Engineering Physics and Mathematics Division

Mathematical Sciences Section

**PUMMA :
PARALLEL UNIVERSAL MATRIX MULTIPLICATION 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: August 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

872

Contents

1	Introduction	1
2	Design Issues	2
3	Algorithms	5
3.1	The Basic Matrix Multiplication Algorithm	5
3.2	Matrix Multiplication Algorithm with Block Scattered Decomposition	7
3.3	Transposed Matrix Multiplication Algorithm, $\mathbf{C} = \mathbf{A}^T \cdot \mathbf{B}$	11
3.4	Multiplication of Transposed Matrices, $\mathbf{C} = \mathbf{A}^T \cdot \mathbf{B}^T$	16
4	Results	19
4.1	Comparison of Three Matrix Multiplication Algorithms	19
4.2	Comparison with Transposed Matrix Multiplication Algorithms	21
4.3	Results with Optimized Communication Routines for the Intel Delta	25
5	Conclusions and Remarks	27
6	References	29

**PUMMA :
PARALLEL UNIVERSAL MATRIX MULTIPLICATION ALGORITHMS
ON DISTRIBUTED MEMORY CONCURRENT COMPUTERS**

Jaeyoung Choi

Jack J. Dongarra

David W. Walker

Abstract

This paper describes the Parallel Universal Matrix Multiplication Algorithms (PUMMA) on distributed memory concurrent computers. The PUMMA package includes not only the non-transposed matrix multiplication routine $\mathbf{C} = \mathbf{A} \cdot \mathbf{B}$, but also transposed multiplication routines $\mathbf{C} = \mathbf{A}^T \cdot \mathbf{B}$, $\mathbf{C} = \mathbf{A} \cdot \mathbf{B}^T$, and $\mathbf{C} = \mathbf{A}^T \cdot \mathbf{B}^T$, for a block scattered data distribution. The routines perform efficiently for a wide range of processor configurations and block sizes. The PUMMA together provide the same functionality as the Level 3 BLAS routine xGEMM. Details of the parallel implementation of the routines are given, and results are presented for runs on the Intel Touchstone Delta computer.

1. Introduction

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). One technique to more effectively exploit the power of such machines is to develop algorithms that maximize reuse of data held in the upper levels of the hierarchy, thereby reducing the need for more expensive accesses to lower levels. For dense linear algebra computations this can be done by using block-partitioned algorithms, that is by recasting algorithms in forms that involve operations on submatrices, rather than individual matrix elements. An example of a block-partitioned algorithm for LU factorization is given in [7,16]. The Level 3 Basic Linear Algebra Subprograms (BLAS) perform a number of commonly-used matrix-matrix operations, and are available in optimized form on most computing platforms ranging from workstations up to supercomputers [13].

The Level 3 BLAS have been successfully used as the building blocks of a number of applications, including LAPACK, a software library that uses block-partitioned algorithms for performing dense and banded linear algebra computations on vector and shared memory computers [2,3,9,11,14]. On shared memory machines block-partitioned algorithms reduce the number of times that data must be fetched from shared memory, while on distributed memory machines they reduce the number of messages required to get data from other processors. Thus, there has been much interest recently in developing versions of the Level 3 BLAS for distributed memory concurrent computers [1,6,18,19].

An important routine in the Level 3 BLAS is xGEMM for performing matrix-matrix multiplication. The general purpose routine performs the following operations:

$$C \leftarrow \alpha A \cdot B + \beta C$$

$$C \leftarrow \alpha A^T \cdot B + \beta C$$

$$C \leftarrow \alpha A \cdot B^T + \beta C$$

$$C \leftarrow \alpha A^T \cdot B^T + \beta C$$

where “ $\cdot$ ” denotes matrix multiplication, A , B and C are matrices, and α and β are scalars.

In this paper, we present the Parallel Universal Matrix Multiplication Algorithms (PUMMA) for performing the above operations on distributed memory concurrent computers. *Universal* means that the PUMMA include all the above multiplication routines and that their performance depends weakly on processor configuration and block size. A block scattered data distribution is used, which can reproduce many of the common data distributions used in dense linear algebra computations [8,16], as discussed in the next section. There have been many im-

plementations of matrix multiplication algorithms on distributed memory machines [20,21,24]. Many of them are limited in their use since they are implemented with a pure block (non-scattered) distribution, or specific (not general-purpose) data distribution, and/or on square processor configurations with a specific number of processors (column and/or row numbers of processors are powers of 2). The PUMMA package eliminates all of these constraints.

The first part of this paper focuses on the design and implementation of the non-transposed matrix multiplication routine on distributed memory concurrent computers. We then deal with the other cases. A parallel matrix transpose algorithm, in which a matrix with a block scattered decomposition is transposed over a two-dimensional processor mesh, is presented in a separate paper [10]. All routines are implemented in Fortran 77 plus message passing and compared on the Intel Touchstone Delta computer.

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 (or block cyclic) 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 [8], a 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 PQ equals the number of processors, N_p . Thus, it is useful to imagine the processors arranged as a $P \times Q$ mesh, or template. Then the processor at position (p, q) ($0 \leq p < P$, $0 \leq q < Q$) in the template is assigned the blocks indexed by,

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

where $i = 0, \dots, \lfloor (M_b - p - 1)/P \rfloor$, $j = 0, \dots, \lfloor (N_b - q - 1)/Q \rfloor$, and $M_b \times N_b$ is the size of the matrix in blocks.

Blocks are scattered in this way so that good load balance can be maintained in algorithms, such as LU factorization [7,16], in which rows and/or columns of blocks of a matrix become eliminated as the algorithm progresses. However, for some of the distributed Level 3 BLAS routines a scattered decomposition does not improve load balance, and may result in higher concurrent overhead. The general matrix-matrix multiplication routine xGEMM is an example of such a routine for which a pure block (i.e., nonscattered) decomposition is optimal when considering the routine in isolation. However, xGEMM may be used in an application for which, overall, a scattered decomposition is best. We are faced with the choice of implementing

a nonscattered distributed version of xGEMM, and transforming the data decomposition to this form if necessary each time xGEMM is called, or of providing a scattered version and thereby avoiding having to transform the data decomposition. We opt for the latter solution because it is more general, and does not *impose* on the user the necessity of potentially costly decomposition transformations. Since the nonscattered decomposition is just a special case of the scattered decomposition in which the block size is given by $r = \lceil M/P \rceil$ and $s = \lceil N/Q \rceil$, where the matrix size is $M \times N$, the user still has the option of using a nonscattered decomposition for the matrix multiplication and transforming between decompositions if necessary. The Basic Linear Algebra Communication Subprograms (BLACS) are intended to perform decomposition transformations of this type [4,12,17].

The decompositions of all matrices involved in a call to a Level 3 BLAS routine must be compatible with respect to the operation performed. To ensure compatibility we impose the condition that all the matrices be decomposed over the same $P \times Q$ processor template. Most distributed Level 3 BLAS routines will also require conditions on the block size to ensure compatibility. For example, in performing the matrix multiplication $C = A \cdot B$, if the block size of **A** is $r \times s$ then that of **B** and **C** must be $s \times t$ and $r \times t$, respectively.

