

Efficient Parallel Out-of-core Matrix Transposition*

Sriram Krishnamoorthy Gerald Baumgartner
Daniel Cociorva Chi-Chung Lam
P. Sadayappan

Department of Computer and Information Science
The Ohio State University

{krishnsr, gb, cociorva, clam, saday}@cis.ohio-state.edu

Abstract

This paper addresses the problem of parallel transposition of large out-of-core arrays. Although algorithms for out-of-core matrix transposition have been widely studied, previously proposed algorithms have sought to minimize the number of I/O operations and the in-memory permutation time. We propose an algorithm that directly targets the improvement of overall transposition time. The I/O characteristics of the system are used to determine the read, write and communication block sizes such that the total execution time is minimized. We also provide a solution to the array redistribution problem for arrays on disk. The solution to the sequential transposition problem and the parallel array redistribution problem are then combined to obtain an algorithm for the parallel out-of-core transposition problem.

1. Introduction

This paper addresses the problem of parallel out-of-core matrix transposition. The problem is viewed in terms of two sub-problems: disk-based array redistribution, followed by concurrent independent uniprocessor transposition of disk-based arrays. The same algebraic framework is used for both steps. We first address the sequential transposition problem, which has been previously studied.

Consider an $N \times N$ matrix that is stored in disk in row-major order. The system has main memory, which can hold M elements, where $M < N^2$, $M = O(N)$. The problem is to transpose the matrix stored in disk, when only a part of the matrix can be brought into memory at any time. Applications that need to access the el-

ements of a matrix in column-major order transpose the matrix in disk and then access the elements in it. Matrix transpose is a key operation in various scientific applications. For example, the multidimensional Fast Fourier transform (FFT) [2, 3] can be implemented as a series of one-dimensional FFTs, one along each dimension. For a matrix stored in disk in row-major order that is too large to fit in memory, the most effective mechanism is to transpose the matrix between the one-dimensional FFTs.

Our primary motivation for addressing the parallel out-of-core matrix transposition problem arises from the domain of electronic structure calculations using ab initio quantum chemistry models such as Coupled Cluster models. We are developing an automatic synthesis system called the Tensor Contraction Engine (TCE) [18], to generate efficient parallel programs from high level expressions for a class of computations expressible as tensor contractions [4, 7, 6, 8]. Often the tensors (essentially multi-dimensional matrices) are too large to fit in memory and must be disk-resident.

The optimized parallel programs synthesized by the tool often have to take as input large disk-resident tensors created by other software packages, such as the NWChem computational chemistry suite [14]. For efficient execution, the TCE-synthesized program might need to store and accesses the disk-resident tensors in a very different order than that used by the producer program. Efficient transformation of the data from the available format to the required format is required through transposition and/or re-blocking. In addition, when TCE-synthesized code is used on different machines, different transformations are required on the data produced by packages like NWChem, requiring efficient out-of-core matrix transposition and transformation algorithms.

This problem has been widely studied in the literature. A simple in-place element-wise approach to trans-

* Supported in part by the National Science Foundation through the Information Technology Research program (CHE-0121676)

pose the matrix is prohibitively expensive. The block transposition algorithm transposes the array in a single pass in $O(N^{3/2})$ I/O operations. An in-place transposition algorithm requiring $O(N \log N)$ disk accesses was proposed by Eklundh [11]. This algorithm requires at least two rows to fit in memory. Extensions to the algorithm for rectangular matrices were presented in [1, 16, 19]. Kaushik et al. [12] proposed an out-of-place algorithm that improves upon these algorithms by reducing the number of read operations. Suh and Prasanna [17] reduced the in-memory permutation time by using *collect* buffers, instead of in-memory permutation, in addition to reducing the number of I/O operations. Their algorithm combines writes and collects the rows to be permuted in subsequent passes.

