

Dartmouth College

## Dartmouth Digital Commons

---

Computer Science Technical Reports

Computer Science

---

1994

# Asymptotically Tight Bounds for Performing BMMC Permutations on Parallel Disk Systems

Thomas H. Cormen  
*Dartmouth College*

Thomas Sundquist  
*Dartmouth College*

Leonard F. Wisniewski  
*Dartmouth College*

Follow this and additional works at: [https://digitalcommons.dartmouth.edu/cs\\_tr](https://digitalcommons.dartmouth.edu/cs_tr)



Part of the [Computer Sciences Commons](#)

---

### Dartmouth Digital Commons Citation

Cormen, Thomas H.; Sundquist, Thomas; and Wisniewski, Leonard F., "Asymptotically Tight Bounds for Performing BMMC Permutations on Parallel Disk Systems" (1994). Computer Science Technical Report PCS-TR94-223. [https://digitalcommons.dartmouth.edu/cs\\_tr/96](https://digitalcommons.dartmouth.edu/cs_tr/96)

This Technical Report is brought to you for free and open access by the Computer Science at Dartmouth Digital Commons. It has been accepted for inclusion in Computer Science Technical Reports by an authorized administrator of Dartmouth Digital Commons. For more information, please contact [dartmouthdigitalcommons@groups.dartmouth.edu](mailto:dartmouthdigitalcommons@groups.dartmouth.edu).

# Asymptotically Tight Bounds for Performing BMMC Permutations on Parallel Disk Systems

Thomas H. Cormen

Thomas Sundquist

Leonard F. Wisniewski

Department of Mathematics and Computer Science

Dartmouth College

## Abstract

We give asymptotically equal lower and upper bounds for the number of parallel I/O operations required to perform bit-matrix-multiply/complement (BMMC) permutations on parallel disk systems. In a BMMC permutation on  $N$  records, where  $N$  is a power of 2, each  $(\lg N)$ -bit source address  $x$  maps to a corresponding  $(\lg N)$ -bit target address  $y$  by the matrix equation  $y = Ax \oplus c$ , where matrix multiplication is performed over  $GF(2)$ . The characteristic matrix  $A$  is  $(\lg N) \times (\lg N)$  and nonsingular over  $GF(2)$ . Under the Vitter-Shriver parallel-disk model with  $N$  records,  $D$  disks,  $B$  records per block, and  $M$  records of memory, we show a universal lower bound of  $\Omega\left(\frac{N}{BD} \left(1 + \frac{\text{rank } \gamma}{\lg(M/B)}\right)\right)$  parallel I/Os for performing a BMMC permutation, where  $\gamma$  is the lower left  $\lg(N/B) \times \lg B$  submatrix of the characteristic matrix. We also present an algorithm that uses at most  $\frac{2N}{BD} \left(\left\lceil \frac{\text{rank } \gamma}{\lg(M/B)} \right\rceil + 2\right)$  parallel I/Os, which asymptotically matches the lower bound and improves upon the BMMC and bit-permute/complement (BPC) algorithms in [4]. When  $\text{rank } \gamma$  is low, this method is an improvement over the general-permutation bound of  $\Theta\left(\frac{N}{BD} \lg(N/B)\right)$ .

We introduce a new subclass of BMMC permutations, called memoryload-dispersal (MLD) permutations, which can be performed in one pass. This subclass, which is used in the BMMC algorithm, extends the catalog of one-pass permutations appearing in [4].

Although many BMMC permutations of practical interest fall into subclasses that might be explicitly invoked within the source code, we show how to detect in at most  $N/BD + \left\lceil \frac{\lg(N/B)+1}{D} \right\rceil$  parallel I/Os whether a given vector of target addresses specifies a BMMC permutation. Thus, one can determine efficiently at run time whether a permutation to be performed is BMMC and then avoid the general-permutation algorithm and save parallel I/Os by using our algorithm.

## 1 Introduction

From both the theoretical and practical points of view, permuting is an interesting and important problem when the data reside on disk. As one of the most basic data-movement operations,

---

Portions of this research were performed while Tom Cormen was at the MIT Laboratory for Computer Science and appear in [3]. He was supported in part by the Defense Advanced Research Projects Agency under Grant N00014-91-J-1698. Other portions of this research were performed while at Dartmouth College and were supported in part by funds from Dartmouth College and in part by the National Science Foundation under Grant CCR-9308667. Tom Sundquist was supported in part by funds from Dartmouth College. Len Wisniewski was supported in part by INFOSEC Grant 3-56666 and in part by the National Science Foundation under Grant CCR-9308667. An extended abstract of this paper appeared in the Proceedings of the 5th Annual ACM Symposium on Parallel Algorithms and Architectures.

permuting is central to the theory of I/O complexity. The problems that we attack with supercomputers are ever-increasing in size, and in several applications matrices and vectors exceed the memory provided by even the largest supercomputers. One solution is to store large matrices and vectors on parallel disk systems. The high latency of disk accesses makes it essential to minimize the number of disk I/O operations. Permuting the elements of a matrix or vector is a common operation, particularly in the data-parallel style of computing, and good permutation algorithms can provide significant savings in disk-access costs over poor ones when the data reside on parallel disk systems.

This paper examines the class of bit-matrix-multiply/complement (BMMC) permutations for parallel disk systems and derives four important results:

1. a universal lower bound for BMMC permutations,
2. an algorithm for performing BMMC permutations whose I/O complexity asymptotically matches the lower bound, thus making it asymptotically optimal,
3. an efficient method for determining at run time whether a given permutation is BMMC, thus allowing us to use the BMMC algorithm if it is, and
4. a new subclass of BMMC permutations, memoryload-dispersal (MLD) permutations, which we show how to perform in one pass.

Depending on the exact BMMC permutation, our asymptotically optimal bound may be significantly lower than the asymptotically optimal bound proven for general permutations. Moreover, the low constant factor in our algorithm makes it very practical.

## Model and previous results

We use the parallel-disk model first proposed by Vitter and Shriver [13, 14], who also gave asymptotically optimal algorithms for several problems including sorting and general permutations. In the Vitter-Shriver model,  $N$  records are stored on  $D$  disks  $\mathcal{D}_0, \mathcal{D}_1, \dots, \mathcal{D}_{D-1}$ , with  $N/D$  records stored on each disk. The records on each disk are organized in *blocks* of  $B$  records each. When a disk is read from or written to, an entire block of records is transferred. Disk I/O transfers records between the disks and a *random-access memory* (which we shall refer to simply as “memory”) capable of holding  $M$  records. Each *parallel I/O operation* transfers up to  $D$  blocks between the disks and memory, with at most one block transferred per disk, for a total of up to  $BD$  records transferred. We assume *independent I/O*, in which the blocks accessed in a single parallel I/O may be at any locations on their respective disks, as opposed to *striped I/O*, which has the restriction that the blocks accessed in a given operation must be at the same location on each disk.

We measure an algorithm’s efficiency by the number of parallel I/O operations it requires. Although this cost model does not account for the variation in disk access times caused by head movement and rotational latency, programmers often have no control over these factors. The number of disk accesses, however, can be minimized by carefully designed algorithms such as those in [3, 4, 7, 9, 10, 11, 13, 14] and this paper.


For convenience, we use the following notation extensively:

$$b = \lg B, \quad d = \lg D, \quad m = \lg M, \quad n = \lg N.$$

We shall assume that  $b, d, m$ , and  $n$  are nonnegative integers, which implies that  $B, D, M$ , and  $N$  are exact powers of 2. In order for the memory to accomodate the records transferred in a parallel I/O operation to all  $D$  disks, we require that  $BD \leq M$ . Also, we assume that  $M < N$ ,

|          | $\mathcal{D}_0$ |    | $\mathcal{D}_1$ |    | $\mathcal{D}_2$ |    | $\mathcal{D}_3$ |    | $\mathcal{D}_4$ |    | $\mathcal{D}_5$ |    | $\mathcal{D}_6$ |    | $\mathcal{D}_7$ |    |
|----------|-----------------|----|-----------------|----|-----------------|----|-----------------|----|-----------------|----|-----------------|----|-----------------|----|-----------------|----|
| stripe 0 | 0               | 1  | 2               | 3  | 4               | 5  | 6               | 7  | 8               | 9  | 10              | 11 | 12              | 13 | 14              | 15 |
| stripe 1 | 16              | 17 | 18              | 19 | 20              | 21 | 22              | 23 | 24              | 25 | 26              | 27 | 28              | 29 | 30              | 31 |
| stripe 2 | 32              | 33 | 34              | 35 | 36              | 37 | 38              | 39 | 40              | 41 | 42              | 43 | 44              | 45 | 46              | 47 |
| stripe 3 | 48              | 49 | 50              | 51 | 52              | 53 | 54              | 55 | 56              | 57 | 58              | 59 | 60              | 61 | 62              | 63 |

**Figure 1:** The layout of  $N = 64$  records in a parallel disk system with  $B = 2$  and  $D = 8$ . Each box represents one block. The number of stripes is  $N/BD = 4$ . Numbers indicate record indices.



**Figure 2:** Parsing the address  $x = (x_0, x_1, \dots, x_{n-1})$  of a record on a parallel disk system. Here,  $n = 13$ ,  $b = 3$ ,  $d = 4$ ,  $m = 8$ , and  $s = 6$ . The least significant  $b$  bits contain the offset of a record within its block, the next  $d$  bits contain the disk number, and the most significant  $s$  bits contain the stripe number. The most significant  $n - m$  bits form the record's memoryload number, and bits  $b, b + 1, \dots, m - 1$  form the relative block number, used in Section 3.

since otherwise we can just perform all operations in memory. These two requirements imply that  $b + d \leq m < n$ .

The Vitter-Shriver model lays out data on a parallel disk system as shown in Figure 1. A *stripe* consists of the  $D$  blocks at the same location on all  $D$  disks. We indicate the *address*, or *index*, of a record as an  $n$ -bit vector  $x$  with the least significant bit first:  $x = (x_0, x_1, \dots, x_{n-1})$ . Record indices vary most rapidly within a block, then among disks, and finally among stripes. As Figure 2 shows, the offset within the block is given by the least significant  $b$  bits  $x_0, x_1, \dots, x_{b-1}$ , the disk number by the next  $d$  bits  $x_b, x_{b+1}, \dots, x_{b+d-1}$ , and the stripe number by the  $s = n - (b + d)$  most significant bits  $x_{b+d}, x_{b+d+1}, \dots, x_{n-1}$ .

Since each parallel I/O operation accesses at most  $BD$  records, any algorithm that must access all  $N$  records requires  $\Omega(N/BD)$  parallel I/Os, and so  $O(N/BD)$  parallel I/Os is the analogue of linear time in sequential computing. Vitter and Shriver showed an upper bound of  $\Theta\left(\min\left(\frac{N}{D}, \frac{N}{BD} \frac{\lg(N/B)}{\lg(M/B)}\right)\right)$  parallel I/Os for general permutations, that is, for arbitrary mappings  $\pi : \{0, 1, \dots, N - 1\} \xrightarrow{1-1} \{0, 1, \dots, N - 1\}$ . The first term comes into play when the block size  $B$  is small, and the second term is the sorting bound  $\Theta\left(\frac{N}{BD} \frac{\lg(N/B)}{\lg(M/B)}\right)$ , which was shown by Vitter and Shriver for randomized sorting and by Nodine and Vitter [9, 10, 11] for deterministic sorting. These bounds are asymptotically tight, for they match the lower bounds proven earlier by Aggarwal and

| Permutation                                      | Characteristic matrix  | Number of passes   |
|--|--|--|
| BMMC<br>(bit-matrix-multiply/<br>complement)     | nonsingular matrix $A$   | $2 \left\lceil \frac{\lg M - r}{\lg(M/B)} \right\rceil + H(N, M, B)$ |
| BPC<br>(bit-permute/<br>complement)              | permutation matrix $A$   | $2 \left\lceil \frac{\rho(A)}{\lg(M/B)} \right\rceil + 1$            |
| MRC<br>(memory-<br>rearrangement/<br>complement) | $\left[ \begin{array}{c c} m & n-m \\ \hline \text{nonsingular} & \text{arbitrary} \\ \hline 0 & \text{nonsingular} \end{array} \right] \begin{array}{l} m \\ n-m \end{array}$ | 1  |

**Table 1:** Classes of permutations, their characteristic matrices, and upper bounds shown in [4] on the number of passes needed to perform them. A pass consists of reading and writing each record exactly once and therefore uses exactly  $2N/BD$  parallel I/Os. For MRC permutations, submatrix dimensions are shown on matrix borders. For BMMC permutations,  $r$  is the rank of the leading  $\lg M \times \lg M$  submatrix of  $A$ , and the function  $H(N, M, B)$  is given by equation (1). For BPC permutations, the function  $\rho(A)$  is defined in equation (3).