Another advantageous aspect of the distributed Level 3 BLAS is that often a distributed routine will call sequential Level 3 BLAS routines. For example, the distributed version of xGEMM, described in Section 3.2, consists of a series of steps in each of which each processor multiplies two local matrices by a call to the sequential version of xGEMM. Since highly optimized assembly-coded versions of the sequential Level 3 BLAS already exist on most processors we can take advantage of these in the distributed implementation.

Figure 1 (a) shows the performance of the DGEMM routine for square matrices on one i860 processor of the Intel Touchstone Delta. In general, performance improves with increasing matrix size and saturates for matrices of size greater than $M = 150$. Figure 1 (b) shows that in our Fortran implementation, for nonsquare matrices, a multiplication a column shape of **A** by a row shape of **B** is more efficient than its opposite. In both the square and nonsquare cases, the size of the matrices multiplied should be maximized in order to optimize performance of the sequential assembly-coded version of xGEMM routines. Thus, in the PUMMA routines, instead of multiplying individual blocks successively on each processor, blocks are conglomerated to form larger matrices which are then multiplied.

The distributed Level 3 BLAS routines have similar argument lists to the sequential Level 3 BLAS routines. In the distributed xGEMM routine, for example, original matrices **A** and **B**, are preserved as in the sequential routine. Users, who are familiar with the sequential routines, should have no difficulty in using the distributed routines.

(a) multiplication of square matrices

(b) multiplication of non-square matrices

Figure 1: Performance of DGEMM on one i860 processor of the Delta. (a) The routine is tested with $A_{M \times M} \cdot B_{M \times M}$, $A_{M \times M/2} \cdot B_{M/2 \times M}$, and $A_{M/2 \times M} \cdot B_{M \times M/2}$, where " $\cdot$ " denotes matrix multiplication, and (b) tested with $A_{500 \times M} \cdot B_{M \times 500}$ and $A_{M \times 500} \cdot B_{500 \times M}$.

Figure 2: A matrix A with 12×12 blocks is distributed over a 2×3 processor template. (a) From the matrix point-of-view. Each shaded and unshaded area represents a different template. 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 global indices of block row and block column, respectively. (b) From the processor point-of-view, each processor has 6×4 blocks.

3. Algorithms

To illustrate the basic parallel algorithm we consider a matrix A distributed over a 2-dimensional processor template as shown in Figure 2 (a), where A with 12×12 blocks is distributed over a 2×3 template. If the matrix distribution is seen from the processor point-of-view as in Figure 2 (b), each processor has several blocks of the matrix and the scattered blocks, $A(0,0), A(2,0), A(4,0), \dots, A(10,0)$ are vertically adjacent in the 2-dimensional array in the first processor P_0 , and can be accessed as one long block column $A(0 : 11 : 2, 0)$. In the same way, $A(0,0), A(0,3), A(0,6), A(0,9)$ are horizontally adjacent in P_0 , and can be accessed as one long block row $A(0, 0 : 11 : 3)$. We exploit this property in implementing the algorithms to deal with larger matrices instead of several small individual blocks. We assume data are stored by column in both our Fortran 77 and message passing implementation. In general, the algorithms are presented from the matrix point-of-view, which is simpler and easier to understand. In dealing with the implementation details, we explain the algorithms from the processor point-of-view.

3.1. The Basic Matrix Multiplication Algorithm

Our matrix multiplication algorithm is a block scattered variant of that of Fox, Hey, and Otto [20], that deals with arbitrary rectangular processor templates.

Suppose the matrix A has M_b block rows and L_b block columns, and the matrix B has L_b


```

DO  $K = 0, L_b - 1$ 
  [Columncast one block of  $\mathbf{B}$  ( $B(I, \text{MOD}(I + K, N_b))$ ),  $I = 0 : L_b$ )
  along each column across template]
  PARDO  $I = 0, M_b - 1$ 
     $KP = \text{MOD}(K + I, L_b)$ 
    PARDO  $J = 0, N_b - 1$ 
       $C(I, J) = C(I, J) + A(I, KP) \cdot B(KP, J)$ 
    END PARDO
  END PARDO
  [Roll  $\mathbf{A}$  leftwards]
END DO

```

Figure 3: A distributed block scattered matrix multiplication algorithm. The PARDO's indicate over which indices the data are decomposed. All indices refer to blocks of elements. Communication phases are indicated in square brackets.

block rows and N_b block columns. Block (I, J) of $\mathbf{C}$ is then given by

$$C(I, J) = \sum_{K=0}^{L_b-1} A(I, K) \cdot B(K, J) \quad (2)$$

where $I = 0, 1, \dots, M_b - 1$, $J = 0, 1, \dots, N_b - 1$. In Equation 2 the order of summation is arbitrary.

Fox et al. initially considered only the case of square matrices in which each processor contains a single row or a single column of blocks. That is, the blocks that start the summation lie along the diagonal. The summation is started at a different point for each block row of $\mathbf{C}$ so that in the phase of the parallel algorithm corresponding to summation index K , $A(I, K)$ and $B(K, J)$ can be multiplied in the processor to which $C(I, J)$ is assigned.

This requires each processor containing a block of $\mathbf{B}$ to be multiplied in step K to broadcast that block along the column of the processor template at the start of the step. Also $\mathbf{A}$ must be rolled leftwards at the end of the step so that each column is overwritten by the one to the right, with the first column wrapping round to overwrite the last column. The pseudocode for this algorithm is shown in Figure 3. Another variant of this algorithm involves broadcasting blocks of $\mathbf{A}$ over rows, and rolling $\mathbf{B}$ upwards.

In Figure 3 and subsequent figures a “columncast” is a communication phase in which one data item (typically a block, or set of blocks) is taken from each block column of the matrix and is broadcast to all the other processors in the same column of the processor template. A “rowcast” is similar, but broadcasts a data item from each block row of the matrix to all processors in the same row of the template.

Figure 4: Snapshot of SDB algorithm. The blocks of the matrix $\mathbf{B}$ communicated in the first two stages of the matrix multiplication algorithm are shown shaded. In this case $P = 2$ and $Q = 3$. In each stage, only one wrapped diagonal is columncast. The total number of stages is L_b .

3.2. Matrix Multiplication Algorithm with Block Scattered Decomposition

We now consider the multiplication of matrices distributed with a block scattered decomposition. The block sizes for matrices $\mathbf{A}$ and $\mathbf{B}$ are $r \times s$ and $s \times t$, respectively, where r , s , and t are arbitrary. In this case the summation in row I starts at $K = I$, so the blocks of $\mathbf{B}$ broadcast in each stage lie along diagonal stripes. The parallel algorithm proceeds in L_b stages, in each of which one block of $\mathbf{B}$ is broadcast along each column of the template, and $\mathbf{A}$ is rolled leftwards. We call this the SDB (Single Diagonal Broadcast) algorithm.