All these studies use the number of I/O operations as the primary optimization metric. Although the execution time of the solution provided has been improved by all these efforts, the total execution time has not been used as the primary metric for optimization. A reduction in the number of I/O operations, in most cases, translates to larger sizes of I/O blocks. The importance given to reducing the number of I/O operations is due to the fact that the seek time for the disk head is very large (of the order of several milliseconds) compared to the per-byte transfer time (of the order of microseconds or less). If the I/O blocks read/written are relatively small, the total number of I/O operations is indeed a suitable optimization metric. However, when the I/O blocks get large, the data transfer time becomes significant and can dominate the total access time. Since previously proposed algorithms for out-of-core transposition have focused on reducing the number of I/O operations, they can become sub-optimal when large block transfers are involved.

Cormen et al. [9] solve the problem based on the parallel disk model (PDM) [20]. PDM handles the read and write block sizes as equivalent, while the I/O characteristics of reads and writes can differ widely. PDM uses the number of I/O operations as the metric, where the size of each I/O is determined by the layout of data on disk. It does not take into account the effect of read-ahead and request reordering in the I/O subsystem.

All the algorithms in the literature determine the fundamental unit of I/O based on the size of the matrix, i.e., they are data-centric. The basic unit of I/O operation in these algorithms is one row of the matrix or a multiple thereof. They do not adapt to the I/O characteristics of the system. In contrast, the approach proposed here takes into account the empirically determined I/O characteristics of the disk and file system in determining the parameters of the algorithm. The basic unit of I/O is not a row, but is determined by the I/O characteristics and the instance of the problem at hand. The execution time of the

algorithm on the system is estimated based on the experimentally observed I/O characteristics. The parameters that minimize the execution time are chosen.

In this paper, we do not discuss the sequential transposition algorithm in detail due to lack of space. For a detailed analysis, please refer to [13].

The paper is organized as follows. The I/O characteristics of two systems are discussed in Section 2. In Section 3 the transposition problem is formulated using the matrix-vector product notation. The sequential transposition algorithm is described in Section 4. The algorithm is extended to parallel systems in Section 5. Experimental results are presented in Section 6. Section 7 concludes the paper.

2. I/O Characteristics

We studied the variation of read and write access times with changes in size and stride of I/O on a commodity PC (henceforth referred to as *PC*) and an HP zx6000 workstation. Their configurations are shown in Table 1. The PC and the HP zx6000 workstation are part of a the IA32 and the Itanium 2 cluster, respectively, at the Ohio Supercomputer Center (OSC) [15].

For both the systems we observe that above a particular block size the stride does not affect the per-byte transfer cost. We expect this observation to hold across a wide variety of systems. These block sizes, above which the per-byte read and write times are not affected by the stride of access, will henceforth be referred to as the *read* and *write thresholds* respectively. These parameters vary depending on the system under consideration and the per-byte read and write costs can saturate at different block sizes.

An important consequence of this observation is that if the thresholds are smaller than N , the size of the matrix, fractions of a row can be read and written without any additional penalty, irrespective of the stride of access. This reduction in the read and write block size in turn decreases the amount of work involved in transposing an array as will be explained later. In the extreme case, if each element is large enough to allow efficient I/O of individual elements, a simple single-pass element-wise transposition would be efficient.

This observation also shows that for I/O sizes above the threshold the number of I/O operations does not reflect the actual performance of the algorithm. An algorithm might involve more I/O operations but be faster than another algorithm with fewer I/O operations due to this effect.

System	Configuration			
	Processor	Memory	OS	Compiler
commodity PC	Dual AMD Athlon MP (1.533 GHz)	2GB	linux 2.4.20	pgcc 4.0-2
HP zx6000	Dual Itanium-2(900 MHz)	4GB	linux 2.4.18	gcc 2.96

Table 1. Configuration of the systems used for I/O characterization.

3. Matrix Vector Product Formulation of Transposition Algorithms

In this section, matrix transposition algorithms are formulated based on the matrix-vector notation used in [10]. This section provides a generic formulation for transposition algorithms.

Transposition of a matrix can be viewed as an interchange of the indices of the matrix.

$$T(i, j) = (j, i)$$

where i is the row index and j is the column index. This is a particular instance of a general class of index transformation algorithms.