Vitter [2] using a model with one disk and  $D$  independent read/write heads, which is at least as powerful as the Vitter-Shriver model.

Specific classes of permutations sometimes require fewer parallel I/Os than general permutations. Vitter and Shriver showed how to transpose an  $R \times S$  matrix ( $N = RS$ ) with only  $\Theta\left(\frac{N}{BD} \left(1 + \frac{\lg \min(B, R, S, N/B)}{\lg(M/B)}\right)\right)$  parallel I/Os. Subsequently, Cormen [4] studied several classes of bit-defined permutations that include matrix transposition as a special case. Table 1 shows some of the classes of permutations examined and the corresponding upper bounds derived in [4].

## BMMC permutations

The most general class considered in [4] is *bit-matrix-multiply/complement*, or *BMMC*, permutations.<sup>1</sup> In a BMMC permutation, we have an  $n \times n$  *characteristic matrix*  $A = (a_{ij})$  whose entries are drawn from  $\{0, 1\}$  and is nonsingular (i.e., invertible) over  $GF(2)$ ,<sup>2</sup> and we have a *complement vector*  $c = (c_0, c_1, \dots, c_{n-1})$  of length  $n$ . Treating a *source address*  $x$  as an  $n$ -bit vector, we perform matrix-vector multiplication over  $GF(2)$  and then form the corresponding  $n$ -bit *target address*  $y$  by complementing some subset of the resulting bits:  $y = Ax \oplus c$ , or

$$\begin{bmatrix} y_0 \\ y_1 \\ y_2 \\ \vdots \\ y_{n-1} \end{bmatrix} = \begin{bmatrix} a_{00} & a_{01} & a_{02} & \cdots & a_{0,n-1} \\ a_{10} & a_{11} & a_{12} & \cdots & a_{1,n-1} \\ a_{20} & a_{21} & a_{22} & \cdots & a_{2,n-1} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ a_{n-1,0} & a_{n-1,1} & a_{n-1,2} & \cdots & a_{n-1,n-1} \end{bmatrix} \begin{bmatrix} x_0 \\ x_1 \\ x_2 \\ \vdots \\ x_{n-1} \end{bmatrix} \oplus \begin{bmatrix} c_0 \\ c_1 \\ c_2 \\ \vdots \\ c_{n-1} \end{bmatrix}.$$

<sup>1</sup>Edelman, Heller, and Johnsson [5] call BMMC permutations *affine transformations* or, if there is no complementing, *linear transformations*.

<sup>2</sup>Matrix multiplication over  $GF(2)$  is like standard matrix multiplication over the reals but with all arithmetic performed modulo 2. Equivalently, multiplication is replaced by logical-and, and addition is replaced by exclusive-or.

We shall generally focus on the matrix-multiplication portion of BMMC permutations rather than on the complement vector. The permutation  $\pi_A$  characterized by a matrix  $A$  is the permutation for which  $\pi_A(x) = A x$  for all source addresses  $x$ .

The following lemma shows the equivalence of multiplying characteristic matrices and composing permutations when the complement vectors are zero. For permutations  $\pi_Y$  and  $\pi_Z$ , the *composition*  $\pi_Z \circ \pi_Y$  is defined by  $(\pi_Z \circ \pi_Y)(x) = \pi_Z(\pi_Y(x))$  for all  $x$  in the domain of  $\pi_Y$ .

**Lemma 1** *Let  $Z$  and  $Y$  be nonsingular  $n \times n$  matrices and let  $\pi_Z$  and  $\pi_Y$  be the permutations characterized by  $Z$  and  $Y$ , respectively. Then the matrix product  $ZY$  characterizes the composition  $\pi_Z \circ \pi_Y$ .*

*Proof:* For any source address  $x$ , we have

$$\begin{aligned} (\pi_Z \circ \pi_Y)(x) &= \pi_Z(\pi_Y(x)) \\ &= \pi_Z(Y x) \\ &= Z(Y x) \\ &= (ZY)x , \end{aligned}$$

and so the matrix product  $ZY$  characterizes the composition  $\pi_Z \circ \pi_Y$ . ■

When we factor a characteristic matrix  $A$  into the product of several nonsingular matrices, each factor characterizes a BMMC permutation. The following corollary describes the order in which we perform these permutations to effect the permutation characterized by  $A$ .

**Corollary 2** *Let the  $n \times n$  characteristic matrix  $A$  be factored as  $A = A^{(k)} A^{(k-1)} A^{(k-2)} \dots A^{(1)}$ , where each factor  $A^{(i)}$  is a nonsingular  $n \times n$  matrix. Then we can perform the BMMC permutation characterized by  $A$  by performing, in order, the BMMC permutations characterized by  $A^{(1)}, A^{(2)}, \dots, A^{(k)}$ . That is, we perform the permutations characterized by the factors of a matrix from right to left.*

*Proof:* The proof is a simple induction, using Lemma 1. ■

The BMMC algorithm in [4] exploits Corollary 2 to factor a characteristic matrix into a product of other characteristic matrices, performing the permutations given by the factors right to left. It uses

$$\frac{2N}{BD} \left( 2 \left\lceil \frac{\lg M - r}{\lg(M/B)} \right\rceil + H(N, M, B) \right)$$

parallel I/Os, where  $r$  is the rank of the leading  $\lg M \times \lg M$  submatrix of the characteristic matrix and

$$H(N, M, B) = \begin{cases} 4 \left\lceil \frac{\lg B}{\lg(M/B)} \right\rceil + 9 & \text{if } M \leq \sqrt{N} , \\ 4 \left\lceil \frac{\lg(N/B)}{\lg(M/B)} \right\rceil + 1 & \text{if } \sqrt{N} < M < \sqrt{NB} , \\ 5 & \text{if } \sqrt{NB} \leq M . \end{cases} \quad (1)$$

One can adapt the lower bound proven in this paper to show that  $\Omega\left(\frac{N}{BD} \frac{\lg M - r}{\lg(M/B)}\right)$  parallel I/Os are necessary (see Section 2.8 of [3] for example), but so far it has been unknown whether the  $\Theta\left(\frac{N}{BD} H(N, M, B)\right)$  term is necessary in all cases. This paper shows that it is not.

## BPC permutations

By restricting the characteristic matrix  $A$  of a BMMC permutation to be a permutation matrix—having exactly one 1 in each row and each column—we obtain the class of *bit-permute/complement*, or *BPC*, permutations.<sup>3</sup> One can think of a BPC permutation as forming each target address by applying a fixed permutation to the source-address bits and then complementing a subset of the resulting bits. The class of BPC permutations includes many common permutations such as matrix transposition, bit-reversal permutations (used in performing FFTs), vector-reversal permutations, hypercube permutations, and matrix reblocking.

Previous work [4] expressed the I/O complexity of BPC permutations in terms of cross-ranks. For any  $n \times n$  permutation matrix  $A$  and for any  $k = 0, 1, \dots, n-1$ , the  $k$ -cross-rank of  $A$  is

$$\rho_k(A) = \text{rank } A_{k..n-1, 0..k-1} = \text{rank } A_{0..k-1, k..n-1} , \quad (2)$$

where, for example,  $A_{k..n-1, 0..k-1}$  denotes the submatrix of  $A$  consisting of the intersection of rows  $k, k+1, \dots, n-1$  and columns  $0, 1, \dots, k-1$ . The *cross-rank* of  $A$  is the maximum of the  $b$ - and  $m$ -cross-ranks:

$$\rho(A) = \max(\rho_b(A), \rho_m(A)) . \quad (3)$$

The BPC algorithm in [4] uses at most

$$\frac{2N}{BD} \left( 2 \left\lceil \frac{\rho(A)}{\lg(M/B)} \right\rceil + 1 \right)$$

parallel I/Os. One can adapt the lower bound we prove in Section 2 for BMMC permutations to show that this BPC algorithm is asymptotically optimal. The BMMC algorithm in Section 5, however, is asymptotically optimal for all BMMC permutations—including those that are BPC—and it reduces the innermost factor of 2 in the above bound to a factor of 1. Not only is the BPC algorithm in [4] improved upon by the results in this paper, but the notion of cross-rank appears to be obviated as well.

## MRC permutations

*Memory-rearrangement/complement*, or *MRC*, permutations are BMMC permutations with the additional restrictions shown in Table 1: both the leading  $m \times m$  and trailing  $(n-m) \times (n-m)$  submatrices of the characteristic matrix are nonsingular, the upper right  $m \times (n-m)$  submatrix can contain any 0-1 values at all, and the lower left  $(n-m) \times m$  submatrix is all 0. Cormen [4] shows that any MRC permutation requires only one pass of  $N/BD$  parallel reads and  $N/BD$  parallel writes. If we partition the  $N$  records into  $N/M$  consecutive sets of  $M$  records each, we call each set a *memoryload*. Each memoryload consists of  $M/BD$  consecutive stripes in which all addresses have the same value in the most significant  $n-m$  bits, as Figure 2 shows. Any MRC permutation can be performed by reading in a memoryload, permuting its records in memory, and writing them out to a (possibly different) memoryload number. The class of MRC permutations includes those characterized by unit upper triangular matrices. As [4] shows, both the standard binary-reflected Gray code and its inverse have characteristic matrices of this form, and so they are MRC permutations.

---

<sup>3</sup>Johnsson and Ho [8] call BPC permutations *dimension permutations*, and Aggarwal, Chandra, and Snir [1] call BPC permutations without complementing *rational permutations*.

## MLD permutations

We define here a new BMMC permutation subclass, which we shall use in our asymptotically optimal BMMC algorithm. We first need to define the *kernel* of any  $p \times q$  matrix  $A$  as the set of  $q$ -vectors that map to 0 when multiplied by  $A$ . That is,

$$\ker A = \{x : Ax = 0\}.$$

A *memoryload-dispersal*, or *MLD*, permutation has a characteristic matrix that is nonsingular and of the following form:

$$\left[ \begin{array}{c|c} \begin{array}{c} m \\ \hline \text{arbitrary} \\ \hline \lambda \\ \hline \mu \end{array} & \begin{array}{c} n-m \\ \text{arbitrary} \end{array} \end{array} \right] \begin{array}{l} b \\ m-b \\ n-m \end{array},$$

subject to the *kernel condition*

$$\ker \lambda \subseteq \ker \mu \tag{4}$$

or, equivalently,  $\lambda x = 0$  implies  $\mu x = 0$ .

As we shall see in Section 3, we can perform any MLD permutation in one pass by reading in each source memoryload, permuting its records in memory, and writing these records out to  $M/BD$  blocks on each disk. Although the blocks read from each memoryload must come from  $M/BD$  consecutive stripes, the blocks written may go to any locations at all, as long as  $M/BD$  blocks are written to each disk. That is, MLD permutations use striped reads and independent writes.

## Outline

The remainder of this paper is organized as follows. Section 2 states and proves the lower bound for BMMC permutations. Section 3 shows how to perform any MLD permutation in one pass and gives some additional properties of MLD permutations. Section 4 previews several of the matrix forms used in Section 5, which presents an algorithm for BMMC permutations whose I/O complexity asymptotically matches the lower bound. Section 6 shows how to detect at run time whether a vector of target addresses describes a BMMC permutation, thus enabling us to determine whether the BMMC algorithm is applicable; this section also presents an easy method for determining whether a nonsingular matrix satisfies the kernel condition (4) and therefore characterizes an MLD permutation. Finally, Section 7 contains some concluding remarks.

The algorithms for MLD and BMMC permutations in Sections 3 and 5 are on-line in the sense that they take little computation time and space. (They do, however, require permutations to be performed in memory, and various architectures may differ in how efficiently they do so.) The data structures are vectors of length  $\lg N$  or matrices of size at most  $\lg N \times \lg N$ . Even serial algorithms for the harder computations take time polynomial in  $\lg N$ , in fact  $O(\lg^3 N)$ .