Figure 4 shows, from the matrix point-of-view, the wrapped diagonal blocks of $\mathbf{B}$ broadcast in the first two stages of the SDB algorithm, where $\mathbf{B}$ with 12×12 blocks is distributed over a 2×3 template. Only one wrapped diagonal is columncast in each stage. In implementing the algorithm, the size of the submatrices multiplied in each processor should be maximized to optimize the performance of the sequential xGEMM routine. From the processor point-of-view, as shown in Figure 2 (b), the first processor P_0 has $A(0 : 11 : 2, 0 : 11 : 3)$ and $B(0 : 11 : 2, 0 : 11 : 3)$, and it will have $C(0 : 11 : 2, 0 : 11 : 3)$ after the computation. In the first stage of Figure 4 ($K = 0$), P_0 multiplies $A(0, 0), A(2, 0), \dots, A(10, 0)$ with $B(0, 0)$. These operations can be combined as one matrix multiplication since blocks of $A(0, 0), A(2, 0), \dots, A(10, 0)$ are vertically adjacent in P_0 . The processor multiplies a long block column of $\mathbf{A}$ ($A(0 : 11 : 2, 0)$) with one block $B(0, 0)$. This is the reason why we prefer a scheme of columnwise broadcasting $\mathbf{B}$ to a scheme of rowwise broadcasting $\mathbf{A}$ in our Fortran implementation, where 2-dimensional arrays are stored by columns.

Denoting the least common multiple of P and Q by LCM , we refer to a square of $LCM \times$

```

DO K1 = 0, LCM - 1
  [Columncast  $L_b/LCM$  blocks of B ( $B(I, J : N_b : LCM)$ ,  $I = 0 : L_b$ ,
   $J = \text{MOD}(I + K1, LCM)$ ) along each column of template]
DO K2 = 0,  $L_b/LCM - 1$ 
   $K = K1 + K2 \times LCM$ 
  PARDO I = 0,  $M_b - 1$ 
     $KP = \text{MOD}(K + \text{MOD}(I, LCM), L_b)$ 
    PARDO J = 0,  $N_b - 1$ 
       $C(I, J) = C(I, J) + A(I, KP) \cdot B(KP, J)$ 
    END PARDO
  END PARDO
END DO
[Roll A leftwards]
END DO

```

Figure 5: MDB1 algorithm, which is a distributed matrix multiplication algorithm suitable for a block scattered decomposition. The outer K loop has been split into loops over $K1$ and $K2$ so that the communication for several steps can be sent in a single message.

LCM blocks as an LCM block. Blocks belong to the same processor if their relative locations are the same in each square LCM block. The concept of the LCM block is very useful, since an algorithm may be developed for the first LCM block, and then be applied to the other LCM blocks, which all have the same structure and 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.

For a block scattered decomposition the communication latency can be reduced by performing multiple instances of the outer K loop (see Figure 3) together. The communication latency is reduced when instances of the outer K loop separated by LCM are grouped together, as shown in Figure 5. We call this the MDB1 (Multiple Diagonal Broadcast 1) algorithm. In this case the parallel algorithm proceeds in LCM stages, in each of which $\lceil L_b/LCM \rceil$ blocks of the **B** matrix are broadcast down each column of the template by a single communication phase in the outer loop. In Figure 6 we show the two ($\lceil L_b/LCM \rceil = 12/6$) wrapped diagonal blocks of **B** broadcast in the first two stages of the algorithm. The size of the submatrices multiplied in each processor cannot be increased and it is the same as in the SDB algorithm.

The communication latency can be reduced even further by noting that the data for matrix **A** returns to the processor in which it started after **A** has been rolled Q times. Thus, we introduce a third variant of the parallel algorithm that proceeds in Q stages, in each of which $\lceil L_b/Q \rceil$ blocks of **B** are broadcast down each template column by a single communication

Figure 6: Snapshot of MDB1 algorithm. In this case $P = 2$, $Q = 3$, and so the LCM of P and Q is 6. In each stage, two ($\lceil L_b/LCM \rceil = 12/6$) wrapped diagonals are columncast. The total number of stages is LCM .

phase in the outer loop. Figure 7 shows the four ($\lceil L_b/Q \rceil = 12/3$) wrapped diagonal blocks of $\mathbf{B}$ broadcast in each stage. The pseudocode for this version of the algorithm is the same as that shown in Figure 5, except that “ LCM ” is replaced by “ Q .” This is called the “MDB2 (Multiple Diagonal Broadcast 2)” algorithm.

In implementing the MDB2 algorithm, the granularity of the algorithm is increased. In the first stage shown in Figure 7 ($K1 = 0$), the first processor P_0 multiplies a column block $\mathbf{A}$ ($A(0 : 11 : 2, 0)$) with $B(0, 0)$, $B(0, 3)$, $B(0, 6)$ and $B(0, 9)$. These blocks of $\mathbf{B}$ are horizontally adjacent in the 2-dimensional submatrix in P_0 , and form a long block row $B(0, 0 : 11 : 3)$. These operations are replaced by one multiplication. P_0 multiplies a long block column of $\mathbf{A}$ ($A(0 : 11 : 2, 0)$) with a long block row of $\mathbf{B}$ ($B(0, 0 : 11 : 3)$). The combined multiplication looks like a *block* version of the *outer product* operation. Since $\lceil L_b/LCM \rceil = 2$, P_0 needs to do another outer product operation at the same step, $A(0 : 11 : 2, 6)$ with $B(6, 0 : 11 : 3)$, as shown in Figure 8 (a).

In MDB2 algorithm, the granularity of the algorithm is maximized. P_0 has two block rows of $\mathbf{B}$ to broadcast ($B(0, 0 : 11 : 3)$ and $B(6, 0 : 11 : 3)$), which are condensed to one large matrix ($B(0 : 11 : 6, 0 : 11 : 3)$) for economical communications. If block columns of $\mathbf{A}$ are presorted with radix LCM in the beginning of the algorithm (or radix LCM/Q in each processor) as shown in Figure 8 (b), two block columns of $\mathbf{A}$ ($A(0 : 11 : 2, 0)$ and $A(0 : 11 : 2, 6)$) are accessed as one large matrix ($A(0 : 11 : 2, 0 : 11 : 6)$). Now, P_0 can complete its operation with *one* large matrix multiplication of $A(0 : 11 : 2, 0 : 11 : 6)$ and $B(0 : 11 : 6, 0 : 11 : 3)$. All processors compute one matrix multiplication in each step instead of $\lceil L_b/LCM \rceil$ multiplications. The

Figure 7: Snapshot of MDB2 algorithm. In each stage, four ($L_b/Q = 12/3$) wrapped diagonals are columncast. The total number of stages is Q .

Figure 8: $C = A \cdot B$ in P_0 from processor point-of-view, where $P = 2$, $Q = 3$ and $M_b = N_b = L_b = 12$. Columns of A are presorted in (b). The shaded area of A and B represents blocks to be multiplied, and that of C represents blocks to be updated by the multiplication.

```

DO I = 0, Mb - 1
  PARDO J = 0, Nb - 1
    IP = MOD(I + J, Mb)
    PARDO K = 0, Lb - 1
      T(K) = [A(K, IP)]T · B(K, J)
    END PARDO
    DO K = 0, Lb - 1
      C(IP, J) = C(IP, J) + T(K)
    END DO
  END PARDO
  [Roll A leftwards]
END DO

```

Figure 9: The basic transposed matrix multiplication algorithm, $\mathbf{C} = \mathbf{A}^T \cdot \mathbf{B}$ for a block scattered decomposition. $[A(K, IP)]^T$ is the transpose of block $A(K, IP)$. This algorithm needs a sequential DO loop to compute $C(IP, J)$ by adding the temporary results $T(K)$ columnwise.

computation is like a *block* version of *matrix-matrix* multiplication.