Each element of the array on disk has a linear address obtained by concatenating the column index bits to the row index bits. This is the address upon which the permutation is applied. The transformation of the address vector using a permutation matrix corresponds to the permutation of the address vector and hence the matrix. The linear address of an element in the array contains $2n$ bits and hence the permutation matrix contains $2n$ rows.

The identity of the transformation is $\begin{pmatrix} I_n & 0 \\ 0 & I_n \end{pmatrix}$. Matrix transposition is defined as the permutation of the address vector i

$$i \rightarrow Ti$$

where T is the transformation matrix $\begin{pmatrix} 0 & I_n \\ I_n & 0 \end{pmatrix}$.

We use the following notation in the discussion.

$$A \oplus B = \begin{pmatrix} A & 0 \\ 0 & B \end{pmatrix} \quad (1)$$

$$L(A, B) = \begin{pmatrix} 0 & B \\ A & 0 \end{pmatrix} \quad (2)$$

Thus, $L(I_n, I_n)$ is the desired permutation. Since the entire array does not fit in memory, $L(I_n, I_n)$ is factorized into a number of permutation matrices such that the transformation effected by each of the matrices can be done with the memory available.

Any out-of-core matrix transposition algorithm consists of three phases — read, permute and write. Each phase corresponds to a permutation matrix. These

phases are repeated on disjoint sets of data in the different steps of each pass. The algorithm might involve many passes, each operating on the entire array. Thus, out-of-core matrix transformation algorithms are of the form

$$T = L(I_n, I_n) = \prod_{i=t-1}^{i=0} W_i P_i R_i$$

where W_i is the permutation matrix corresponding to a write, R_i is a permutation matrix corresponding to a read and P_i corresponds to in-memory permutation. The algorithms under this formulation read some data, permute it in memory, and write the data to disk before reading data for the next step in the same pass. t specifies the number of passes. Thus each algorithm is defined by the parameters t, W_i, P_i and R_i , where the suffix i is used to refer to the permutations in the i th pass. Each algorithm can also have additional parameters.

Some restrictions apply to the possible values of W_i, P_i and R_i . These restrictions are induced by the memory constraint involved in the algorithm. Each permutation matrix must correspond to a transformation of the given matrix that can be done with the memory available. Thus, each step of the algorithm can operate on at most M elements. In particular, W_i, P_i and R_i must be expressed as

$$R_i = A_{2*n-r} \oplus I_r \quad r \leq m$$

$$P_i = I_{2*n-m} \oplus B_m$$

$$W_i = C_{2*n-w} \oplus I_w \quad w \leq m$$

The algorithm reads $R = 2^r$ elements and writes $W = 2^w$ elements in one I/O operation. Henceforth, we use the term read(write) threshold to refer to $R(W)$ and the least significant $r(w)$ rows of the permutation matrix, interchangeably. The reference will be clear from the context. Also note that P_i can permute only data corresponding to elements in memory. Given these parameters for an algorithm it can be implemented as

Algorithm 1: Generic Transposition Algorithm

- (1) **for** $i = 0$ **to** $t - 1$
- (2) **for** $j = 0$ **to** $2^{2*n-m} - 1$
- (3) Read M elements at address $R_i^{-1}(j)$
/*Might involve multiple I/O operations*/
- (4) Permute data in memory according to P_i .
- (5) Write M elements at address $W_i(j)$
/*Might involve multiple I/O operations*/

The read (write) may involve multiple I/O operations each of size at least $2^r(2^w)$ elements.

For a discussion of performance of different transposition algorithms based on these I/O characteristics refer to [13].

4. Sequential Transposition Algorithm

Our algorithm is based on estimating the total transposition time and choosing parameters for the algorithm that optimize it. The observation that an increase in I/O size beyond the threshold does not influence the performance of the algorithm is exploited. There is a trade-off between the I/O size and the number of passes the algorithm requires. The smaller the I/O size, the more the algorithms approach the block-transposition algorithm and hence run in a smaller number of passes. However, reducing the I/O size below the threshold increases the I/O time above the minimum possible.