We shall use several notational conventions in this paper, as in equation (2). Matrix row and column numbers are indexed from 0 starting from the upper left. Vectors are indexed from 0, too. We index rows and columns by sets to indicate submatrices, using “.” notation to indicate sets of contiguous numbers. When a matrix is indexed by just one set rather than two, the set indexes column numbers; the submatrix consists of entire columns. When a submatrix index is a singleton set, we shall often omit the enclosing braces. We denote an identity matrix by  $I$  and a matrix whose entries are all 0s by 0; the dimensions of such matrices will be clear from their contexts. All matrix



and vector elements are drawn from  $\{0, 1\}$ , and all matrix and vector arithmetic is over  $GF(2)$ . When convenient, we interpret bit vectors as the integers they represent in binary. Vectors are treated as 1-column matrices in context.

Some readers familiar with linear algebra may notice that a few of the lemmas in this paper are special cases of standard linear-algebra properties restricted to  $GF(2)$ . We include the proofs here for completeness.

## 2 A universal lower bound for BMMC permutations

In this section, we state and prove the lower bound for BMMC permutations. After stating the lower bound, we briefly discuss its significance before presenting the full proof. The lower bound is given by the following theorem.

**Theorem 3** *Any algorithm that performs a non-identity BMMC permutation with characteristic matrix  $A$  requires*

$$\Omega\left(\frac{N}{BD}\left(1 + \frac{\text{rank } \gamma}{\lg(M/B)}\right)\right)$$

*parallel I/Os, where  $\gamma$  is the submatrix  $A_{b..n-1, 0..b-1}$  of size  $\lg(N/B) \times \lg B$ .* ■

This lower bound is *universal* in the sense that it applies to all inputs other than the identity permutation, which of course requires no data movement at all. In contrast, lower bounds such as the standard  $\Omega(N \lg N)$  lower bound for sorting  $N$  items on a sequential machine are *existential*: they apply to worst-case inputs, but for some inputs an algorithm may be able to do better.

Section 5 presents an algorithm that achieves the bound given by Theorem 3, and so this algorithm is asymptotically optimal. The BPC algorithm of [4] is asymptotically optimal as well; see [3] for details.

### Technique

To prove Theorem 3, we rely heavily on the technique used by Aggarwal and Vitter [2] for a lower bound on I/Os in matrix transposition; their proof is based in turn on a method by Floyd [6]. We prove the lower bound for the case in which  $D = 1$ ; the general case follows by dividing by  $D$ . We consider only I/Os that are *simple*. An input is simple if each record read is removed from the disk and moved into an empty location in memory. An output is simple if the records are removed from the memory and written to empty locations on the disk. When all I/Os are simple, exactly one copy of each record exists at any time during the execution of an algorithm. The following lemma, proven by Aggarwal and Vitter, allows us to consider only simple I/Os when proving lower bounds.

**Lemma 4** *For each computation that implements a permutation of records, there is a corresponding computation strategy involving only simple I/Os such that the total number of I/Os is no greater.* ■

The basic scheme of the proof of Theorem 3 uses a potential-function argument. *Time  $q$*  is the time interval starting when the  $q$ th I/O completes and ending just before the  $(q + 1)$ st I/O starts. We define a potential function  $\Phi$  so that  $\Phi(q)$  is the *potential* at time  $q$ . This potential measures how close the current record ordering is to the desired permutation order. Higher potentials indicate that the current ordering is closer to the desired permutation. We compute the initial and final potentials and bound the amount that the potential can increase in each I/O operation. The lower bound then follows.

To be more precise, we start with some definitions. For  $i = 0, 1, \dots, N/B - 1$ , we define the  $i$ th *target group* to be the set of records that belong in block  $i$  according to the given BMBC permutation. We denote by  $g_{\text{block}}(i, k, q)$  the number of records in the  $i$ th target group that are in block  $k$  on disk at time  $q$ , and  $g_{\text{mem}}(i, q)$  denotes the number of records in the  $i$ th target group that are in memory at time  $q$ . We define the continuous function

$$f(x) = \begin{cases} x \lg x & \text{if } x > 0, \\ 0 & \text{if } x = 0, \end{cases}$$

and we define *togetherness functions*

$$G_{\text{block}}(k, q) = \sum_{i=0}^{N/B-1} f(g_{\text{block}}(i, k, q))$$

for each block  $k$  at time  $q$  and

$$G_{\text{mem}}(q) = \sum_{i=0}^{N/B-1} f(g_{\text{mem}}(i, q))$$

for memory at time  $q$ . Finally, we define the *potential* at time  $q$ , denoted  $\Phi(q)$ , as the sum of the togetherness functions:

$$\Phi(q) = G_{\text{mem}}(q) + \sum_{k=0}^{N/B-1} G_{\text{block}}(k, q).$$

Aggarwal and Vitter embed the following lemmas in their lower-bound argument. The first lemma is based on the observation that the number of parallel I/Os is at least the total increase in potential over all parallel I/Os divided by the maximum increase in potential (denoted  $\Delta\Phi_{\text{max}}$ ) in any single parallel I/O.

**Lemma 5** *Let  $D = 1$ , and consider any algorithm that uses  $t$  parallel I/Os to perform a permutation. Then  $t = \Omega\left(\frac{\Phi(t) - \Phi(0)}{\Delta\Phi_{\text{max}}}\right)$ .* ■

**Lemma 6** *Let  $D = 1$ , and consider any permutation that can be performed with  $t$  parallel I/Os. Then  $\Phi(t) = N \lg B$  and  $\Delta\Phi_{\text{max}} = O(B \lg(M/B))$ . Therefore, any algorithm that performs a permutation uses  $\Omega\left(\frac{N \lg B - \Phi(0)}{B \lg(M/B)}\right)$  parallel I/Os.* ■

Observe that these lemmas imply lower bounds that are universal. No matter what the input, the initial potential is  $\Phi(0)$ , the final potential is  $\Phi(t)$ , the increase in potential per parallel I/O is at most  $\Delta\Phi_{\text{max}}$ , and so  $\Omega\left(\frac{\Phi(t) - \Phi(0)}{\Delta\Phi_{\text{max}}}\right)$  parallel I/Os are required.

## Ranges and preimages

To prove lower bounds for BMBC and BPC permutations, we shall examine ranges of matrices and preimages of vectors under matrix multiplication.

For a  $p \times q$  matrix  $A$  with 0-1 entries, we define the *range* of  $A$  by

$$\mathcal{R}(A) = \{y : y = Ax \text{ for some } x \in \{0, 1, \dots, 2^q - 1\}\},$$

that is,  $\mathcal{R}(A)$  is the set of  $p$ -vectors that can be produced by multiplying all  $q$ -vectors with 0-1 entries (interpreted as integers in  $\{0, 1, \dots, 2^q - 1\}$ ) by  $A$  over  $GF(2)$ . We also adopt the notation

$$\mathcal{R}(A) \oplus c = \{z : z = y \oplus c \text{ for some } y \in \mathcal{R}(A)\} ,$$

that is,  $\mathcal{R}(A) \oplus c$  is the exclusive-or of the range of  $A$  and a fixed vector  $c$ .

**Lemma 7** *Let  $A$  be a  $p \times q$  matrix whose entries are drawn from  $\{0, 1\}$ , let  $c$  be any  $p$ -vector whose entries are drawn from  $\{0, 1\}$ , and let  $r = \text{rank } A$ . Then  $|\mathcal{R}(A) \oplus c| = 2^r$ .*

*Proof:* Let  $S$  index a maximal set of linearly independent columns of  $A$ , so that  $S \subseteq \{0, 1, \dots, q-1\}$ ,  $|S| = r$ , the columns of the submatrix  $A_S$  are linearly independent, and for any column number  $j \notin S$ , the column  $A_j$  is linearly dependent on the columns of  $A_S$ . We claim that  $\mathcal{R}(A) = \mathcal{R}(A_S)$ . Clearly,  $\mathcal{R}(A_S) \subseteq \mathcal{R}(A)$ , since  $\mathcal{R}(A)$  includes the sum (over  $GF(2)$ ) of each subset of columns of  $A$ . To see that  $\mathcal{R}(A) \subseteq \mathcal{R}(A_S)$ , consider any  $q$ -vector  $y \in \mathcal{R}(A)$ . There is some set  $T$  of column indices such that  $y = \bigoplus_{j \in T} A_j$ . For each column index  $j \in T - S$ , let  $S_j \subseteq S$  index the columns of  $A_S$  that  $A_j$  depends on:  $A_j = \bigoplus_{k \in S_j} A_k$ . Then we have

$$\begin{aligned} y &= \bigoplus_{j \in T} A_j \\ &= \left( \bigoplus_{j \in T \cap S} A_j \right) \oplus \left( \bigoplus_{j \in T - S} A_j \right) \\ &= \left( \bigoplus_{j \in T \cap S} A_j \right) \oplus \left( \bigoplus_{j \in T - S} \left( \bigoplus_{k \in S_j} A_k \right) \right) , \end{aligned}$$

and so  $y$  is a linear combination of columns of  $A_S$ . Thus,  $y \in \mathcal{R}(A_S)$ , which in turn proves that  $\mathcal{R}(A) \subseteq \mathcal{R}(A_S)$  and consequently  $\mathcal{R}(A) = \mathcal{R}(A_S)$ .

We have  $|\mathcal{R}(A_S)| = 2^{|S|} = 2^r$ , since each column index in  $S$  may or may not be included in a sum of the columns. Thus,  $|\mathcal{R}(A)| = 2^r$ . Exclusive-oring the result of the matrix multiplication by a constant  $p$ -vector does not change the cardinality of the range. Therefore,  $|\mathcal{R}(A) \oplus c| = |\mathcal{R}(A)| = 2^r$ .  $\blacksquare$

For a  $p \times q$  matrix  $A$  and a  $p$ -vector  $y \in \mathcal{R}(A)$ , we define the *preimage* of  $y$  under  $A$  by

$$\text{Pre}(A, y) = \{x : Ax = y\} .$$

That is,  $\text{Pre}(A, y)$  is the set of  $q$ -vectors  $x$  that map to  $y$  when multiplied by  $A$ .

**Lemma 8** *Let  $A$  be a  $p \times q$  matrix whose entries are drawn from  $\{0, 1\}$ , let  $y$  be any  $p$ -vector in  $\mathcal{R}(A)$ , and let  $r = \text{rank } A$ . Then  $|\text{Pre}(A, y)| = 2^{q-r}$ .*

*Proof:* Let  $S$  index a maximal set of linearly independent columns of  $A$ , so that  $S \subseteq \{0, 1, \dots, q-1\}$ ,  $|S| = r$ , the columns of the submatrix  $A_S$  are linearly independent, and for any column number  $j \notin S$ , the column  $A_j$  is linearly dependent on the columns of  $A_S$ . Let  $S' = \{0, 1, \dots, q-1\} - S$ .

We claim that for any value  $i \in \{0, 1, \dots, 2^{q-r} - 1\}$ , there is a unique  $q$ -vector  $x^{(i)}$  for which  $x_{S'}^{(i)}$  is the binary representation of  $i$  and  $y = Ax^{(i)}$ . Why? We have  $y = A_S x_S^{(i)} \oplus A_{S'} x_{S'}^{(i)}$  or, equivalently,

$$y \oplus A_{S'} x_{S'}^{(i)} = A_S x_S^{(i)} . \tag{5}$$

The columns of  $A_S$  span  $\mathcal{R}(A)$ , which implies that for all  $z \in \mathcal{R}(A)$ , there is a unique  $r$ -vector  $w$  such that  $z = A_S w$ . Letting  $z = y \oplus A_{S'} x_{S'}^{(i)}$ , we see that there is a unique  $r$ -vector  $x_S^{(i)}$  that satisfies equation (5), which proves the claim.

Thus, we have shown that  $|\text{Pre}(A, y)| \geq 2^{q-r}$ . If we had  $|\text{Pre}(A, y)| > 2^{q-r}$ , then because  $y$  is arbitrarily chosen from  $\mathcal{R}(A)$ , we would have that  $\sum_{y' \in \mathcal{R}(A)} |\text{Pre}(A, y')| > 2^q$ . But this inequality contradicts there being only  $2^q$  possible preimage vectors. We conclude that  $|\text{Pre}(A, y)| = 2^{q-r}$ . ■

We can now show a trivial lower bound for all non-identity BMMC permutations.

**Lemma 9** *If  $D = 1$ , any algorithm that performs a BMMC permutation requires  $\Omega(N/B)$  parallel I/Os whenever the permutation is not the identity permutation.*

*Proof:* Consider a BMMC permutation with characteristic matrix  $A$  and complement vector  $c$ . It is the identity permutation if and only if  $A = I$  and  $c = 0$ , so we shall assume that either  $A \neq I$  or  $c \neq 0$ .