The communication scheme of the MDB2 algorithm can be changed to rowwise broadcasting of $[L_b/P]$ blocks of $\mathbf{A}$ and columnwise shifting of presorted $\mathbf{B}$ without decreasing its performance. The two schemes have the same number of steps and the same amount of computation per processor in each step, but they have different communication strategies.

3.3. Transposed Matrix Multiplication Algorithm, $\mathbf{C} = \mathbf{A}^T \cdot \mathbf{B}$

We now describe the multiplication of transposed matrices, that is, multiplications of the form, $\mathbf{C} = \mathbf{A}^T \cdot \mathbf{B}$ and $\mathbf{C} = \mathbf{A} \cdot \mathbf{B}^T$. The multiplication algorithm of two transposed matrices, $\mathbf{C} = \mathbf{A}^T \cdot \mathbf{B}^T$, is presented in Section 3.4. Lin and Snyder [24] has given an algorithm computing $\mathbf{C} = \mathbf{A} \cdot \mathbf{B}$ based on a block distribution, that first transposes one of the matrices and then uses a series of block multiplication and reduction steps to evaluate $\mathbf{C}$.

Consider first $\mathbf{C} = \mathbf{A}^T \cdot \mathbf{B}$, where $\mathbf{A}$ and $\mathbf{B}$ are $L_b \times M_b$ and $L_b \times N_b$ blocks, respectively, and they are distributed with a block scattered decomposition. $C(I, J)$ is then computed by

$$C(I, J) = \sum_{K=0}^{L_b-1} [A(K, I)]^T \cdot B(K, J) \quad (3)$$

where $I = 0, 1, \dots, M_b - 1$, $J = 0, 1, \dots, N_b - 1$ and $[A(K, I)]^T$ is the transposed block of $A(K, I)$. As in Equation 2, block indices are used, and the order of summation is arbitrary.

Figure 9 gives the pseudocode of the basic transposed matrix multiplication algorithm. The algorithm proceeds in L_b steps, in each of which blocks of $\mathbf{C}$ lying along a wrapped diagonal

```

DO I1 = 0, Q - 1
  DO I2 = 0, Mb/Q - 1
    I = I1 + I2 × Q
    PARDO J = 0, Nb - 1
      IP = MOD(I + MOD(J, Q), Mb)
      PARDO K1 = 0, P - 1
        T(K1) = 0.0
        DO K2 = 0, Lb/P - 1
          K = K1 + K2 × P
          T(K1) = T(K1) + [A(K, IP)]T · B(K, J)
        END DO
      END PARDO
      DO K1 = 0, P - 1
        C(IP, J) = C(IP, J) + T(K1)
      END DO
    END PARDO
  END DO
  [Roll A leftwards]
END DO

```

Figure 10: The transposed matrix multiplication algorithm, $C = A^T \cdot B$. The outer loop has been split into loops over $I1$ and $I2$ so that the communication for several steps can be sent in a single message.

are evaluated. Each step consists of block matrix multiplication to form contributions to a wrapped diagonal block of C , followed by summation over columns. Finally, a communication phase shifts A to the left by one block.

As in the MDB1 matrix multiplication algorithm of Section 3.2, the communication latency is reduced by simultaneously performing multiple instances of the outer I loop separated by LCM . Again the communication latency is reduced further when instances of the outer loop separated by Q are executed together as in the MDB2 algorithm. The blocks of A return to the same processor from which they started after they have been rolled Q times. So the algorithm proceeds in Q stages, in each of which $\lceil L_b/Q \rceil$ wrapped diagonal blocks of C are computed. The pseudocode of the modified algorithm is shown in Figure 10.

The transposed matrix multiplication algorithm is conceptually simpler than the non-transposed matrix multiplication algorithm. In $C = A^T \cdot B$, processors in the same column of the template compute and add their products, and distribute the summations to the appropriate positions. The most difficult aspect when implementing the algorithm is how to add and distribute the products efficiently.

As an example, consider the matrix multiplication $C = A^T \cdot B$ where matrices A and B ,

Figure 11: Snapshot of $C = A^T \cdot B$ when $P = Q = 3$ and $M_b = N_b = L_b = 6$. (a) From the matrix point-of-view, the computed blocks of the matrix C in the first two stages of the transposed matrix multiplication algorithm are shaded. (b) Snapshot of the first stage from the processor point-of-view. The shaded area of A and B represents blocks to be multiplied, and that of C denotes blocks computed from the multiplication. Only diagonal processors have results in the first stage. After each stage, A is shifted to the left.

each consisting of 6×6 blocks, are distributed over a 3×3 processor template as shown in Figure 11. In each stage, every Q -th wrapped block diagonal of C is computed. In the first stage, as shown in Figure 11 (b), the processors in the first column of the template, P_0 , P_3 , and P_6 , multiply the zeroth and third block columns of A ($A(:, 0 : 5 : 3)$) with the zeroth and third block columns of B ($B(:, 0 : 5 : 3)$). They compute their own portion of multiplications and add them to obtain 2×2 blocks of C ($C(0 : 5 : 2, 0 : 5 : 2)$), which are placed in P_0 . In this example, where the template is square, only the diagonal processors P_0 , P_4 , and P_8 have the computed blocks of C for each column of the template. After the first stage, A shifts to the left. The next wrapped diagonal processors P_2 , P_3 , and P_7 have the computed blocks of C in the second stage.

Figure 12 shows the case of $P = 3$, $Q = 2$, where C is computed in two stages. The first

Figure 12: Snapshot of $C = A^T \cdot B$ when $P = 3$, $Q = 2$, and $M_b = N_b = L_b = 6$. (a) From matrix point-of-view, the computed blocks of the matrix C in the first two stages of the transposed matrix multiplication algorithm are shaded. (b) Snapshot of the first stage from processor point-of-view. If P and Q are relatively prime, the computed blocks of C are scattered over all processors in each stage.

column of processors, P_0 , P_2 , and P_4 , compute 3×3 blocks of C ($C(0 : 5 : 2, 0 : 5 : 2)$), by multiplying the zeroth, second and fourth block columns of A ($A(:, 0 : 5 : 2)$) with the zeroth, second and fourth block columns of B ($B(:, 0 : 5 : 2)$). After summing over columns they have computed their own row blocks of C .

When Q is smaller than P , processors need more memory to store the partial products, if they compute their own products first and then add them together. Imagine the case when $P = 4$, $Q = 1$ and $M_b = N_b = L_b = 4$. Each processor has 1×4 blocks of A and B , and it has 1×4 blocks of C after the computation. But processors need 4×4 blocks to store their own partial products. Thus, memory requirements do not scale well.

Processors can multiply one block column of A with whole blocks of B in each step to avoid nonscalable memory use. In the first step of Figure 12, P_0 , P_2 , and P_4 compute $C(0 : 5 : 2)$ by multiplying $A(:, 0)$ with $B(:, 0 : 5 : 2)$. The computed blocks of C are placed in P_0 . These

(a) $C = A^T \cdot B$ in P_0 from processor point-of-view

(b) computation in P_0 with presorted A^T

Figure 13: $C = A^T \cdot B$ in P_0 from processor point-of-view, where $P = 3$, $Q = 2$ and $M_b = N_b = L_b = 12$. The shaded area of A and B represents blocks to be multiplied. And that of C stands for the result blocks to be placed after multiplication and summation processes over the column of the template.

processors then compute $C(2, 0 : 5 : 2)$, which is placed in P_4 , and finally compute $C(4, 0 : 5 : 2)$, which is placed in P_2 . After this stage A is shifted to the left. With this scheme, the processors require three steps to compute $C(0 : 5 : 2, 0 : 5 : 2)$ for the first stage of the algorithm. This procedure is less efficient, but needs less memory to hold partial products.