The transposition time can be written as

$$Time_{total} = \sum_{i=0}^{i=t-1} Time_{Read} + Time_{Permute} + Time_{Write}$$

The read and write times for each pass can be computed from the stride and block size of the I/O operation. Estimating the permutation time is more difficult as it depends on the exact permutation involved. Unlike the I/O characteristics of a system, which can be determined independent of any specific algorithm, the permutation characteristic for each algorithm has to individually determined. Here, we determine the best parameters for the algorithm that optimize the total I/O time. The characteristics of the algorithm allow for optimizing the in-memory permutation, as will be discussed later.

The algorithm has three parameters, namely the factors to be permuted in each pass, the read and write block size and the number of passes. The read and write block sizes are chosen close to the threshold in order to optimize the total I/O time. The most common case in which the I/O block size is chosen to be smaller than the threshold is when such a choice reduces the number of passes and offsets the additional cost incurred due to the smaller I/O size.

The I/O permutations are of the form $A \oplus I_n$, while the required permutation $L(I_n, I_n)$ involves exchanging the upper and lower n address elements in the address vector. The nature of the I/O permutation prevents any effective permutation from being done in the read and write phases. The I/O phases ‘gather’ data to be permuted and ‘scatter’ the result of the permutation. In previous algorithms, the basic unit of I/O was a row. In our algorithm, the I/O block size could be smaller than N , say $B = 2^b$, in which case the exchange $(0 : b - 1) \leftrightarrow (n : n + b - 1)$ can be done in the read and/or write phases. This reduces the number of address vector elements to be permuted in the permutation phase and might result in a reduction in the number of passes, and hence significantly reduce transposition time.

Our algorithm is formulated as shown below. The unit of each read and write is 2^r and 2^w elements respectively. Except in the first pass, the algorithm reads M elements in each read operation. In the first pass, the read and write phases permute the address vector elements $(n : n + r - 1)$ and $(r : r + w - 1)$, respectively, to their appropriate positions.

The conditions to be satisfied are:

$$\begin{aligned} n &\geq w \\ m &\geq r \geq w \\ m &> w \end{aligned}$$

$i = 0$

Case 1: $r \leq n$

$$\begin{aligned} s_0 &= \min((m - n), w) \\ s &= s_0 \\ R_0 &= (I_{2n}) \\ P_0 &= (I_{n-s} \oplus L(I_s, L(I_{n-s}, I_s))) \\ W_0 &= (L(L(I_{n-s}, I_s), I_{n-w}) \oplus I_w) \end{aligned}$$

Case 2: $r > n$

$$\begin{aligned} s_0 &= \min((m - r), w) \\ s &= s_0 \\ R_i &= (I_{n-s} \oplus L(I_s, I_{n-r}) \oplus I_r) \\ P_i &= (I_{2*n-(r+s)} \oplus L(I_s, L(I_{r-s}, I_s))) \\ W_i &= (L(I_{n-s}, I_{n-r} \oplus L(I_s, I_{r-w})) \oplus I_w) \end{aligned}$$

$0 < i \leq t - 1$

$$\begin{aligned} s_i &= (w - s_0) \bmod (m - w) \text{ if } i = t - 1 \text{ and} \\ &\quad (w - s_0) \bmod (m - w) \neq 0 \\ s_i &= \lfloor (w - s_0) / (m - w) \rfloor \text{ otherwise} \\ sp_i &= \prod_{j=0}^{j=i} s_j \\ R_i &= (I_{2n}) \\ P_i &= (I_{2*n-(w+s_i)} \oplus L(I_{s_i}, L(I_{w-sp_i-s_i}, I_{s_i}))) \oplus \\ &\quad I_{sp_i} \\ W_i &= (I_{n-w} \oplus L(I_{s_{i-1}} \oplus I_{n-s_{i-1}-s_i}, I_{s_i}) \oplus \\ &\quad I_w) \end{aligned}$$

With increasing memory size, a modification of the I/O parameters provides diminishing improvements, unless it results in a reduction in the number of passes. Greater improvements can be obtained if the additional memory available is used to improve permutation time. Kaushik et al. perform an in-place in-memory permutation. Suh and Prasanna use collect buffers to collect data to be written in each write operation. The locality of the permutation operation can be improved by optimizations such as blocking.