A *fixed point* of the BMMC permutation is a source address  $x$  for which

$$A x \oplus c = x . \quad (6)$$

If a record's source address is not a fixed point, its source block must be read and its target block must be written. We shall show that for any non-identity BMMC permutation, at least  $N/2$  addresses are not fixed points. Even if these records are clustered into as few source blocks as possible, then at least half the source blocks, or  $N/2B$ , must be read. The lemma then follows.

To show that at least  $N/2$  addresses are not fixed points, we shall show that at most  $N/2$  addresses are. Rewriting equation (6) as  $(A \oplus I)x = c$ , we see that we wish to bound the size of  $\text{Pre}(A \oplus I, c)$ . If  $c \notin \mathcal{R}(A \oplus I)$ , then this size is 0. Otherwise, by Lemma 8, this size is  $2^{n-\text{rank}(A \oplus I)}$ . If  $A \neq I$ , then  $\text{rank}(A \oplus I) \geq 1$ , which implies that  $|\text{Pre}(A \oplus I, c)| \leq 2^{n-1} = N/2$ . If  $A = I$ , then  $A \oplus I$  is the 0 matrix, and the only vector in its range is 0. But  $A = I$  and  $c = 0$  yields the identity permutation, which we specifically disallow. ■

### Proof of Theorem 3

To prove Theorem 3, we prove the lower bound for the case in which  $D = 1$ ; the general case follows by dividing by  $D$ . We work with characteristic matrix  $A$  and complement vector  $c$ . We assume that all I/Os are simple and transfer exactly  $B$  records, some possibly empty. Since all records start on disk and I/Os are simple, memory is initially empty.

We need to compute the initial potential in order to apply Lemma 6. The initial potential depends on the number of records that start in the same source block and are in the same target group. A record with source address  $x = (x_0, x_1, \dots, x_{n-1})$  is in source block  $k$  if and only if

$$k = x_{b..n-1} , \quad (7)$$

interpreting  $k$  as an  $(n - b)$ -bit binary number. This record maps to target block  $i$  if and only if

$$\begin{aligned} i &= A_{b..n-1, 0..n-1} x_{0..n-1} \oplus c_{b..n-1} \\ &= A_{b..n-1, 0..b-1} x_{0..b-1} \oplus A_{b..n-1, b..n-1} x_{b..n-1} \oplus c_{b..n-1} , \end{aligned} \quad (8)$$

also interpreting  $i$  as an  $(n - b)$ -bit binary number. The following lemma gives the exact number of records that start in each source block and are in the same target group.

**Lemma 10** *Let  $r = \text{rank } A_{b..n-1,0..b-1}$ , and consider any source block  $k$ . There are exactly  $2^r$  distinct target blocks that some record in source block  $k$  maps to, and for each such target block, exactly  $B/2^r$  records in source block  $k$  map to it.*

*Proof:* For a given source block  $k$ , all source addresses fulfill condition (7), and so they map to target block numbers given by condition (8) but with  $x_{b..n-1}$  fixed at  $k$ . The range of target block numbers is thus  $\mathcal{R}(A_{b..n-1,0..b-1}) \oplus (A_{b..n-1,b..n-1} k \oplus c_{b..n-1})$  which, by Lemma 7, has cardinality  $2^r$ .

Now we determine the set of source addresses in source block  $k$  that map to a particular target block  $i$  in  $\mathcal{R}(A_{b..n-1,0..b-1}) \oplus (A_{b..n-1,b..n-1} k \oplus c_{b..n-1})$ . Again fixing  $x_{b..n-1} = k$  in condition (8) and exclusive-oring both sides by  $A_{b..n-1,b..n-1} k \oplus c_{b..n-1}$ , we see that this set is precisely  $\text{Pre}(A_{b..n-1,0..b-1}, i \oplus A_{b..n-1,b..n-1} k \oplus c_{b..n-1})$ . By Lemma 8, this set has cardinality exactly  $2^{b-r}$ , which equals  $B/2^r$ . ■

We can interpret Lemma 10 as follows. Let  $r = \text{rank } A_{b..n-1,0..b-1}$ , and consider a particular source block  $k$ . Then there are exactly  $2^r$  target blocks  $i$  for which  $g_{\text{block}}(i, k, 0)$  is nonzero, and for each such nonzero target block, we have  $g_{\text{block}}(i, k, 0) = B/2^r$ .

Now we can compute  $\Phi(0)$ . Since memory is initially empty,  $g_{\text{mem}}(i, 0) = 0$  for all blocks  $i$ , which implies that  $G_{\text{mem}}(0) = 0$ . We have

$$\begin{aligned}
\Phi(0) &= G_{\text{mem}}(0) + \sum_{k=0}^{N/B-1} G_{\text{block}}(k, 0) \\
&= 0 + \sum_{k=0}^{N/B-1} \sum_{i=0}^{N/B-1} f(g_{\text{block}}(i, k, 0)) \\
&= \sum_{k=0}^{N/B-1} 2^r \frac{B}{2^r} \lg \frac{B}{2^r} \quad (\text{by Lemma 10}) \\
&= \frac{N}{B} B \lg \frac{B}{2^r} \\
&= N(\lg B - r) .
\end{aligned} \tag{9}$$

Combining Lemmas 6 and 9 with equation (9), we get a lower bound of

$$\Omega \left( \frac{N}{B} + \frac{N \lg B - N(\lg B - r)}{B \lg(M/B)} \right) = \Omega \left( \frac{N}{B} \left( 1 + \frac{\text{rank } A_{b..n-1,0..b-1}}{\lg(M/B)} \right) \right)$$

parallel I/Os. Dividing through by  $D$  yields a lower bound of

$$\Omega \left( \frac{N}{BD} \left( 1 + \frac{\text{rank } A_{b..n-1,0..b-1}}{\lg(M/B)} \right) \right) ,$$

which completes the proof of Theorem 3.

### 3 MLD permutations

In this section, we describe how to perform any MLD permutation in only one pass. This section also discusses additional properties of MLD and MRC permutations. Section 6 shows how to determine whether a given matrix characterizes an MLD permutation.

## How the kernel condition implies a one-pass permutation

We shall show in three steps that the kernel condition implies that, for a given source memoryload, the source records are permuted into full target blocks spread evenly across the disks. To do so, we first need to define the notion of relative block number, as shown in Figure 2. For a given  $n$ -bit record address  $x_{0..n-1}$ , the *relative block number* of  $x$  is the  $m - b$  bits  $x_{b..m-1}$ . The relative block number ranges from 0 to  $M/B - 1$  and determines the number of a block within its memoryload. We shall prove that *for a given source memoryload*, the following properties hold:

1. Its records map to all  $M/B$  relative block numbers, and each relative block number that has any record mapping to it has exactly  $B$  records mapping to it.
2. Records that map to the same relative block number map to the same target memoryload number as well.

The first two properties imply that the records of each source memoryload map to exactly  $M/B$  target blocks and that each such target block is full.

3. These  $M/B$  target blocks are distributed evenly among the disks, with  $M/BD$  mapping to each disk.

Given these properties, we can perform an MLD permutation in one pass. Like the other one-pass permutations described in [4], we allow the permutation to map records from one set of  $N/BD$  stripes (the “source portion” of the parallel disk system) to a different set of  $N/BD$  stripes (the “target portion”). One can think of addresses as relative to the beginning of the appropriate portion. In this way, we need not be concerned with overwriting source records before we get a chance to read them. Note that when we chain passes together, as in the BMMC algorithm of Section 5 and the BPC algorithm of [4], we can avoid allocating a new target portion in each pass by reversing the roles of the source and target portions between passes.

We perform an MLD permutation by processing source memoryload numbers from 0 to  $N/M - 1$ . For each source memoryload, we first read into memory its  $M/BD$  consecutive stripes from the source portion. We then permute its records in memory, clustering them into  $M/B$  full target blocks that are distributed evenly among the disks. We then write out these target blocks using  $M/BD$  independent writes to the target portion. After processing all  $N/M$  source memoryloads, we have read each record from the source portion and written it to where it belongs in the target portion. Thus, we have performed the MLD permutation in one pass.

Before proving the three properties above, we define the *row space* of a matrix  $A$ , written  $\text{row } A$ , as the span of the rows of  $A$ . We start by proving the following lemma about the relationship between kernels and row spaces, which we shall subsequently use to prove properties resulting from the kernel condition.

**Lemma 11** *Let  $K$  and  $L$  be  $q$ -column matrices for which  $\ker K \subseteq \ker L$ . Then  $\text{row } L \subseteq \text{row } K$ .*

*Proof:* It is a well-known fact from linear algebra (see Strang [12, p. 138] for example) that the row space is the orthogonal complement of the kernel. That is, for any  $q$ -vectors  $u \in \ker L$  and  $v \in \text{row } L$ , the inner product  $u \cdot v$  is 0. We write  $u \perp v$  to denote this property, and we extend this notation to sets:  $v \perp \ker L$  means that for all  $u \in \ker L$ ,  $v \perp u$ .

Consider any  $q$ -vector  $v \in \text{row } L$ . Then  $v \perp \ker L$  and, because  $\ker K \subseteq \ker L$ , we have that  $v \perp \ker K$  as well. Again using the orthogonality of the row space and the kernel, we have that  $v \in \text{row } K$ . Thus,  $\text{row } L \subseteq \text{row } K$ . ■

Although the kernel condition does not explicitly mention the leading  $m \times m$  submatrix, the following lemma gives an important property of this submatrix.

**Lemma 12** *Let  $A$  characterize an MLD permutation. Then the leading  $m \times m$  submatrix of  $A$  is nonsingular.*

*Proof:* Since  $\ker \lambda \subseteq \ker \mu$ , Lemma 11 implies that  $\text{row } \mu \subseteq \text{row } \lambda$ . Thus, every row of  $\mu$  is linearly dependent on some rows of  $\lambda$ . Because the rank of the leftmost  $m$  columns of  $A$  (i.e., the submatrix  $A_{0..n-1,0..m-1}$ ) is  $m$  and each row of  $\mu$  is a linear combination of the rows of the leading  $m \times m$  submatrix of  $A$ , all rows of this leading  $m \times m$  submatrix must be linearly independent. Therefore, the leading  $m \times m$  submatrix of  $A$  is nonsingular. ■

We now prove property 1.

**Lemma 13** *The records of each source memoryload in an MLD permutation map to exactly  $M/B$  relative block numbers. Moreover, for a given source memoryload, each relative block number that has any record mapping to it has exactly  $B$  records mapping to it.*

*Proof:* Let  $A$  characterize an MLD permutation with complement vector  $c$ . By Lemma 12,  $\text{rank } A_{b..m-1,0..m-1} = m - b$ . The target relative block number  $y_{b..m-1}$  corresponding to a source address  $x$  is given by the equation

$$y_{b..m-1} = A_{b..m-1,0..m-1} x_{0..m-1} \oplus A_{b..m-1,m..n-1} x_{m..n-1} \oplus c_{b..m-1}. \quad (10)$$

The value of  $x_{m..n-1}$  is fixed for a given source memoryload, and so the  $(m - b)$ -vector  $A_{b..m-1,m..n-1} x_{m..n-1} \oplus c_{b..m-1}$  has the same value for all records. By Lemma 7,  $y_{b..m-1}$  takes on  $2^{\text{rank } A_{b..m-1,0..m-1}} = 2^{m-b} = M/B$  different values for the  $M$  different values of  $x_{0..m-1}$ . That is, the records of each source memoryload map to exactly  $M/B$  different relative block numbers.

Now consider some relative block number  $y_{b..m-1}$  that some source address in a memoryload maps to. Using equation (10), the number of source addresses  $x_{0..m-1}$  within that memoryload that map to  $y_{b..m-1}$  is equal to  $|\text{Pre}(A_{b..m-1,0..m-1}, y_{b..m-1} \oplus A_{b..m-1,m..n-1} x_{m..n-1} \oplus c_{b..m-1})|$ . By Lemma 8, this number is equal to  $2^{m - \text{rank } A_{b..m-1,0..m-1}} = 2^{m-(m-b)} = B$ . ■

Property 2 follows directly from the kernel condition. Although we use kernel notation for its simplicity of expression, the following lemma shows that the kernel condition is equivalent to requiring that, for a given source memoryload, every record destined for a particular relative block number must also be destined for the same target memoryload.