The loss of efficiency can be offset by overlapping computation and communication. Consider a modified algorithm in which the blocks of C rotate downwards over the processor template after each stage. Each processor computes its own products and updates the received blocks. The processors receive their own desired blocks of C after $P - 1$ communications. If P and Q are relatively prime as shown in Figure 12, all processors have their own blocks of C in each stage. They receive partial products from the processor above, add their contributions to the partial products, and then send them to the processor below. If processors are waiting to receive the products before multiplying some processors have to wait a long time when $P = Q$ as in Figure 11 (or P and Q are not relatively prime). For these cases, processors compute their own multiplications first, and then add them after they receive the products. This can be implemented effectively with *asynchronous message passing* to minimize processors' waiting time to receive the products.

As an example, consider Figure 13 (a), where 12×12 block matrices are distributed over a

3×2 processor template. P_0 computes two ($\lceil M_b/LCM \rceil$) transposed matrix multiplications of block columns of $\mathbf{A}$ ($A(0:11:3,0)$ and $A(0:11:3,6)$) with its own submatrix $\mathbf{B}$ ($B(0:11:3,0:11:2)$), and generates two block rows of $\mathbf{C}$ ($C(0,0:11:2)$ and $C(6,0:11:2)$). The two rows of $\mathbf{C}$ are condensed for fast communications as in the MDB2 algorithm in Section 3.2. If block columns of $\mathbf{A}$ are presorted with radix LCM (or radix LCM/Q for each processor) at the beginning of the algorithm, processors compute *one* transposed matrix multiplication in each step instead of $\lceil L_b/LCM \rceil$ multiplications as shown in Figure 13 (b). Again, the computation is like a *block* version of (*transposed*) *matrix-matrix* multiplication.

The case $\mathbf{C} = \mathbf{A} \cdot \mathbf{B}^T$ is similar to the $\mathbf{C} = \mathbf{A}^T \cdot \mathbf{B}$ algorithm, but the partial result blocks of $\mathbf{C}$ rotate horizontally in each step, and $\mathbf{B}^T$ shifts upwards after each stage.

3.4. Multiplication of Transposed Matrices, $\mathbf{C} = \mathbf{A}^T \cdot \mathbf{B}^T$

Suppose we need to compute $\mathbf{C} = \mathbf{A}^T \cdot \mathbf{B}^T$, where $\mathbf{A}$ is $L_b \times M_b$ blocks, $\mathbf{B}$ is $N_b \times L_b$ blocks, and $\mathbf{C}$ is $M_b \times N_b$ blocks. One approach is to evaluate the product

$$C(I, J) = \sum_{K=0}^{L_b-1} [A(K, I)]^T \cdot [B(J, K)]^T, \quad (4)$$

directly using a variant of the matrix multiplication routine in Section 3.2, but in which blocks of $\mathbf{A}$ are columncast in each step, and blocks of $\mathbf{B}$ are rotated leftwards. The resultant matrix then has to be blockwise transposed, i.e., block $C(I, J)$ must be swapped with block $C(J, I)$, in order to obtain $\mathbf{C}$. Thus, for this approach the algorithm is as follows,

1. locally transpose each block of $\mathbf{A}$ and $\mathbf{B}$,
2. multiply $\mathbf{A}$ and $\mathbf{B}$ using variant of parallel algorithm,
3. do a blockwise transpose of the result to get $\mathbf{C}$.

In an actual implementation, the local transpose in (1) can be performed within the calls to the assembly-coded sequential xGEMM routine.

Another approach is to evaluate $\mathbf{C}^T = \mathbf{B} \cdot \mathbf{A}$ and then transpose the resulting matrix to obtain $\mathbf{C}$. In this case the algorithm is as follows,

1. multiply $\mathbf{B}$ and $\mathbf{A}$ using the parallel algorithm in Section 3.2,
2. locally transpose each block of result,
3. do a blockwise transpose to get $\mathbf{C}$.

These last two steps together transpose $\mathbf{C}^T$, and may be done in any order. The performance of both approaches is very nearly the same, but the second approach has the advantage of using

(a) block size = 5

(b) block size = 10

Figure 14: Performance comparison of the three matrix multiplication routines on an 8×8 processor template.

(a) block size = 5

(b) block size = 10

Figure 15: Performance comparison of the three matrix multiplication routines on a 9×8 processor template.

96 processors		64 processors		48 processors	
$P \times Q$	Gflops	$P \times Q$	Gflops	$P \times Q$	Gflops
6×16	1.972	4×16	1.373	4×12	1.101
8×12	2.007	8×8	1.447	6×8	1.181
12×8	2.008	16×4	1.444	8×6	1.200
16×6	2.002			12×4	1.130

Table 1: Dependence of performance on template configuration ($M = N = L = 1600$).

the existing algorithm for finding $\mathbf{B} \cdot \mathbf{A}$, as described in Section 3.2, without any modification being necessary. Parallel matrix transpose algorithms are described in [10], and are used to compute $\mathbf{C} = \alpha \mathbf{A}^T \cdot \mathbf{B}^T + \beta \mathbf{C}$ as described above in two steps: $\mathbf{T} = \alpha \mathbf{B} \cdot \mathbf{A}$, then $\mathbf{C} = \mathbf{T}^T + \beta \mathbf{C}$.

4. Results

In this section we present performance results for the PUMMA package on the Intel Touchstone Delta system. 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.

4.1. Comparison of Three Matrix Multiplication Algorithms

We first compared the three matrix multiplication algorithms, SDB, MDB1, and MDB2 on two fixed processor templates. Figures 14 and 15 show the performance of the algorithms on a square processor template (8×8 , $P = Q$) and a nonsquare template (9×8 , P and Q are relatively prime), respectively. Two different block sizes are considered to see how block size affects the performance of the algorithms for a number of different sized matrices.

The performance of the SDB and MDB1 algorithms improves as the block size is increased from 5 to 10, but this change of the block size has almost no effect on the performance of the MDB2 algorithm, since in MDB2 the size of the submatrices multiplied in each processor (using the assembly-coded Level 3 BLAS) is independent of block size. For a square template, the number of communication steps is the same in the MDB1 and MDB2 algorithms since $LCM = Q$, but there is a big difference in their performance. This difference arises because the basic operation of the MDB1 algorithm is a multiplication of a block column of $\mathbf{A}$ with a single block of $\mathbf{B}$, where as, in the MDB2 algorithm, larger matrices are multiplied in each step, as explained in Section 3.2.

The block size is selected by the user. In most cases, the optimal block size is determined by the size and shape of the processor template, floating-point performance of the processor,

(a) Performance of MDB2

(b) Isogranularity Plot

Figure 16: Performance of MDB2 algorithm. (a) Performance in gigaflops as a function of matrix size for different numbers of processors. (b) Isogranularity curves in the (G, N_p) plane. The curves are labeled by the granularity g in units of 10^3 matrix elements per processor.

communication bandwidth between processors, and the size of the matrices. However, for the MDB2 algorithm, the performance is independent of the block size. We adopted a block size of 5×5 in all subsequent runs of the matrix multiplication routines.