We use collect operations to perform the permutation, as this was empirically found to take less time than in-memory permutation. The permutation involved in the first pass is similar to transposition. Since the naive element-wise approach or the collect operation have poor cache performance, the permutation was done out-of-place in-memory. The I/O size was further reduced in order to maintain the number of passes.

5. Parallel Out-of-Core Matrix Transposition

In this section, the problem of transposing an out-of-core array distributed among multiple processors is discussed. Each processor has a local disk and the array is distributed among the processors in a row-blocked fashion. The required distribution of the transposed array among the processors is specified.

In the following discussion, we first formulate the representation of an array distributed among multiple processors. Then an algorithm is provided for redistributing out-of-core arrays in a parallel system. To our knowledge the problem of parallel out-of-core array redistribution has not been addressed previously.

The array redistribution mechanism and the sequential transposition algorithm are combined to describe the out-of-core transposition algorithm for arrays distributed among multiple processors.

5.1. Formulation for Arrays Distributed among Multiple Processors

The arrays are assumed to be distributed in a regular fashion so that some of the elements in the address vector represent the processor identifier. This corresponds to a mapping of the elements of the array to a sequence of processors. A row-blocked distribution is obtained when the most significant elements in the address vector represent the processor identifier. A cyclic distribution is obtained when the least significant elements of the address vector represent the processor identifier.

We define the linear address vector of an element in the array to be the concatenation of the address vector of

the element in the local disk to the processor identifier. This view preserves the notion of contiguity of elements which differ in the lower most elements of the address vector, analogous to the sequential formulation. Hence the formulation can represent read and write thresholds in the address vector and the access pattern that can take advantage of prefetching as well.

Given that the uppermost elements in the linear address vector correspond to the processor identifier, the distribution of the array among multiple processors corresponds to choosing a set of elements in the address vector to become the uppermost elements. Hence array distribution among multiple processors can be viewed as a permutation of the linear address space of the array. The identity of array distribution is I_n , which corresponds to a row-blocked distribution. Any other distribution of data among processors is viewed as a permutation on the row-blocked distribution. For example, a cyclic distribution of an array among two processors corresponds to the following permutation:

$$\begin{pmatrix} 0 & 1 \\ I_{2^{*n-1}} & 0 \end{pmatrix}$$

5.2. Array Redistribution Problem

The array redistribution problem is stated as follows: Given an array distributed among processors, represented by a permutation matrix, achieve a target distribution corresponding to a new permutation.

The array redistribution problem brings with it another cost factor in the form of communication. Communication cost varies linearly and is modeled as $T_s + m * T_b$, where T_s is the startup cost, m the message size and T_b the per-byte transfer cost. Depending on the parameters T_s and T_b of a communication protocol, beyond a message size m , the transfer cost dominates the startup cost and the average per-byte cost converges to a constant. The message size beyond which there is little change in the communication cost is called the communication threshold 2^c . This is the minimum message size for communication. Note that as in the case of the read and write thresholds, the message size chosen for a specific instance of an algorithm may be below the threshold, if it cannot be improved upon. The communication characteristics of various systems have been widely studied and we do not discuss them here. For the following discussion, it is assumed that there are 2^p processors. The uppermost p rows of any permutation matrix correspond to the elements that constitute the processor identifier. The lowermost c elements of the address vector correspond to the communication threshold. The terms read, write and communication thresholds will be used interchangeably to refer to the size of I/O and r , w and

c least significant elements in the address vector respectively. The reference will be clear from the context.