**Lemma 14** *Let  $K$  and  $L$  be matrices with  $q$  columns. Then  $\ker K \subseteq \ker L$  if and only if for all  $q$ -vectors  $x$  and  $y$ ,  $Kx = Ky$  implies  $Lx = Ly$ .*

*Proof:* Suppose that  $\ker K \subseteq \ker L$  and  $Kx = Ky$ . Then  $K(x \oplus y) = 0$ , which implies that  $L(x \oplus y) = 0$ , which in turn implies  $Lx = Ly$ .

Conversely, suppose that  $Kx = Ky$  implies  $Lx = Ly$  for all  $q$ -vectors  $x$  and  $y$ , and consider any  $q$ -vector  $z \in \ker K$ . We have  $Kz = 0 = K \cdot 0$ , which implies  $Lz = L \cdot 0 = 0$ . Thus,  $z \in \ker L$ . ■

For an MLD permutation, since  $\ker \lambda \subseteq \ker \mu$ , we apply Lemma 14 with  $K = \lambda$  and  $L = \mu$ . Thus, any two source records  $x$  and  $y$  from the same source memoryload that are mapped to relative block number  $\lambda x_{0..m-1}$  are also mapped to the same target memoryload  $\mu x_{0..m-1}$ .

Property 3 follows from property 1. Each source memoryload maps to relative block numbers  $0, 1, \dots, M/B - 1$ . As Figure 2 shows, the number of the disk that a block resides on is encoded in the least significant  $d$  bits of its relative block number. The  $M/B$  relative block numbers, therefore, are evenly distributed among the  $D$  disks, with  $M/BD$  residing on each disk.

Thus, we have proven the following.

**Theorem 15** *Any MLD permutation can be performed in one pass.* ■

## Additional properties of MLD and MRC permutations

We conclude this section with a study of some additional properties of MLD permutations. The first property bounds the rank of the submatrix  $\mu$  as another consequence of the kernel condition.

**Lemma 16** *In the characteristic matrix for an MLD permutation, the submatrix  $\mu$  has rank at most  $m - b$ .*

*Proof:* By Lemma 11 and the kernel condition,  $\text{row } \mu \subseteq \text{row } \lambda$ , which in turn implies that  $\dim(\text{row } \mu) \leq \dim(\text{row } \lambda)$ , where the dimension of a vector space is the size of any basis for it. Thus,  $\text{rank } \mu \leq \text{rank } \lambda = m - b$ . ■

Thus, if the lower left  $(n - m) \times m$  submatrix of a characteristic matrix has rank more than  $m - b$ , the matrix cannot characterize an MLD permutation.

**Theorem 17** *Let the matrix  $Y$  characterize an MLD permutation, and let the matrix  $X$  characterize an MRC permutation. Then the matrix product  $YX$  characterizes an MLD permutation.*

*Proof:* Write the nonsingular matrix  $Y$  as

$$Y = \left[ \begin{array}{c|c} m & n-m \\ \hline \alpha & \beta \\ \hline \gamma & \delta \end{array} \right] \begin{array}{l} m \\ n-m \end{array},$$

where

$$\ker \alpha_{b..m-1,0..m-1} \subseteq \ker \gamma. \quad (11)$$

Write the nonsingular matrix  $X$  as

$$X = \left[ \begin{array}{c|c} m & n-m \\ \hline \phi & \sigma \\ \hline 0 & \nu \end{array} \right] \begin{array}{l} m \\ n-m \end{array},$$

where  $\phi$  and  $\nu$  are nonsingular. We now show that the product

$$YX = \left[ \begin{array}{c|c} m & n-m \\ \hline \alpha\phi & \alpha\sigma \oplus \beta\nu \\ \hline \gamma\phi & \gamma\sigma \oplus \delta\nu \end{array} \right] \begin{array}{l} m \\ n-m \end{array}.$$

characterizes an MLD permutation.

First, the product  $YX$  is nonsingular because  $Y$  and  $X$  are each nonsingular.

We must also prove that the kernel condition (4) holds for the product, i.e., that  $\ker(\alpha\phi)_{b..m-1,0..m-1} \subseteq \ker(\gamma\phi)$ . For an  $m \times m$  matrix  $\tau$ , note that  $\tau_{b..m-1,0..m-1} = I_{b..m-1,0..m-1} \tau$ , where  $I$  is the usual  $m \times m$  identity matrix. We have that  $x \in \ker(\alpha\phi)_{b..m-1,0..m-1}$  implies  $(I_{b..m-1,0..m-1} \alpha\phi)x = 0$  (taking  $\alpha\phi$  as  $\tau$ ), which in turn implies that  $\phi x \in \ker(I_{b..m-1,0..m-1} \alpha) = \ker \alpha_{b..m-1,0..m-1} \subseteq \ker \gamma$ , by property (11). Thus,  $\gamma\phi x = 0$ , and so  $x \in \ker(\gamma\phi)$ . We conclude that  $\ker(\alpha\phi)_{b..m-1,0..m-1} \subseteq \ker(\gamma\phi)$ , which completes the proof. ■

Theorem 17 shows that the composition of an MLD permutation with an MRC permutation is an MLD permutation. Since we have seen how to perform MLD permutations, we can gain an intuitive understanding of why Theorem 17 holds. An MRC permutation permutes memoryload



numbers, with records that start together within a source memoryload remaining together in a target memoryload. We perform an MLD permutation by reading in entire memoryloads. Thus, to perform the composition as an MLD permutation, we only have to remap the source memoryload numbers and adjust the in-memory permutations accordingly. Furthermore, as the following lemma shows, the composition of two MRC permutations is merely the composition of their memoryload mappings with the in-memory permutations adjusted accordingly.

**Theorem 18** *The class of MRC permutations is closed under composition and inverse. That is, if a matrix  $A$  characterizes an MRC permutation, then so does the matrix  $A^{-1}$ , and if  $A^{(1)}$  and  $A^{(2)}$  characterize MRC permutations, then so does the product  $A^{(1)} A^{(2)}$ .*

*Proof:* We first show that MRC permutations are closed under inverse. Let the matrix

$$A = \left[ \begin{array}{c|c} m & n-m \\ \hline \alpha & \beta \\ \hline 0 & \delta \end{array} \right] \begin{array}{l} m \\ n-m \end{array}$$

characterize an MRC permutation, so that the leading submatrix  $\alpha$  and trailing submatrix  $\delta$  are nonsingular. The inverse of this matrix is

$$A^{-1} = \left[ \begin{array}{c|c} m & n-m \\ \hline \alpha^{-1} & \alpha^{-1}\beta\delta^{-1} \\ \hline 0 & \delta^{-1} \end{array} \right] \begin{array}{l} m \\ n-m \end{array},$$

where the leading  $m \times m$  submatrix  $\alpha^{-1}$  and trailing  $(n-m) \times (n-m)$  submatrix  $\delta^{-1}$  are nonsingular. Thus, the matrix  $A^{-1}$  characterizes an MRC permutation.

We now show that MRC permutations are closed under composition. Consider MRC characteristic matrices

$$A^{(1)} = \left[ \begin{array}{c|c} m & n-m \\ \hline \alpha^{(1)} & \beta^{(1)} \\ \hline 0 & \delta^{(1)} \end{array} \right] \begin{array}{l} m \\ n-m \end{array},$$

$$A^{(2)} = \left[ \begin{array}{c|c} m & n-m \\ \hline \alpha^{(2)} & \beta^{(2)} \\ \hline 0 & \delta^{(2)} \end{array} \right] \begin{array}{l} m \\ n-m \end{array},$$

where the submatrices  $\alpha^{(1)}$ ,  $\alpha^{(2)}$ ,  $\delta^{(1)}$ , and  $\delta^{(2)}$  are nonsingular. Then their product is

$$A^{(1)} A^{(2)} = \left[ \begin{array}{c|c} m & n-m \\ \hline \alpha^{(1)} \alpha^{(2)} & \alpha^{(1)} \beta^{(2)} \oplus \beta^{(1)} \delta^{(2)} \\ \hline 0 & \delta^{(1)} \delta^{(2)} \end{array} \right] \begin{array}{l} m \\ n-m \end{array}.$$

Because  $\alpha^{(1)}$  and  $\alpha^{(2)}$  are nonsingular, so is their product  $\alpha^{(1)} \alpha^{(2)}$ . Similarly, the product  $\delta^{(1)} \delta^{(2)}$  is nonsingular. The product  $A^{(1)} A^{(2)}$ , therefore, characterizes an MRC permutation. ■

On the other hand, the composition of two MLD permutations is not necessarily an MLD permutation. We can see this fact in two ways. First, since we perform an MLD permutation by reading in entire memoryloads but writing blocks independently, it may not be possible to remap the source memoryload numbers. Second, consider the product of two matrices, each of which

characterizes an MLD permutation. Although the rank of the lower left  $(n - m) \times m$  submatrix of each factor is at most  $m - b$ , it may be the case that the rank of the lower left  $(n - m) \times m$  submatrix of the product exceeds  $m - b$ . If so, then by Lemma 16, the product cannot characterize an MLD permutation.

Moreover, the composition of an MRC permutation with an MLD permutation (that is, reversing the order of the factors in Theorem 17) is not necessarily an MLD permutation. A simple example is the product

$$\begin{array}{ccc} \begin{array}{c|c|c} b & m-b & n-m \\ \hline 0 & I & 0 \\ \hline I & 0 & 0 \\ \hline 0 & 0 & I \end{array} & \begin{array}{c|c|c} b & m-b & n-m \\ \hline I & 0 & 0 \\ \hline 0 & I & 0 \\ \hline 0 & I & I \end{array} & = \begin{array}{c|c|c} b & m-b & n-m \\ \hline 0 & I & 0 \\ \hline I & 0 & 0 \\ \hline 0 & I & I \end{array} \begin{array}{l} b \\ m-b \\ n-m \end{array}, \\ \text{MRC} & \text{MLD} & \text{not MLD} \end{array}$$

with  $b = m - b = n - m$ . This product is not MLD since an  $m$ -vector with 0s in the first  $b$  positions and 1s in the last  $m - b$  positions is a vector in  $\ker \lambda$ , but it is not a vector in  $\ker \mu$ .

Finally, we note that any MRC permutation is an MLD permutation. Observe that the lower left  $(n - m) \times m$  submatrix of an MRC permutation must be 0, which implies that its kernel is the set of all  $m$ -vectors. No matter what  $\ker \lambda$  is, it is a subset of this set.

## 4 Matrix-column operations

In this section, we classify forms of matrices which, as factors, have the effect of adding columns of other matrices to yield a product. We shall use matrices of this form in Section 5 to transform the characteristic matrix for any BMCM permutation into an MRC permutation. This section shows the structure and useful properties of specific characteristic matrix forms we shall use.

### Column additions

We define a *column-addition* matrix as a matrix  $Q$  such that the product  $A' = AQ$  is a modified form of  $A$  in which specified columns of  $A$  have been added into others. Denoting the  $k$ th column of  $A$  by  $A_k$ , we define the matrix  $Q = (q_{ij})$  by

$$q_{ij} = \begin{cases} 1 & \text{if } i = j, \\ 1 & \text{if column } A_i \text{ is added into column } A_j, \\ 0 & \text{otherwise.} \end{cases}$$

For example,

$$\begin{array}{ccc} \begin{bmatrix} 1 & 0 & 1 & 1 \\ 0 & 1 & 1 & 0 \\ 1 & 1 & 0 & 0 \\ 0 & 1 & 0 & 1 \end{bmatrix} & \begin{bmatrix} 1 & 1 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 1 \end{bmatrix} & = \left[ A_0 \left| A_0 \oplus A_1 \oplus A_3 \right| A_0 \oplus A_2 \right| A_3 \right] = \begin{bmatrix} 1 & 0 & 0 & 1 \\ 0 & 1 & 1 & 0 \\ 1 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \\ A & Q & A' \end{array}$$

Column-addition matrices are also subject to a *dependency restriction* that if column  $i$  is added into column  $j$ , then column  $j$  cannot be added into any other column. That is, if  $q_{ij} = 1$ , then  $q_{jk} = 0$  for all  $k \neq j$ . The following lemma shows that any column-addition matrix is the product of two nonsingular matrices, and so any column-addition matrix is also nonsingular.

**Lemma 19** *Any column-addition matrix is nonsingular.*

*Proof:* We shall prove by induction on the matrix size  $n$  that any column-addition matrix  $Q$  is the product of two nonsingular matrices  $L$  and  $U$ . Thus, the matrix  $Q$  is also nonsingular.