We next considered how, for a fixed number of processors $N_p = P \times Q$, performance depended on the configuration of the processor template. Some typical results are presented in Table 1 from which it may be seen that the template configuration does have a small effect on performance, with squarer templates giving better performance than long, thin templates. For a fixed number for processors, a larger value of Q increases the number of outer loops performed, but reduces the time to broadcast blocks of $\mathbf{B}$ across the template. The relative importance of these two factors determines the optimal template configuration. For rectangular templates with different aspect ratios, those with small Q show better performance than those with small P . For a fixed processor template with small P , an MDB2 algorithm, in which $\mathbf{A}$ is broadcast rowwise and $\mathbf{B}$ is shifted columnwise, is preferable to the version described in Section 3.2, in which $\mathbf{B}$ is broadcast columnwise and $\mathbf{A}$ is shifted rowwise.

Figure 16 (a) shows the performance of the MDB2 algorithm on the Intel Touchstone Delta as a function of problem size for different numbers of processors for up to 256 processors. In all cases a square processor template was used, i.e. $P = Q$, the block size was fixed at 5×5 elements, and the test matrices were of size up to 400×400 elements per processor.

In Figure 16 (b) we show how performance depends on the number of processors for a fixed grain size. The fact that these isogranularity plots are almost linear indicates that the distributed matrix multiplication routine scales well on the Delta, even for small granularity.

4.2. Comparison with Transposed Matrix Multiplication Algorithms

We compared the performance of the MDB2 version of the matrix multiplication routine $\mathbf{C} = \mathbf{A} \cdot \mathbf{B}$ with that of the transposed matrix multiplication routines, $\mathbf{C} = \mathbf{A}^T \cdot \mathbf{B}$, and $\mathbf{C} = \mathbf{A}^T \cdot \mathbf{B}^T$. For $\mathbf{C} = \mathbf{A} \cdot \mathbf{B}$, we adopted a routine with rowwise broadcasting of $\mathbf{A}$ and columnwise shifting of $\mathbf{B}$. $\mathbf{C} = \mathbf{A}^T \cdot \mathbf{B}$ is implemented as described in Section 3.3. For $\mathbf{C} = \mathbf{A}^T \cdot \mathbf{B}^T$, $\mathbf{B}$ is directly multiplied with $\mathbf{A}$ to form $\mathbf{B} \cdot \mathbf{A}$, which is then transposed to give $\mathbf{C}$.

Figures 17, 18, 19, and 20 show the performance of the algorithms on 8×8 , 8×9 , 8×10 , and 8×12 templates, respectively. In all cases the block size is fixed at 5×5 elements. The solid and the dashed lines show the performance of $\mathbf{A} \cdot \mathbf{B}$ and $\mathbf{A}^T \cdot \mathbf{B}^T$, respectively. The difference of the two lines is due to the matrix transpose routine used in evaluating $\mathbf{A}^T \cdot \mathbf{B}^T$. In most cases, the performance of the $\mathbf{A}^T \cdot \mathbf{B}$ algorithm, which is drawn with the dot-dashed lines, lies between that of the $\mathbf{A} \cdot \mathbf{B}$ and $\mathbf{A}^T \cdot \mathbf{B}^T$ algorithms, but for the square template in Figure 17, its performance is worse than that of $\mathbf{A}^T \cdot \mathbf{B}^T$. In the $\mathbf{A}^T \cdot \mathbf{B}$ routine, processors in the same column of the template sequentially update their own $\mathbf{C}$. Some of the processors have to wait

Figure 17: Performance comparison of three routines on an 8×8 template. $P = Q = LCM = 8$, and $GCD = 8$

Figure 18: Performance comparison of three routines on an 8×9 template. $P = 8, Q = 9, LCM = 72$, and $GCD = 1$

Figure 19: Performance comparison of three routines on an 8×10 template. $P = 8, Q = 10, LCM = 40$, and $GCD = 2$

Figure 20: Performance comparison of three routines on an 8×12 template. $P = 8, Q = 12, LCM = 24$, and $GCD = 4$

$P \times Q$	Matrix Size	Block Size	$\mathbf{A} \cdot \mathbf{B}$	$\mathbf{A}^T \cdot \mathbf{B}$	$\mathbf{A}^T \cdot \mathbf{B}^T$
8×8	2400×2400	1×1	1.640	1.529	1.607
		5×5	1.641	1.530	1.619
		300×300	1.643	1.531	1.618
8×9	2520×2520	1×1	1.902	1.904	1.732
		5×5	1.924	1.939	1.850
		35×35	1.926	1.946	1.860
8×10	2400×2400	1×1	2.085	2.067	1.961
		5×5	2.107	2.110	2.033
		60×60	2.096	2.123	2.028
8×12	2400×2400	1×1	2.374	2.121	2.265
		5×5	2.389	2.310	2.306
		100×100	2.397	2.338	2.317

Table 2: Dependence of performance on block size (Unit: Gflops)

$P \times Q$	Matrix Size	$\mathbf{A} \cdot \mathbf{B}$		$\mathbf{A}^T \cdot \mathbf{B}$		$\mathbf{A}^T \cdot \mathbf{B}^T$	
1×1	400×400	36.21	(100.0)	35.54	(100.0)	34.58	(100.0)
8×8	3200×3200	27.77	(76.7)	25.86	(72.8)	27.36	(79.1)
8×9	3240×3240	29.00	(80.1)	28.56	(80.4)	28.10	(81.3)
8×10	3200×3200	28.25	(78.0)	27.74	(78.1)	27.47	(79.4)
8×12	3200×3200	28.44	(78.5)	27.55	(77.5)	27.48	(79.5)

Table 3: Performance per node in Mflops. Block size is fixed at 5×5 elements. Entries for the 1×1 template case give the performance of the assembly-coded Level 3 BLAS matrix multiplication routine. Numbers in parenthesis are concurrent efficiency.

a long time to receive the partial products if $P = Q$.

Table 2 shows how the block size has an effect on the performance of the algorithms. It includes three cases of the block size, two extreme cases – the smallest and largest possible block sizes – and 5×5 block of elements. The algorithms depend only weakly on the block size. Even for the case of the smallest block size (1×1 element), the algorithms show good performance.

Performance per node is shown in Table 3. The 1×1 template gives the performance of the assembly-coded Level 3 BLAS matrix multiplication routine. The numbers in parentheses are concurrent efficiency, which is the relative performance of nodes compared with the maximum performance of the assembly-coded Level 3 BLAS routine. Approximately 77% efficiency is achieved for $\mathbf{A} \cdot \mathbf{B}$, 73% for $\mathbf{A}^T \cdot \mathbf{B}$, and 79% for $\mathbf{A}^T \cdot \mathbf{B}^T$ if $P = Q$. The routines perform better on templates for which $P \neq Q$. More than 80% efficiency is achieved for all cases if P and Q are relatively prime.

Figure 21: Two rotating schemes. (a) Nodes first send to the left and then receive from the right. (b) In the first step, odd-numbered processors send data blocks and even-numbered processors receive them. In the next step, even-numbered processors send and odd-numbered processors receive. Odd-even rotating is faster on Delta, but simultaneous rotating is faster on iPSC/860 hypercube.

