, CA 94550
90. William D. Gropp, Mathematics and Computer Science Division, Argonne National Laboratory, 9700 South Cass Avenue, Argonne, IL 60439
91. Eric Grosse, AT&T Bell Labs 2T 504, Murray Hill, NJ 07974
92. John Gustafson, 236 Wilhelm, Ames Laboratory, Iowa State University, Ames, IA 50011

93. James J. Hack, National Center for Atmospheric Research, P. O. Box 3000, Boulder, CO 80307
94. Robert M. Haralick, Department of Electrical Engineering, Director, Intelligent Systems Lab, University of Washington, 402 Electrical Engineering Building, FT-10, Seattle, WA 98195
95. Michael T. Heath, Center for Supercomputing Research and Development, 305 Talbot Laboratory, University of Illinois, 104 South Wright Street, Urbana, IL 61801-2932
96. Michael Henderson, Los Alamos National Laboratory, Group T-3, Los Alamos, NM 87545
97. Fred Howes, Office of Scientific Computing, ER-7, Applied Mathematical Sciences, Office of Energy Research, U. S. Department of Energy, Washington, DC 20585
98. Gary Johnson, Office of Scientific Computing, ER-7, Applied Mathematical Sciences, Office of Energy Research, U. S. Department of Energy, Washington, DC 20585
99. Lennart Johnsson, Thinking Machines Inc., 245 First Street, Cambridge, MA 02142-1214
100. Malvyn Kalos, Cornell Theory Center, Engineering and Theory Center Bldg., Cornell University, Ithaca, NY 14853-3901
101. Hans Kaper, Mathematics and Computer Science Division, Argonne National Laboratory, 9700 S. Cass Avenue, Bldg. 221 Argonne, IL 60439
102. Alan H. Karp, IBM Scientific Center, 1530 Page Mill Road, Palo Alto, CA 94304
103. Kenneth Kennedy, Department of Computer Science, Rice University, P. O. Box 1892, Houston, Texas 77001
104. Tom Kitchens, ER-7, Applied Mathematical Sciences, Scientific Computing Staff, Office of Energy Research, Office G-437 Germantown, Washington, DC 20585
105. Peter D. Lax, Courant Institute of Mathematical Sciences, New York University, 251 Mercer Street, New York, NY 10012
106. James E. Leiss, Rt. 2, Box 142C, Broadway, VA 22815
107. Rich Loft, National Center for Atmospheric Research, P. O. Box 3000, Boulder, CO 80307
108. Michael C. MacCracken, Lawrence Livermore National Laboratory, L-262, P. O. Box 808, Livermore, CA 94550
109. Robert Malone, Los Alamos National Laboratory, C-3, Mail Stop B265, Los Alamos, NM 87545

110. Len Margolin, Los Alamos National Laboratory, Los Alamos, NM 87545
111. Frank McCabe, Department of Computing, Imperial College of Science and Technology, 180 Queens Gate, London SW7 2BZ, ENGLAND
112. James McGraw, Lawrence Livermore National Laboratory, L-306, P. O. Box 808, Livermore, CA 94550
113. Paul C. Messina, Mail Code 158-79, California Institute of Technology, 1201 E. California Blvd. Pasadena, CA 91125
114. Neville Moray, Department of Mechanical and Industrial Engineering, University of Illinois, 1206 West Green Street, Urbana, IL 61801
115. David Nelson, Director of Scientific Computing, ER-7, Applied Mathematical Sciences, Office of Energy Research, U. S. Department of Energy, Washington, DC 20585
116. V. E. Oberacker, Department of Physics, Vanderbilt University, Box 1807, Station B, Nashville, TN 37235
117. Joseph Oliger, Computer Science Department, Stanford University, Stanford, CA 94305
118. Robert O'Malley, Department of Mathematical Sciences, Rensselaer Polytechnic Institute, Troy, NY 12180-3590
119. James M. Ortega, Department of Applied Mathematics, Thornton Hall, University of Virginia, Charlottesville, VA 22901
120. Ron Peierls, Applied Mathematical Department, Brookhaven National Laboratory, Upton, NY 11973
121. Richard Pelz, Dept. of Mechanical and Aerospace Engineering, Rutgers University, Piscataway, NJ 08855-0909
122. Paul Pierce, Intel Scientific Computers, 15201 N.W. Greenbrier Parkway, Beaverton, OR 97006
123. Robert J. Plemmons, Departments of Mathematics and Computer Science, North Carolina State University, Raleigh, NC 27650
124. Jesse Poore, Computer Science Department, University of Tennessee, Knoxville, TN 37996-1300
125. Andrew Priestley, Institute for Computational Fluid Dynamics, Reading University, Reading RG6 2AX, ENGLAND

126. Daniel A. Reed, Computer Science Department, University of Illinois, Urbana, IL 61801
127. Lee Riedinger, Director, The Science Alliance Program, University of Tennessee, Knoxville, TN 37996
128. Garry Rodrigue, Numerical Mathematics Group, Lawrence Livermore National Laboratory, Livermore, CA 94550
129. Ahmed Sameh, University of Illinois at Urbana-Champaign, Center for Supercomputer R&D, 469 CSRL, 1308 West Main St., Urbana, IL 61801
130. Dave Schneider University of Illinois at Urbana-Champaign, Center for Supercomputing Research and Development, 319E Talbot - 104 S. Wright Street Urbana, IL 61801
131. David S. Scott, Intel Scientific Computers, 15201 N.W. Greenbrier Parkway, Beaverton, OR 97006
132. Robert Schreiber, RIACS, MS 230-5, NASA Ames Research Center, Moffet Field, CA 94035
133. William C. Skamarock, 3973 Escuela Court, Boulder, CO 80301
134. Richard Smith, Los Alamos National Laboratory, Group T-3, Mail Stop B2316, Los Alamos, NM 87545
135. Peter Smolarkiewicz, National Center for Atmospheric Research, MMM Group, P. O. Box 3000, Boulder, CO 80307
136. Jurgen Steppeler, DWD, Frankfurterstr 135, 6050 Offenbach, WEST GERMANY
137. Rick Stevens, Mathematics and Computer Science Division, Argonne National Laboratory, 9700 South Cass Avenue, Argonne, IL 60439
138. Paul N. Swarztrauber, National Center for Atmospheric Research, P. O. Box 3000, Boulder, CO 80307
139. Wei Pai Tang, Department of Computer Science, University of Waterloo, Waterloo, Ontario, Canada N2L 3G1
140. Harold Trease, Los Alamos National Laboratory, Mail Stop B257, Los Alamos, NM 87545
141. Robert G. Voigt, ICASE, MS 132-C, NASA Langley Research Center, Hampton, VA 23665
142. Mary F. Wheeler, Rice University, Department of Mathematical Sciences, P. O. Box 1892, Houston, TX 77251

- 143. Andrew B. White, Los Alamos National Laboratory, P. O. Box 1663, MS-265, Los Alamos, NM 87545
- 144. David L. Williamson, National Center for Atmospheric Research, P. O. Box 3000, Boulder, CO 80307
- 145. Samuel Yee, Air Force Geophysics Lab, Department LYP, Hancom AFB, Bedford, MA 01731
- 146. Office of Assistant Manager for Energy Research and Development, U.S. Department of Energy, Oak Ridge Operations Office, P. O. Box 2001, Oak Ridge, TN 37831-8600
- 147-148. Office of Scientific & Technical Information, P. O. Box 62, Oak Ridge, TN 37830

**END**

**DATE  
FILMED**

**12/28/93**