For the basis, when  $n = 2$ , the only column-addition matrices are  $\begin{bmatrix} 1 & 0 \\ 1 & 1 \end{bmatrix}$ ,  $\begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix}$ , and  $\begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$ . Since each of these matrices is nonsingular, each of them is the product of itself and the identity matrix.

For the inductive step, we assume that every  $(n - 1) \times (n - 1)$  column-addition matrix is the product of two nonsingular matrices. We partition an arbitrary  $n \times n$  column-addition matrix  $Q$  as

$$Q = \left[ \begin{array}{c|c} 1 & n-1 \\ \hline 1 & \psi \\ \theta & \chi \end{array} \right] \begin{array}{c} 1 \\ n-1 \end{array}.$$

The trailing  $(n - 1) \times (n - 1)$  submatrix  $\chi$  is a column-addition matrix because all of its diagonal elements are 1s and, as a submatrix of  $Q$ , it obeys the dependency restriction. By our inductive hypothesis, therefore, the submatrix  $\chi$  is the product of two  $(n - 1) \times (n - 1)$  nonsingular matrices, say  $\nu$  and  $\eta$ . By the dependency restriction, if there are any 1s in  $\theta$ , then there cannot be any 1s in  $\psi$ . Therefore, either  $\psi$  or  $\theta$  is a zero submatrix, and consequently we can factor  $Q$  as

$$Q = \underbrace{\left[ \begin{array}{c|c} 1 & n-1 \\ \hline 1 & 0 \\ \theta & \nu \end{array} \right]}_L \underbrace{\left[ \begin{array}{c|c} 1 & n-1 \\ \hline 1 & \psi \\ 0 & \eta \end{array} \right]}_U \begin{array}{c} 1 \\ n-1 \end{array}.$$

The rightmost  $n - 1$  columns of  $L$  are linearly independent since the submatrix  $\nu$  is nonsingular and the upper right  $1 \times (n - 1)$  submatrix is 0. The leftmost column is linearly independent of the rightmost  $n - 1$  columns since its top entry is 1 and the top entry of each of the rightmost  $n - 1$  columns is 0. Thus,  $L$  is nonsingular. Similarly, because the submatrix  $\eta$  is nonsingular and the lower left  $(n - 1) \times 1$  submatrix of  $U$  is 0, the matrix  $U$  is nonsingular. Thus, any column-addition matrix is the product of two nonsingular matrices, and therefore is also nonsingular. ■

In fact, the factors  $L$  and  $U$  in the proof of Lemma 19 are unit lower-triangular and unit upper-triangular matrices, respectively. Thus, we can factor the example above as

$$Q = \begin{bmatrix} 1 & 1 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 1 \end{bmatrix} = \begin{bmatrix} 1 & 1 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 1 \end{bmatrix} = L U.$$

### Partitioning the matrix to represent operations between column sections

In Section 5, we shall factor nonsingular matrices into column-addition matrices and matrices that characterize MRC permutations. These matrices will be of various block forms, and to classify these forms, we use the following block representation. We partition a matrix into three sections: left, middle, and right. The *left section* includes the leftmost  $b$  columns, the *middle section* includes the middle  $m - b$  columns and the *right section* includes the rightmost  $n - m$  columns. When the

form of a particular submatrix is known, we label that block accordingly. Otherwise, we place an asterisk (\*) in blocks containing indeterminate submatrices.

For column-addition operations, the characteristic matrix has the following form. Every entry on the diagonal is 1. We place an asterisk in each submatrix that contains any non-diagonal 1s as defined by the operation. Returning to the example above, if  $b = 1$  and  $m = 2$ , the form of  $Q$  is

$$Q = \left[ \begin{array}{c|c|c} b=1 & m-b=1 & n-m=2 \\ \hline I & * & * \\ \hline 0 & I & 0 \\ \hline 0 & * & I \end{array} \right] \begin{array}{l} b=1 \\ m-b=1 \\ n-m=2 \end{array} .$$

We define several column-addition operations and MRC permutations by the form of their characteristic matrices. Each of these forms is nonsingular and characterizes a one-pass permutation. We shall show that the inverse of each of these one-pass permutations falls into a specific class of one-pass permutations.

### Trailer matrix form

In Section 5, we shall need to transform a nonsingular matrix into one that has a nonsingular trailing  $(n-m) \times (n-m)$  submatrix. We shall create the nonsingular trailing submatrix by adding some columns from the left and middle sections to the right section. We define a *trailer matrix* as a column-addition matrix that adds some columns from the left and middle sections into the columns of the right section. The matrix  $T$  for this operation is of the form

$$T = \left[ \begin{array}{c|c|c} b & m-b & n-m \\ \hline I & 0 & * \\ \hline 0 & I & * \\ \hline 0 & 0 & I \end{array} \right] \begin{array}{l} b \\ m-b \\ n-m \end{array} .$$

The trailer matrix form characterizes an MRC permutation.

### Reducer matrix form

Once we have a matrix with a nonsingular trailing submatrix, we need an operation that puts the matrix into “reduced form.” (We shall define reduced form precisely in Section 5.) We convert a matrix into reduced form by adding columns from the left and middle sections into other columns in the left and middle sections while respecting the dependency restriction. Thus, a *reducer matrix*  $R$  is a column-addition matrix of the form

$$R = \left[ \begin{array}{c|c|c} b & m-b & n-m \\ \hline * & * & 0 \\ \hline * & * & 0 \\ \hline 0 & 0 & I \end{array} \right] \begin{array}{l} b \\ m-b \\ n-m \end{array} .$$

Since the dependency restriction is obeyed, the leading  $m \times m$  submatrix of  $R$  is nonsingular. Thus, the matrix  $R$  characterizes an MRC permutation.

We can multiply the forms  $T$  and  $R$  to create another matrix form that also characterizes a one-pass permutation. The product  $TR$  results in a matrix of the form

$$P = \left[ \begin{array}{c|c|c} b & m-b & n-m \\ \hline * & * & * \\ \hline * & * & * \\ \hline 0 & 0 & I \end{array} \right] \begin{array}{l} b \\ m-b \\ n-m \end{array} .$$

Since both of the matrix forms  $T$  and  $R$  characterize MRC permutations, by Theorem 18, so does the matrix form  $P$  and its inverse.

### Swapper matrix form

We shall also need to transform the columns in the lower left and lower middle submatrices into columns of zeros. To do so, we must move the nonzero columns in the lower left submatrix into the lower middle submatrix positions by swapping at most  $m-b$  columns at a time from the left section with those in the middle section. Thus, the swap operation is a permutation of the leftmost  $m$  columns. A *swapper matrix* is of the form

$$S = \left[ \begin{array}{c|c} m & n-m \\ \hline \text{permutation} & 0 \\ \hline 0 & I \end{array} \right] \begin{array}{l} m \\ n-m \end{array}$$

so that the leading  $m \times m$  submatrix is a permutation matrix, which dictates the permutation of the leftmost  $m$  columns. The matrix form  $S$  characterizes an MRC permutation and, by Theorem 18, so does its inverse.

### Erasure matrix form

The last operation used in Section 5 is an erasure operation to zero out columns in the lower middle submatrix. To perform this operation, we add columns from the right section into columns in the middle section. Thus, an *erasure matrix* form is defined as

$$E = \left[ \begin{array}{c|c|c} b & m-b & n-m \\ \hline I & 0 & 0 \\ \hline 0 & I & 0 \\ \hline 0 & * & I \end{array} \right] \begin{array}{l} b \\ m-b \\ n-m \end{array} .$$

This matrix form characterizes an MLD permutation because the kernel of  $E_{b..m-1,0..m-1}$  includes only those  $m$ -vectors  $x$  with  $x_{b..m-1} = 0$ , and each such vector is also in the kernel of  $E_{m..n-1,0..m-1}$ . Moreover, observe that any matrix of this form is its own inverse. Consequently, the inverse of such a matrix characterizes an MLD permutation.

## 5 An asymptotically optimal BMBC algorithm

In this section, we present an algorithm to perform any BMBC permutation by factoring its characteristic matrix into matrices which characterize one-pass permutations. We assume that the BMBC permutation is given by an  $n \times n$  characteristic matrix  $A$  and a complement vector  $c$  of length  $n$ . We show that the number of parallel I/Os to perform any BMBC permutation is at

most  $\frac{2N}{BD} \left( \left\lceil \frac{\text{rank } \gamma}{\lg(M/B)} \right\rceil + 2 \right)$  parallel I/Os, where  $\gamma$  is the submatrix  $A_{b..n-1,0..b-1}$ , which appears in the lower bound given by Theorem 3.

Our strategy is to factor the matrix  $A$  into a product of matrices, each of which characterizes an MRC or MLD permutation. For now, we ignore the complement vector  $c$ . According to Corollary 2, we read the factors right to left to determine the order in which to perform the permutations.

To obtain the factoring for  $A$ , we multiply  $A$  by matrices of the forms described in Section 4. By applying these matrix-column operations, we transform the matrix  $A$  into a matrix  $F$  that characterizes an MRC permutation. Multiplying  $F$  by the inverse of each of the matrix-column factors yields the factoring.

### Creating a nonsingular trailing submatrix

We start to transform the characteristic matrix  $A$  by creating a nonsingular matrix  $A^{(1)}$  which has a nonsingular trailing  $(n - m) \times (n - m)$  submatrix. We represent the matrix  $A$  as

$$A = \left[ \begin{array}{c|c} m & n - m \\ \hline \alpha & \beta \\ \phi & \delta \end{array} \right] \begin{array}{l} m \\ n - m \end{array} .$$

Our algorithm depends on the structure of  $\phi$  rather than  $\gamma$ . The following lemma allows us to consider  $\text{rank } \phi$  instead of  $\text{rank } \gamma$  with only a minor difference.

**Lemma 20** *For any matrix  $A$ ,*

$$\text{rank } A_{b..n-1,0..b-1} - \lg(M/B) \leq \text{rank } A_{m..n-1,0..m-1} \leq \text{rank } A_{b..n-1,0..b-1} + \lg(M/B) .$$

*Proof:* Because the rank of a submatrix is the maximum number of linearly independent rows or columns, we have

$$\text{rank } A_{m..n-1,0..b-1} \leq \text{rank } A_{b..n-1,0..b-1} \leq \text{rank } A_{m..n-1,0..b-1} + \lg(M/B) , \quad (12)$$

$$\text{rank } A_{m..n-1,0..b-1} \leq \text{rank } A_{m..n-1,0..m-1} \leq \text{rank } A_{m..n-1,0..b-1} + \lg(M/B) . \quad (13)$$

Subtracting  $\lg(M/B)$  from the right-hand inequality of (12) and combining the result with the left-hand inequality of (13) yields

$$\text{rank } A_{b..n-1,0..b-1} - \lg(M/B) \leq \text{rank } A_{m..n-1,0..b-1} \leq \text{rank } A_{m..n-1,0..m-1} . \quad (14)$$

Adding  $\lg(M/B)$  to the left-hand inequality of (12) and combining the result with the right-hand inequality of (13) yields

$$\text{rank } A_{m..n-1,0..m-1} \leq \text{rank } A_{m..n-1,0..b-1} + \lg(M/B) \leq \text{rank } A_{b..n-1,0..b-1} + \lg(M/B) . \quad (15)$$

Combining inequalities (14) and (15) proves the lemma. ■

By Lemma 20, therefore,

$$\text{rank } \phi \leq \text{rank } \gamma + \lg(M/B) . \quad (16)$$

We shall use this fact later in the analysis of the algorithm to express the bound in terms of  $\text{rank } \gamma$ .

We make the trailing  $(n - m) \times (n - m)$  submatrix nonsingular by adding columns in  $\phi$  into those in  $\delta$ . Consider  $\delta$  as a set of  $n - m$  columns and  $\phi$  as a set of  $m$  columns. We use Gaussian

elimination as described in [3] to determine a maximal set  $V$  of linearly independent columns in  $\delta$  and a set  $W$  of  $n - m - \text{rank } \delta$  columns in  $\phi$  that, along with  $V$ , comprise a set of  $n - m$  linearly independent columns. Denoting by  $\overline{V}$  the  $n - m - \text{rank } \delta$  columns of  $\delta$  not in  $V$ , we make the trailing submatrix of  $A$  nonsingular by pairing up columns of  $W$  with columns of  $\overline{V}$  and adding each column in  $W$  into its corresponding column in  $\overline{V}$ .