The formulation of the parallel redistribution involves four permutation matrices — read, write, in-memory permutation and communication. Extending the template for the formulation of read, write and in-memory permutation discussed in Section 3 to the parallel domain we get

$$\begin{aligned} R_i &= I_p \oplus A_{2*n-r-p} \oplus I_r \quad r \leq m \\ P_i &= I_{2*n-m} \oplus B_m \\ W_i &= I_p \oplus C_{2*n-w-p} \oplus I_w \quad w \leq m \end{aligned}$$

which indicates that R_i , W_i and P_i cannot permute the elements corresponding to the processor identifier. Only communication can permute the elements corresponding to the processor identifier. The permutation corresponding to communication is of the form

$$C = \begin{pmatrix} D_{2*n-c} & 0 \\ 0 & I_c \end{pmatrix}$$

where D describes the permutations done by communication.

Note that there are some restrictions on C similar to those on R_i , W_i and P_i as discussed in Section 3. C cannot permute between address elements corresponding to in-memory and out-of-memory data (the elements corresponding to the processor identifier are special and will be discussed below). Any permutation except those involving the the processor identifier can be performed by P_i and W_i . Therefore, we place additional restrictions on C , so that it can only involve permutations required to change the processor identifier. In most cases, c is smaller than r and w and we assume the same.

Array redistribution can involve permutations of three kinds. First is the exchange of elements that are part of the processor identifier. This effect is achieved by an exchange of all the data between processors. An equivalent effect could be achieved by relabeling the processors. But this does not obviate the problem as the same situation arises when there are multiple arrays which are aligned with respect to one another. This, or other constraints, might involve such an exchange that cannot be handled by relabeling.

Second is the exchange involved when elements within the communication threshold are to become part of the processor identifier. Any permutation involving the elements beyond the communication threshold is performed by an all-to-all personalized collective communication operation. If the number of elements within the communication threshold that are to become elements corresponding to the processor identifier is greater than $m - c$, then a sequence of in-memory permutation and communication operations are carried

out. Each in-memory permutation operation moves as many elements from within the communication threshold to be beyond the threshold as possible. These elements are then made part of the processor identifier by a scatter operation. This process is repeated until there are no more elements in the least significant c address elements that are to be part of the processor identifier.

Thus any element already part of the processor identifier or within the least significant m elements (memory size), that is to be part of the processor identifier can be made part of the processor identifier in a single pass, where a single pass is defined as all operations between the read of a data element from disk and its write to disk.

A more complicated operation is required when trying to permute the elements corresponding to the processor identifier and the address elements which are beyond the least significant m elements. This involves a collect operation by each processor. The difference in handling this case and the previous two cases is that in the previous two cases all processors do the same operations throughout each pass. In this case, each processor collects all the data in memory from certain other processors in turn, in different iterations of the loop. But since all the collected data cannot be stored in memory, the data received from every processor is written to disk. This breaks the clear demarcation between the communication and write operations as they become interleaved. Since handling this case essentially involves writing the data to disk, this case is handled last.

Hence all communication required to handle array redistribution can be done in a single pass. But note that this may not be the most efficient way of performing the array redistribution. In handling the last case, each processor might receive data from a different set of processors in different iterations. Each receive is separated by a write to disk. Hence the communication and write times cannot overlap and could lead to very poor execution time especially when the number of processors is large.

We handle the last case by a sequence of in-memory permutation and write operations which move the address elements above the communication threshold but within the memory size. These are now handled as in the second case. This would increase the number of passes, but the total execution time might be reduced as the communication and write times overlap. The exact operation is determined by the relative cost of communication and I/O in the system under consideration and by determining the sequence of operations that minimizes the overall execution time.

5.3. Combining Array Redistribution and Sequential Matrix Transposition

In this section, we combine the mechanisms considered until now to derive an algorithm for transposing out-of-core matrices which are distributed in a row-blocked fashion among multiple processors.

The only major change is in the first pass, into which the array distribution phase is merged. Subsequent passes are similar to that in the sequential case and involve only local permutation. The first pass for the parallel algorithm is as follows:

- Read as in sequential case (R_i).
- Perform in-memory permutation as in sequential case P_i .
- Perform array distribution, handling the different cases discussed above.
- Write data to disk, if this has not yet been done as part of the previous step.

Note that the parallel case does not lead to an increase in the number of passes in the form of additional reads or writes.