4.3. Results with Optimized Communication Routines for the Intel Delta

For the implementation of the PUMMA package, blocking and nonblocking communication schemes were used. In this section, we modify the algorithms with optimized communication schemes specifically for the Intel Touchstone Delta.

First, *force type* communications [22] are incorporated for faster communications. A force type message bypasses the normal flow control mechanism, and is not delayed by clogged message buffers on a processor. A force type message is discarded if no receive has been posted on the destination processor prior to its arrival. If force types are not used on the Delta, the routine can accommodate matrices up to 400×400 elements per processor without encountering problems arising from system buffer overflow [23]. With force type communication, the routines can handle larger matrices, up to 500×500 per processor, where the maximum size is determined by the available memory per processor rather than by system buffer constraints.

A block rotating scheme is used to shift $\mathbf{A}$ rowwise in the MDB2 algorithm of Section 3.2 and in the $\mathbf{A}^T \cdot \mathbf{B}$ routine of Section 3.3. A simultaneous rotating scheme, shown in Figure 21 (a), may be used on the Intel iPSC/860 hypercube. However, an odd-even rotating scheme is preferable on the Delta [25]. This scheme performs the communication in two steps as shown in Figure 21 (b). In the first step, odd-numbered processors send their own data blocks and even-numbered processors receive them. In the next step, even-numbered processors send and odd-numbered processors receive.

In the original MDB2 algorithm, blocks of $\mathbf{B}$ are broadcast in each column of the template based on a ring communication scheme. In the Delta-specific MDB2 algorithm, messages are broadcast based on a minimum spanning tree. A special broadcasting routine is desirable for the Delta, which differs from that used on hypercubes [5]. Consider broadcasting a message on a linear array of $p = 7$ processors as shown in Figure 22, where nodes are numbered 0 through 6.

Figure 22: Broadcasting on linear array of $p = 7$, where nodes are numbered 0 through 6. P_2 is a root node.

Figure 23: Collecting on linear array. P_2 is a root node.

In the hypercube scheme, the root node P_2 , which has the message to be broadcast, first sends the message to P_3 , whose least significant bit (LSB) is different from the root node. Then the message is delivered by toggling successive bits from LSB to the most significant bit (MSB). On a mesh topology such as the Delta, the network traffic becomes congested as the broadcast proceeds, as shown in Figure 22 (a).

In order to avoid network contention, the root node sends the message to the first node in the other half of the processors. By recursing for $\lceil \log_2 P \rceil$ similar steps, the message is delivered to all nodes without any contention as shown in Figure 22 (b). In general, each column of the template has P/GCD root nodes in a stage, which broadcast their blocks of $\mathbf{B}$ over GCD processors of the column, where GCD denotes the greatest common divisor of P and Q . These operations are a form of *group communication* [15].

For $\mathbf{A}^T \cdot \mathbf{B}$ in Section 3.3, the partial products in the same column of the processors are combined and the sum is stored in the root (destination) node. A special collecting scheme has been developed for the Delta to avoid network contention. The new collecting scheme on a linear array shown in Figure 23 (b) is based on the broadcasting scheme in Figure 22 (b). The partial products are sent and added in nodes which are nearer to the root node. Generally, in each stage of the algorithm, each column of the template has P/GCD root nodes to collect the partial products. Partial products of a group of GCD processors are added first with

$P \times Q$	Matrix Size	$\mathbf{A} \cdot \mathbf{B}$		$\mathbf{A}^T \cdot \mathbf{B}$	
		RING	TREE	RING	TREE
4×4	2000×2000	0.515	0.530	0.488	0.496
6×6	3000×3000	1.130	1.159	1.081	1.073
8×8	4000×4000	1.985	2.056	1.908	1.901
12×12	6000×6000	4.438	4.507	4.260	4.150
16×16	8000×8000	7.844	8.001	7.542	7.325
8×9	3960×3960	2.326	2.326	2.321	2.321
8×10	4400×4400	2.561	2.641	2.493	2.486
8×12	4800×4800	2.965	3.091	2.956	2.918
8×16	5600×5600	3.858	4.022	3.707	3.622

Table 4: Performance in Gflops with optimized communication routines on two structures, ring and spanning tree. Block size is fixed to 5×5 . The routine for $\mathbf{A} \cdot \mathbf{B}$ is faster for a tree structure, but the routine for $\mathbf{A}^T \cdot \mathbf{B}$ has better performance for a ring structure.

the products of other group(s) in the same column. After $P/GCD - 1$ communications and additions, the partial products in each group of GCD processors are effectively added to the root nodes.

The $\mathbf{A} \cdot \mathbf{B}$ and $\mathbf{A}^T \cdot \mathbf{B}$ routines have been implemented with the optimized communications for the Delta based on both the ring and the minimum spanning tree structure for broadcasts. Performance results are shown in Table 4. The non-transposed matrix multiplication routine for 8000×8000 matrices on 16×16 nodes performs at about 8.00 Gflops for the tree structure, and the transposed multiplication routine executes at about 7.54 Gflops for the ring structure. They obtain about 31.25 Mflops and 29.46 Mflops per processor, respectively, which correspond to concurrent efficiencies of 86% and 83%, respectively.

If P and Q are relatively prime, there is no performance difference between tree and ring versions. The $\mathbf{A} \cdot \mathbf{B}$ algorithm performs well for the tree structure. Though broadcasting a message to the entire column of the processors on the ring is slow, the overall performance is not influenced since the stages of the algorithm are pipelined. That is, processors directly proceed to the next stage as soon as they finish their multiplication at the current stage.

In a single stage of the $\mathbf{A}^T \cdot \mathbf{B}$ routine, collecting the partial products in a column of the processor template is faster for the tree algorithm. However, overall the ring algorithm is preferred for the $\mathbf{A}^T \cdot \mathbf{B}$ routine, since stages of the algorithm can be pipelined.

5. Conclusions and Remarks

We have presented a parallel matrix multiplication routine and its variants for the block scattered decomposition over a two dimensional processor template. We have described how to develop the algorithms for distributed memory concurrent computers from a matrix point-of-

view, and given implementation details from a processor point-of-view. Finally we have shown how to adapt the communications for a specific target machine, the Intel Touchstone Delta computer, by exploiting its communication characteristics. The general purpose matrix multiplication routines developed are universal algorithms that can be used for arbitrary processor configuration and block size.

In general, the first dimension of the data matrix may be different from the number of rows of the matrix in a processor. That means, when shifting **A** in the MDB2 routine, **A** needs to be copied before it is sent out. Instead of a direct copy, the block columns of **A** are presorted so that each processor performs a block version of matrix-matrix multiplication in each step. Without this presorting, processors compute multiplications as a block version of the outer product operation, i.e., a column of blocks is multiplied by a row of blocks. The outer product operation performs well and its performance is almost the same as the routine with presorting for blocks larger than 5×5 elements. But for the case of small block sizes, presorting improves performance. If the first dimension of matrix **A** is the same as the number of rows, the presorting is not necessary, and **A** can be sent out directly, since after Q shifts of **A** processors have their original blocks, and **A** is unchanged. This scheme may also save communication buffer space. For the transposed matrix multiplication routines ($\mathbf{A}^T \cdot \mathbf{B}$ and $\mathbf{A} \cdot \mathbf{B}^T$), the presorting process improves the performance more than 10% for a block size of 5×5 .