We must express the above transformation as a column-addition operation. Although we focused above on adding columns of  $\phi$  to columns of  $\delta$ , column-addition operations add entire columns, and so we must also add the corresponding columns of  $\alpha$  to the corresponding columns of  $\beta$ . Since we add columns from the leftmost  $m$  columns of  $A$  to the rightmost  $n - m$  columns, the characteristic matrix of this operation has the trailer matrix form  $T$  described in Section 4. The matrix product is now

$$A^{(1)} = AT = \left[ \begin{array}{c|c} m & n-m \\ \hline \alpha & \hat{\beta} \\ \phi & \hat{\delta} \end{array} \right] \begin{array}{l} m \\ n-m \end{array},$$

where  $\hat{\delta}$  is nonsingular. Since the matrices  $A$  and  $T$  are nonsingular, the matrix  $A^{(1)}$  is nonsingular.

### Transforming the matrix into reduced form

The next step is to transform the matrix  $A^{(1)}$  into reduced form. For our purposes, a matrix is in *reduced form* when there are  $\text{rank } \phi$  linearly independent columns and  $m - \text{rank } \phi$  columns of zeros in the lower left  $(n - m) \times m$  submatrix and the trailing  $(n - m) \times (n - m)$  submatrix is nonsingular. Once again, we use Gaussian elimination to determine a set  $U$  that indexes  $\text{rank } \phi$  linearly independent columns of  $\phi$ . To perform the reduction, we determine for each linearly dependent column  $\phi_j$  a set of column indices  $U_j \subseteq U$  such that  $\phi_j = \oplus_{k \in U_j} \phi_k$ . Adding the set of columns of  $\phi$  indexed by  $U_j$  into  $\phi_j$  zeros it out. We add linearly independent columns from the left and middle sections into the linearly dependent columns of these sections. Since we never add a linearly dependent column into any other column and there are no column additions into the linearly independent columns, we respect the dependency restriction. The matrix  $R$  that reduces the matrix  $A^{(1)}$  has the reducer matrix form described in Section 4. Thus, the matrix product  $TR$  is of the form  $P$  also described in Section 4, and it characterizes an MRC permutation. We now have the product

$$A^{(2)} = A^{(1)}R = ATR = AP = \left[ \begin{array}{c|c} m & n-m \\ \hline \hat{\alpha} & \hat{\beta} \\ \hat{\phi} & \hat{\delta} \end{array} \right] \begin{array}{l} m \\ n-m \end{array},$$

with a nonsingular trailing submatrix  $\hat{\delta}$  and a lower left submatrix  $\hat{\phi}$  in reduced form. Since  $A$  and  $P$  are nonsingular, the matrix  $A^{(2)}$  is also nonsingular.

### Zeroing out the lower left submatrix

Our eventual goal is to transform the original matrix  $A$  into a matrix  $F$  that characterizes an MRC permutation. At this point, the matrix  $A$  has been transformed into the nonsingular matrix  $A^{(2)}$ . Thus, our final task is to multiply  $A^{(2)}$  by a series of matrices that transform the  $\text{rank } \phi$  nonzero columns in the lower left  $(n - m) \times m$  submatrix  $\hat{\phi}$  into columns of zeros. We do so by multiplying  $A^{(2)}$  by matrices of the swapper and erasure matrix forms described in Section 4. Let us further

partition the leftmost  $m$  columns of  $A^{(2)}$  into the leftmost  $b$  columns and the middle  $m - b$  columns:

$$A^{(2)} = \left[ \begin{array}{c|c|c} b & m-b & n-m \\ \hline \hat{\alpha}' & \hat{\alpha}'' & \hat{\beta} \\ \hline \hat{\phi}' & \hat{\phi}'' & \hat{\delta} \end{array} \right] \begin{array}{l} m \\ n-m \end{array} .$$

Our strategy is to repeatedly use swapper matrix forms to move at most  $m - b$  columns from the left section into the middle section and then zero out those columns using erasure matrix forms.

We begin by swapping  $m - b - \text{rank } \hat{\phi}''$  columns from  $\hat{\phi}'$  with the zero columns of  $\hat{\phi}''$ . Multiplying the matrix  $A^{(2)}$  by a nonsingular matrix  $S_1$  of the swapper matrix form described in Section 4, we swap at most  $m - b$  columns from the left section with the appropriate columns in the middle section. After performing the swap operation on the matrix  $A^{(2)}$ , the lower left submatrix has  $m - b - \text{rank } \hat{\phi}''$  additional zero columns and the lower middle submatrix has full rank. The above discussion assumes that  $\text{rank } \hat{\phi}' \geq m - b - \text{rank } \hat{\phi}''$ ; if the opposite holds, we swap  $\text{rank } \hat{\phi}'$  columns and the lower left submatrix becomes all zeros.

Our next step is to transform the  $m - b$  columns in the lower middle submatrix into columns of zeros. Since the nonsingular trailing  $(n - m) \times (n - m)$  submatrix  $\hat{\delta}$  forms a basis for the columns of the lower  $n - m$  rows of matrix  $A^{(2)}$ , we zero out each nonzero column in the lower middle submatrix by adding columns of  $\hat{\delta}$ . Since we add columns from the rightmost  $n - m$  columns into the middle  $m - b$  columns, we perform the matrix-column operation characterized by a nonsingular matrix  $E_1$  of the erasure matrix form described in Section 4. After multiplying by the erasure matrix  $E_1$ , the original matrix is transformed into a nonsingular matrix

$$A^{(3)} = A P S_1 E_1 ,$$

which has zero columns in the lower middle  $(n - m) \times (m - b)$  submatrix and possibly some more nonzero columns in the lower left  $(n - m) \times b$  submatrix.

If there are still nonzero columns in the lower left  $(n - m) \times b$  submatrix of the matrix  $A^{(3)}$ , then those columns must also be swapped into the lower middle  $(n - m) \times (m - b)$  submatrix by a swapper matrix and transformed into zero columns by an erasure matrix. We repeatedly swap in up to  $m - b$  nonzero columns of the lower left submatrix into the lower middle submatrix. Each time we perform a swap operation, we multiply the current product by a matrix  $S_i$  of the swapper matrix form. Note that we swap entire columns here, not just the portions in the lower submatrices. After we perform each matrix-column operation  $S_i$ , we zero out the lower middle submatrix by multiplying the current product by a matrix  $E_i$  of the erasure matrix form.

After repeatedly swapping and erasing each of the nonzero columns in the lower left  $(n - m) \times m$  submatrix, the lower left submatrix will contain only zero columns. This matrix is the matrix  $F$  mentioned at the beginning of this section. Since the matrix  $A^{(2)}$  is in reduced form, there are at most  $\text{rank } \phi$  columns in the submatrix  $\hat{\phi}$  that need to be transformed into zero columns. Thus, at most

$$g = \left\lceil \frac{\text{rank } \phi}{m - b} \right\rceil \quad (17)$$

pairs of swap and erasure operations transform all the columns in the lower left  $(n - m) \times m$  submatrix into zero columns. Since each matrix-column operation that we performed on the original matrix  $A$  to transform it into  $F$  is nonsingular, the resulting matrix product

$$F = A P S_1 E_1 S_2 E_2 \cdots S_g E_g$$

is a nonsingular matrix that characterizes an MRC permutation. Multiplying both sides by the inverses of the factors yields the desired factorization of  $A$ :

$$A = F E_g^{-1} S_g^{-1} E_{g-1}^{-1} S_{g-1}^{-1} \cdots E_1^{-1} S_1^{-1} P^{-1} . \quad (18)$$



## Analysis

We now apply several properties that we have gathered to complete the analysis of our BMMC permutation factoring method.

**Theorem 21** *We can perform any BMMC permutation with characteristic matrix  $A$  and complement vector  $c$  in at most*

$$\frac{2N}{BD} \left( \left\lceil \frac{\text{rank } \gamma}{\lg(M/B)} \right\rceil + 2 \right)$$

*parallel I/Os, where  $\gamma = A_{b..n-1, 0..b-1}$ .*

*Proof:* Ignore the complement vector  $c$  for the moment. In the factorization (18) of  $A$ , both factors  $S_1^{-1}$  and  $P^{-1}$  characterize MRC permutations. By Theorem 18, therefore, so does the product  $S_1^{-1} P^{-1}$ . As we saw in Section 4, each factor  $E_i^{-1}$  characterizes an MLD permutation. Applying Theorem 17, each grouping of factors  $E_1^{-1} S_1^{-1} P^{-1}$  and  $E_i^{-1} S_i^{-1}$ , for  $i = 2, 3, \dots, g$ , characterizes an MLD permutation. By Theorem 15, and adding in one more pass for the MRC permutation characterized by  $F$ , we can perform  $A$  with  $g + 1$  passes.

If the complement vector  $c$  is nonzero, we include it as part of the MRC permutation characterized by the leftmost factor  $F$ .

Regardless of the complement vector, therefore, we can perform the BMMC permutation with  $g + 1$  passes. Combining equation (17) and inequality (16), we obtain a bound of

$$\begin{aligned} g + 1 &= \left\lceil \frac{\text{rank } \phi}{\lg(M/B)} \right\rceil + 1 \\ &\leq \left\lceil \frac{\text{rank } \gamma + \lg(M/B)}{\lg(M/B)} \right\rceil + 1 \\ &= \left\lceil \frac{\text{rank } \gamma}{\lg(M/B)} \right\rceil + 2 \end{aligned}$$

passes for a total of at most

$$\frac{2N}{BD} \left( \left\lceil \frac{\text{rank } \gamma}{\lg(M/B)} \right\rceil + 2 \right)$$

parallel I/Os. ■

## 6 Detecting BMMC permutations at run time

In practice, we wish to run the BMMC algorithm of Section 5 whenever possible to reap the savings over having run the more costly algorithm for general permutations. For that matter, we wish to run even faster algorithms for any of the special cases of BMMC permutations (MRC, MLD, or block BMMC [4]) whenever possible as well. We must know the characteristic matrix  $A$  and complement vector  $c$ , however, to run any of these algorithms. If  $A$  and  $c$  are specified in the source code, before running the algorithm we only need to check that  $A$  is of the correct form, e.g., that it is nonsingular for a BMMC permutation, of the MLD or MRC form, etc. Later in this section, we show how to check the kernel condition for MLD permutations. If instead the permutation is given by a vector of  $N$  target addresses, we can detect at run time whether it is a BMMC permutation by the following procedure:

1. Check that  $N$  is a power of 2.

2. Form a candidate characteristic matrix  $A$  and complement vector  $c$  such that if the permutation is BMMC, then  $A$  and  $c$  must be the correct characterizations. This section shows how to do so with only  $\left\lceil \frac{\lg(N/B)+1}{D} \right\rceil$  parallel reads.
3. Check that the characteristic matrix is of the correct form.
4. Verify that all  $N$  target addresses are described by the candidate characteristic matrix and complement vector. If for any source address  $x$  and its corresponding target address  $y$  we have  $y \neq Ax \oplus c$ , the permutation is not BMMC and we can terminate verification. If  $y = Ax \oplus c$  for all  $N$  source-target pairs, the permutation is BMMC. Verification uses at most  $N/BD$  parallel reads.

The total number of parallel I/Os is at most

$$\frac{N}{BD} + \left\lceil \frac{\lg(N/B) + 1}{D} \right\rceil ,$$

all of which are reads, and it is usually far fewer when the permutation turns out not to be BMMC.

One benefit of run-time BMMC detection is that the programmer might not realize that the desired permutation is BMMC. For example, as noted in Section 1, the standard binary reflected Gray code and its inverse are both MRC permutations. Yet the programmer might not know to call a special MRC or BMMC routine. Even if the system provides an entry point to perform the standard Gray code permutation and this routine invokes the MRC algorithm, variations on the standard Gray code may foil this approach. For example, a standard Gray code with all bits permuted the same (i.e., a characteristic matrix of  $\Pi G$ , where  $\Pi$  is a permutation matrix and  $G$  is the MRC matrix that characterizes the standard Gray code) is BMMC but not necessarily MRC. It might not be obvious enough that the permutation characterized by  $\Pi G$  is BMMC for the programmer to invoke the BMMC algorithm explicitly.

### Forming the candidate characteristic matrix and complement vector

The method for forming the candidate characteristic matrix  $A$  and candidate complement vector  $c$  is based on two observations. First, if the permutation is BMMC, then the complement vector  $c$  must be the target address corresponding to source address 0. This relationship holds because  $x = 0$  and  $y = Ax \oplus c$  imply that  $y = c$ .