6. Experimental Results

In this section, we discuss the results obtained from implementing the parallel transposition algorithm. For results pertaining to sequential transposition, please refer to [13]. The transposition times were measured on the Itanium 2 cluster and on the IA32 cluster at the Ohio Supercomputer Center. Each machine in the Itanium 2 cluster is an HP zx6000 workstation and each machine in the IA32 cluster is the commodity PC discussed in Section 2. Both clusters use the Myrinet [5] interconnection network. The implementation was out-of-place and used an auxiliary array.

The transposition time for different memory sizes and numbers of processors was measured. Tables 2 and 3 show the transposition times on the Itanium 2 cluster for array sizes of 16GB ($N=64K$) and 64GB ($N=128K$). Table 4 shows the transposition times on the IA32 cluster for an array size of 16GB ($N=64K$).

In both systems the read threshold was much higher than N . So the execution was influenced mainly by the write threshold. The results show that the effect of memory size on the execution time is more pronounced for specific memory sizes. This is due to a reduction in the number of passes required. Note that even in cases without a significant change in execution time, corresponding to an increase in memory size the execution time improves. This is due to a reduction in the stride of writes.

#procs	Memory size (MB)					
	16	32	64	128	256	512
1	3406	3322	2265	2230	2003	2079
2	1536	1127	962	949	984	1006
4	740	542	484	483	475	474

Table 2. Execution time, in seconds, on the Itanium 2 cluster. Array size is 16GB ($N=64K$).

#procs	Memory size (MB)					
	16	32	64	128	256	512
4	3448	3252	3213	2102	2907	2801
8	1470	1533	1469	921	985	1007

Table 3. Execution time, in seconds, on the Itanium 2 cluster. Array size is 64GB ($N=128K$).

The write block size is reduced to be below the write threshold if it reduces the total execution time, for example in Table 2 for memory sizes of 128MB and 64MB and one processor.

The parallel algorithm scales well with the number of processors. A slightly superlinear speedup can be seen in some cases. This is due to improved cache locality in access. Note that for the 4 processor case, the portion of each array in a processor is 4GB, equal to its memory size. But since there are three arrays the arrays are not fully cached in memory, making the results dependent on the caching mechanism. In some cases, an increase in the number of processors reduces the number of passes thus significantly reducing the execution time. This effect can be observed in Table 3 for a memory size of 64MB.

#procs	Memory size (MB)					
	16	32	64	128	256	512
1	6448	5680	3229	4138	4253	4056
2	4009	2059	2185	2182	2280	2171
4	2008	981	1135	1167	1186	1140
8	995	583	664	688	638	517

Table 4. Execution time, in seconds, on the IA32 cluster. Array size is 16GB ($N=64K$).

7. Conclusions

In this paper, we have addressed the efficient parallel out-of-core transposition of matrices that are too large to fit in main memory. The problem was cast as a composition of two sub-problems: disk-based array redistribution, followed by concurrent independent uniprocessor transposition of disk-based arrays. The same algebraic framework was used for both steps. By viewing the transposition problem as an index permutation on the addresses of matrix elements, effective use was made of available main memory in optimizing the overall transposition time, rather than reducing the number of I/O operations, as previous algorithms have done. A solution to the out-of-core array redistribution problem was then provided using the same algebraic framework, combining to provide an algorithm for parallel out-of-core matrix transposition. Experimental measurements were provided, demonstrating the scalability of the proposed approach and the limited communication overhead. Extensions of this framework are being pursued for efficient index permutation of multi-dimensional arrays on parallel systems.

Acknowledgments

We would like to thank the Ohio Supercomputer Center (OSC) for the use of their computing facilities.