In some cases, the transposed matrix multiplication algorithm may be slower than the two combined routines, matrix transposition and matrix multiplication. That is, $\mathbf{C} \leftarrow \mathbf{A}^T \cdot \mathbf{B}$ can be implemented with two steps, ($\mathbf{T} \leftarrow \mathbf{A}^T$, $\mathbf{C} \leftarrow \mathbf{T} \cdot \mathbf{B}$), where extra memory space for **T** is necessary. Users can choose the best routine according to their machine specifications and their application. The performance of the routines not only depends on the machine characteristics, but also the processor configuration and the problem size.

The performance of the PUMMA package can be improved with optimized assembly-coded routines, if available, such as a two dimensional buffer copy routine ($\mathbf{T} \leftarrow \text{op}(\mathbf{A})$, where $\text{op}(\mathbf{A}) = \mathbf{A}$ or $\mathbf{A}^T$), and a two dimensional addition routine ($\mathbf{T} \leftarrow \alpha \mathbf{A} + \beta \mathbf{T}$).

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 following message "send index from pumma" to netlib@ornl.gov.

Acknowledgments

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] P. R. Amestoy, M. J. Dayde, I. S. Duff, and P. Morere. Linear algebra calculations on the BBN TC2000. In G. Goos and J. Hartmanis, editors, *Proceedings of Second Joint International Conference on Vector and Parallel Processing*, pages 319-330. Springer-Verlag, 1992.
- [2] E. Anderson, Z. Bai, C. Bischof, J. Demmel, J. Dongarra, J. DuCroz, A. Greenbaum, S. Hammarling, A. McKenney, and D. Sorensen. LAPACK: A portable linear algebra library for high-performance computers. In *Proceedings of Supercomputing '90*, pages 1-10. IEEE Press, 1990.
- [3] E. Anderson, Z. Bai, J. Demmel, J. Dongarra, J. DuCroz, A. Greenbaum, S. Hammarling, A. McKenney, S. Ostrouchov, and D. Sorensen. *LAPACK Users' Guide*. SIAM Press, Philadelphia, PA, 1992.
- [4] E. Anderson, A. Benzoni, J. Dongarra, S. Moulton, S. Ostrouchov, B. Tourancheau, and R. van de Geijn. Basic Linear Algebra Communication Subprograms. In *Sixth Distributed Memory Computing Conference Proceedings*, pages 287-290. IEEE Computer Society Press, 1991.
- [5] M. Barnett, D. G. Payne, and R. van de Geijn. Optimal minimum spanning tree broadcasting in mesh-connected architecture. Technical Report TM-91-38, The University of Texas at Austin, December 1991.
- [6] P. Berger, M. J. Dayde, and P. Morere. Implementation and use of Level 3 BLAS kernels on a transputer T800 ring network. Technical Report TR/PA/91/54, CERFACS, June 1991.
- [7] 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.
- [8] 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.

- [9] J. Choi, J. J. Dongarra, and D. W. Walker. Level 3 BLAS 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.
- [10] J. Choi, J. J. Dongarra, and D. W. Walker. Parallel matrix transpose algorithms on distributed memory concurrent computers. Technical Report TM-12309, Oak Ridge National Laboratory, Mathematical Sciences Section, August 1993.
- [11] J. Demmel, J. J. Dongarra, J. Du Croz, A. Greenbaum, S. Hammarling, and D. Sorensen. Prospectus for the development of a linear algebra library for high performance computers. Technical Report 97, Argonne National Laboratory, Mathematics and Computer Science Division, September 1987.
- [12] J. J. Dongarra. Workshop on the BLACS. LAPACK Working Note 34, Technical Report CS-91-134, University of Tennessee, 1991.
- [13] J. J. Dongarra, I. Duff, J. Du Croz, and S. Hammarling. A set of level 3 basic linear algebra subprograms. *ACM TOMS*, 16:1-17, March 1990.
- [14] J. J. Dongarra, I. S. Duff, D. C. Sorensen, and H. A. van der Vorst. *Solving Linear Systems on Vector and Shared Memory Computers*. SIAM, Philadelphia, PA, 1990.
- [15] J. J. Dongarra, R. Hempel, A. J. G. Hey, and D. W. Walker. A proposal for a user-level, message passing interface in a distributed memory environment. Technical Report TM-12231, Oak Ridge National Laboratory, March 1993.
- [16] 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.
- [17] J. J. Dongarra and R. A. van de Geijn. Two dimensional basic linear algebra communication subprograms. LAPACK Working Note 37, Technical Report CS-91-138, University of Tennessee, 1991.
- [18] A. C. Elster. Basic matrix subprograms for distributed memory systems. In D. W. Walker and Q. F. Stout, editors, *Proceedings of the Fifth Distributed Memory Computing Conference*, pages 311-316. IEEE Press, 1990.
- [19] R. D. Falgout, A. Skjellum, S. G. Smith, and C. H. Still. The multicomputer toolbox approach to concurrent BLAS and LACS. In *Proceedings of the 1992 Scalable High Performance Computing Conference*, pages 121-128. IEEE Press, 1992.

- [20] G. C. Fox, S. W. Otto, and A. J. G. Hey. Matrix algorithms on a hypercube I: Matrix multiplication. *Parallel Computing*, 4:17-31, 1987.
- [21] S. Huss-Lederman, E. M. Jacobson, A. Tsao, and G. Zhang. Matrix multiplication on the Intel Touchstone Delta. Technical report, Supercomputing Research Center, 1993. in preparation.
- [22] Intel Corporation. *Touchstone Delta Fortran Calls Reference Manual*, April 1991.
- [23] Intel Corporation. *Touchstone Delta System User's Guide*, October 1991.
- [24] C. Lin and L. Snyder. A matrix product algorithm and its comparative performance on hypercubes. In *Proceedings of the 1992 Scalable High Performance Computing Conference*, pages 190-194. IEEE Press, 1992.
- [25] 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.

ORNL/TM-12252

INTERNAL DISTRIBUTION

1. B. R. Appleton	24-28. S. A. Raby
2. C. Bottcher	29-33. R. F. Sincovec
3. B. A. Carreras	34. G. M. Stocks
4-8. J. Choi	35. M. R. Strayer
9-10. T. S. Darland	36-40. D. W. Walker
11-15. J. J. Dongarra	41-45. R. C. Ward
16. J. B. Drake	46. P. H. Worley
17. T. H. Dunigan	47. Central Research Library
18. W. R. Emanuel	48. ORNL Patent Office
19. R. E. Flanery	49. K-25 Applied Technology Library
20. W. F. Lawkins	50. Y-12 Technical Library
21. M. R. Leuze	51. Laboratory Records Department - RC
22. R. Mann	52-53. Laboratory Records Department
23. C. E. Oliver	

EXTERNAL DISTRIBUTION

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

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

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

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

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

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

144. Andrew B. White, Los Alamos National Laboratory, P. O. Box 1663, MS-265, Los Alamos, NM 87545
145. David L. Williamson, National Center for Atmospheric Research, P. O. Box 3000, Boulder, CO 80307
146. Samuel Yee, Air Force Geophysics Lab, Department LYP, Hancom AFB, Bedford, MA 01731
147. 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
- 148-149. Office of Scientific & Technical Information, P. O. Box 62, Oak Ridge, TN 37830

END

DATE
FILMED
10 / 7 / 93