The second observation is as follows. Consider a source address  $x = (x_0, x_1, \dots, x_{n-1})$ , and suppose that bit position  $k$  holds a 1, i.e.,  $x_k = 1$ . Let us denote the  $j$ th column for matrix  $A$  by  $A_j$ . Also, let  $S_k$  denote the set of bit positions in  $x$  other than  $k$  that hold a 1:  $S_k = \{j : j \neq k \text{ and } x_j = 1\}$ . If  $y = Ax \oplus c$ , then we have

$$y = \left( \bigoplus_{j \in S_k} A_j \right) \oplus A_k \oplus c , \quad (19)$$

since only the bit positions  $j$  for which  $x_j = 1$  contribute a column of  $A$  to the sum of columns that forms the matrix-vector product. If we know the target address  $y$ , the complement vector  $c$  and the columns  $A_j$  for all  $j \neq k$ , we can rewrite equation (19) to yield the  $k$ th column of  $A$ :

$$A_k = y \oplus \left( \bigoplus_{j \in S_k} A_j \right) \oplus c . \quad (20)$$

We shall compute the complement vector  $c$  first and then the columns of the characteristic matrix  $A$  one at a time, from  $A_0$  up to  $A_{n-1}$ . When computing  $A_k$ , we will have already computed  $A_0, A_1, \dots, A_{k-1}$ , and these will be the only columns we need in order to apply equation (20). In other words,  $S_k \subseteq \{0, 1, \dots, k-1\}$ . Recall that as Figure 2 shows, the lower  $b$  bits of a record's address give the record's offset within its block, the middle  $d$  bits give the disk number, and the upper  $s = n - (b + d)$  bits give the stripe number.

From equation (20), it would be easy to compute  $A_k$  if  $S_k$  were empty. The set  $S_k$  is empty if the source address is a unit vector, with its only 1 in position  $k$ . If we look at these addresses, however, we find that the target addresses for a disproportionate number—all but  $d$  of them—reside on disk  $\mathcal{D}_0$ . The block whose disk and stripe fields are all zero contains  $b$  such addresses, so they can be fetched in one disk read. A problem arises for the  $s$  source addresses with one 1 in the stripe field: their target addresses all reside on different blocks of disk  $\mathcal{D}_0$ . Each must be fetched in a separate read. The total number of parallel reads to fetch all the target addresses corresponding to all unit-vector source addresses is  $s + 1 = \lg(N/BD) + 1$ .

To achieve only  $\left\lceil \frac{\lg(N/B)+1}{D} \right\rceil$  parallel reads, each read fetches one block from each of the  $D$  disks. The first parallel read determines the complement vector, the first  $b + d$  columns, and the next  $D - d - 1$  columns. Each subsequent read determines another  $D$  columns, until all  $n$  columns have been determined.

In the first parallel read, we do the same as above for the first  $b + d$  bits. That is, we fetch blocks containing target addresses whose corresponding source addresses are unit vectors with one 1 in the first  $b + d$  positions. As before,  $b$  of them are in the same block on disk  $\mathcal{D}_0$ . This block also contains address 0, which we need to compute the complement vector. The remaining  $d$  are in stripe number 0 of disks  $\mathcal{D}_1, \mathcal{D}_2, \mathcal{D}_4, \mathcal{D}_8, \dots, \mathcal{D}_{D/2}$ . Having fetched the corresponding target addresses, we have all the information we need to compute the complement vector  $c$  and columns  $A_0, A_1, \dots, A_{b+d-1}$ .

The columns we have yet to compute correspond to bit positions in the stripe field. If we were to compute these columns in the same fashion as the first  $b + d$ , we would again encounter the problem that all the blocks we need to read are on disk  $\mathcal{D}_0$ . In the first parallel read, the only unused disks remaining are those whose numbers are not a power of 2 ( $\mathcal{D}_3, \mathcal{D}_5, \mathcal{D}_6, \mathcal{D}_7, \mathcal{D}_9, \dots$ ). The key observation is that we have already computed all  $d$  columns corresponding to the disk field, and we can thus apply equation (20). For example, let us compute column  $A_{b+d}$ , which corresponds to the first bit of the stripe number. We read stripe 1 on disk  $\mathcal{D}_3$  and find the first target address  $y$  in this block. Disk number 3 corresponds to the first two disk-number columns,  $A_b$  and  $A_{b+1}$ . Applying equation (20) with  $S_{b+d} = \{b, b+1\}$ , we compute  $A_{b+d} = y \oplus A_b \oplus A_{b+1} \oplus c$ . The next column we compute is  $A_{b+d+1}$ . Reading the block at stripe 2 on disk  $\mathcal{D}_5$ , we fetch a target address  $y$  and then compute  $A_{b+d+1} = y \oplus A_b \oplus A_{b+2} \oplus c$ . Continuing on in this fashion, we compute a total of  $D - d - 1$  stripe-bit columns from the first parallel read.

The remaining parallel reads compute the remaining stripe-bit columns. We follow the stripe-bit pattern of the first read, but we use all disks, not just those whose disk numbers are not powers of 2. Each block read fetches a target address  $y$ , which we exclusive-or with a set of columns from the disk field and with the complement vector to compute a new column from the stripe field. The first parallel read computes  $b + D - 1$  columns and all subsequent parallel reads compute  $D$  columns. The total number of parallel reads is thus

$$\begin{aligned} 1 + \left\lceil \frac{n - (b + D - 1)}{D} \right\rceil &= 1 + \left\lceil \frac{\lg(N/B) - D + 1}{D} \right\rceil \\ &= \left\lceil \frac{\lg(N/B) + 1}{D} \right\rceil. \end{aligned}$$

## Checking the kernel condition for MLD permutations

In practice, we would like a simple procedure to verify that a given matrix characterizes an MLD permutation. By Gaussian elimination, it is easy to verify that a candidate matrix  $A$  is nonsingular. It may not be obvious how to verify that  $\ker \lambda \subseteq \ker \mu$  when the matrix is blocked into  $\lambda$  and  $\mu$ . The following two steps suffice:

- Find a basis of  $m$ -vectors for  $\ker \lambda$ . Since  $\dim(\text{row } \lambda) = m - b$  for MLD permutations and  $\lambda$  has  $m$  columns,  $\dim(\ker \lambda) = b$ . Any basis of  $\ker \lambda$ , therefore, consists of exactly  $b$  vectors. If the basis found contains more than  $b$  vectors, the matrix cannot characterize an MLD permutation.
- Letting  $\{x^{(0)}, x^{(1)}, \dots, x^{(b-1)}\}$  be the basis found for  $\ker \lambda$ , verify that  $\mu x^{(i)} = 0$  for  $i = 0, 1, \dots, b - 1$ .

This method checks that any vector in  $\ker \lambda$  is also a vector in  $\ker \mu$ . For all  $x \in \ker \lambda$ ,  $x$  is a sum of some subset of the basis vectors of  $\ker \lambda$ . Thus, if all the basis vectors map to 0 by the linear transformation  $\mu$ , then all of the vectors  $x \in \ker \lambda$  map to 0 by the linear transformation  $\mu$ . For all  $x \in \ker \lambda$ , therefore,  $\mu x = 0$  implies  $x \in \ker \mu$ , which in turn implies that  $\ker \lambda \subseteq \ker \mu$ .

## 7 Conclusions

This paper has shown an asymptotically tight bound on the number of parallel I/Os required to perform BMMC permutations on parallel disk systems. It is particularly satisfying that the tight bound was achieved not by raising the lower bound proven here and in [3], but by decreasing the upper bound in [4]. The multiplicative and additive constants in the I/O complexity of our algorithm are small, which is especially fortunate in light of the expense of disk accesses.

One can adapt the proof by Aggarwal and Vitter [2] of Lemma 6 to bound  $\Delta \Phi_{\max}$  precisely, rather than just asymptotically. In particular, it is a straightforward exercise to derive the bound

$$\Delta \Phi_{\max} \leq B \left( \frac{2}{e \ln 2} + \lg(M/B) \right) .$$

Moreover, the potential change is at most zero for write operations, and so the potential increases only during read operations. If all I/Os are simple, then the total number of blocks read equals the total number of blocks written. Therefore, we can modify the lower bound of Lemma 5 to  $2 \frac{\Phi(t) - \Phi(0)}{\Delta \Phi_{\max}}$ , with which we can derive a lower bound of

$$\frac{2N}{BD} \frac{\text{rank } \gamma}{\frac{2}{e \ln 2} + \lg(M/B)}$$

parallel I/Os for any BMMC permutation. Since the quantity  $2/(e \ln 2)$  is approximately 1.06, this lower bound is quite close to the exact upper bound given by Theorem 21.

We have also shown how to detect BMMC permutations at run time, given a vector of target addresses. Detection is inexpensive and, when successful, permits the execution of our BMMC algorithm or possibly a faster algorithm for a more restricted permutation class.

What other permutations can be performed quickly? Several  $O(1)$ -pass permutation classes appear in [3], and this paper has added one more (MLD permutations in Section 3). One can also show that the inverse of any one-pass permutation is a one-pass permutation. Moreover, the composition of an MLD permutation with the inverse of an MLD permutation is a one-pass permutation. We also ask what other useful permutation classes can we show to be BMMC?

Finally, we ask whether the lower bound of  $\Omega\left(\frac{\Phi(t)-\Phi(0)}{\Delta\Phi_{\max}}\right)$  parallel I/Os is universally tight for all permutations, not just those that are BMMC. One possible approach is to design an algorithm that explicitly manages the potential. If each pass increases the potential by  $\Theta\left(\frac{N}{BD}(\Delta\Phi_{\max})\right)$ , the algorithm's I/O count would match the lower bound.

## Acknowledgments

Thanks to C. Esther Jeserum and Michael Klugerman for their helpful suggestions.

## References

- [1] A. Aggarwal, A. K. Chandra, and M. Snir. Hierarchical memory with block transfer. In *Proceedings of the 28th Annual Symposium on Foundations of Computer Science*, pages 204–216, Oct. 1987.
- [2] A. Aggarwal and J. S. Vitter. The input/output complexity of sorting and related problems. *Commun. ACM*, 31(9):1116–1127, Sept. 1988.
- [3] T. H. Cormen. *Virtual Memory for Data-Parallel Computing*. PhD thesis, Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, 1992. Available as Technical Report MIT/LCS/TR-559.
- [4] T. H. Cormen. Fast permuting in disk arrays. *Journal of Parallel and Distributed Computing*, 17(1–2):41–57, Jan. and Feb. 1993.
- [5] A. Edelman, S. Heller, and S. L. Johnsson. Index transformation algorithms in a linear algebra framework. Technical Report LBL-31841, Lawrence Berkeley Laboratory, 1992.
- [6] R. W. Floyd. Permuting information in idealized two-level storage. In R. E. Miller and J. W. Thatcher, editors, *Complexity of Computer Computations*, pages 105–109. Plenum Press, 1972.
- [7] M. T. Goodrich, J.-J. Tsay, D. E. Vengroff, and J. S. Vitter. External-memory computational geometry. In *Proceedings of the 34th Annual Symposium on Foundations of Computer Science*, pages 714–723, Nov. 1993.
- [8] S. L. Johnsson and C.-T. Ho. Generalized shuffle permutations on boolean cubes. Technical Report TR-04-91, Harvard University Center for Research in Computing Technology, Feb. 1991.
- [9] M. H. Nodine and J. S. Vitter. Greed sort: An optimal external sorting algorithm for multiple disks. Technical Report CS-91-20, Department of Computer Science, Brown University, 1991.
- [10] M. H. Nodine and J. S. Vitter. Large-scale sorting in parallel memories. In *Proceedings of the 3rd Annual ACM Symposium on Parallel Algorithms and Architectures*, pages 29–39, July 1991.
- [11] M. H. Nodine and J. S. Vitter. Optimal deterministic sorting on parallel disks. Technical Report CS-92-08, Department of Computer Science, Brown University, 1992.
- [12] G. Strang. *Linear Algebra and Its Applications*. Harcourt Brace Jovanovich, San Diego, third edition, 1988.

- [13] J. S. Vitter and E. A. M. Shriver. Optimal disk I/O with parallel block transfer. In *Proceedings of the Twenty Second Annual ACM Symposium on Theory of Computing*, pages 159–169, May 1990.
- [14] J. S. Vitter and E. A. M. Shriver. Algorithms for parallel memory I: Two-level memories. Technical Report CS-92-04, Department of Computer Science, Brown University, Aug. 1992. Revised version of Technical Report CS-90-21.