References

- [1] W. O. Alltop. A computer algorithm for transposing nonsquare arrays. *IEEE Transactions on Computers*, 24(10):1038–1040, 1975.
- [2] G. L. Anderson. A stepwise approach to computing the multidimensional fast Fourier transform of large arrays. *IEEE Transactions on Acoustics and Speech Signal Processing*, 28(3):280–284, 1980.
- [3] D. H. Bailey. FFTs in external or hierarchical memory. *Journal of Supercomputing*, 4(1):23–35, 1990.
- [4] G. Baumgartner, D. Bernholdt, D. Cociorva, R. Harrison, S. Hirata, C. Lam, M. Nooijen, R. Pitzer, J. Ramanujam, and P. Sadayappan. A high-level approach to synthesis of high-performance codes for quantum chemistry. In *Proceedings of Supercomputing 2002*, 2003.
- [5] N. J. Boden, D. Cohen, R. E. Felderman, A. E. Kulawik, C. L. Seitz, J. N. Seizovic, and W. Su. Myrinet: A gigabit-per-second local area network. *IEEE Micro*, 15(1):29–36, February 1995.
- [6] D. Cociorva, G. Baumgartner, C. Lam, P. Sadayappan, J. Ramanujam, M. Nooijen, D. Bernholdt, , and R. Harrison. Space-time trade-off optimization for a class of electronic structure calculations. In *Proc. of ACM SIG-PLAN 2002 Conference on Programming Language Design and Implementation (PLDI)*, 2002.
- [7] D. Cociorva, X. Gao, S. Krishnan, G. Baumgartner, C. Lam, P. Sadayappan, and J. Ramanujam. Global communication optimization for tensor contraction expressions under memory constraints. In *Proc. of 17th International Parallel & Distributed Processing Symposium (IPDPS)*, 2003.
- [8] D. Cociorva, J. Wilkins, G. Baumgartner, P. Sadayappan, J. Ramanujam, M. Nooijen, D. Bernholdt, and R. Harrison. Towards automatic synthesis of high-performance codes for electronic structure calculations: Data locality optimization. In *Proc. of the Intl. Conf. on High Performance Computing*, 2001.
- [9] T. H. Cormen, T. Sundquist, and L. F. Wisniewski. Asymptotically tight bounds for performing BMMC permutations on parallel disk systems. *SIAM Journal on Computing*, 28(1):105–136, 1998.
- [10] A. Edelman, S. Heller, and S. L. Johnsson. Index transformation algorithms in a linear algebra framework. *IEEE Transactions on Parallel and Distributed Systems*, 5(12):1302–1309, 1994.
- [11] J. O. Eklundh. A fast computer method for matrix transposing. *IEEE Transactions on Computers*, 20(7):801–803, 1972.
- [12] S. D. Kaushik, C.-H. Huang, R. W. Johnson, P. Sadayappan, and J. R. Johnson. Efficient transposition algorithms for large matrices. In *Proceedings of the 1993 ACM/IEEE conference on Supercomputing*, pages 656–665. ACM Press, 1993.
- [13] S. Krishnamoorthy, G. Baumgartner, D. Cociorva, C. Lam, and P. Sadayappan. On efficient out-of-core matrix transposition. Technical Report OSU-CIRSC-9/03-T52, School of Computer and Information Science, The Ohio State University, Sept 2003.
- [14] NWChem. <http://www.emsl.pnl.gov:2080/docs/nwchem/nwchem.html>.
- [15] Ohio Supercomputing Center. <http://www.osc.edu>.
- [16] H. K. Ramapriyan. A generalization of Eklundh’s algorithm for transposing large matrices. *IEEE Transactions on Computers*, 24(12):1221–1226, 1975.
- [17] J. Suh and V. K. Prasanna. An efficient algorithm for out-of-core matrix transposition. *IEEE Transactions on Computers*, 51(4):420–438, April 2002.
- [18] Synthesis of High-Performance Algorithms for Electronic Structure Calculations. <http://www.cis.ohio-state.edu/saday/TCE/index.html>.
- [19] R. E. Twogood and M. P. Ekstrom. An extension of Eklundh’s matrix transposition algorithm and its application to digital signal processing. *IEEE Transactions on Computers*, 25(12):950–952, 1976.
- [20] J. S. Vitter and E. A. M. Shriver. Algorithms for parallel memory I: Two-level memories. *Algorithmica*, 12(2–3):110–147, 1994.