NPAC Technical Report SCCS–622, Sept. 1994

Data Access Reorganizations in Compiling Out-of-core Data Parallel Programs on Distributed Memory Machines*

Rajesh Bordawekar **Alok Choudhary** **Rajeev Thakur**
Dept. of Elect. and Comp. Eng. and NPAC
Syracuse University

Northeast Parallel Architectures Center
111 College Place, Rm. 3-228
Syracuse University
Syracuse, NY 13244-4100
Tel : (315) 443-4061
Fax : (315) 443-1973

Abstract

This paper describes techniques for translating out-of-core programs written in a data parallel language like HPF to message passing node programs with explicit parallel I/O. We describe the basic compilation model and various steps involved in the compilation. The compilation process is explained with the help of an out-of-core matrix multiplication program. We first discuss how an out-of-core program can be translated by extending the method used for translating in-core programs. We demonstrate that *straightforward extension of in-core compiler does not work for out-of-core programs.* We then describe how the compiler can optimize the code by (1) estimating the I/O costs associated with different array access patterns, (2) reorganizing array accesses, (3) selecting the method with the least I/O cost, and (4) allocating memory according to access cost for competing out-of-core arrays. These optimizations can reduce the amount of I/O by as much as an order of magnitude. Performance results on the Intel Touchstone Delta are presented and analyzed.

*This work was supported in part by NSF Young Investigator Award CCR-9357840, a grant from Intel SSD, and IBM, and in part by USRA CESDIS Contract # 5555-26. This work was performed in part using the Intel Touchstone Delta System operated by Caltech on behalf of the Concurrent Supercomputing Consortium. Access to this facility was provided by CRPC.

1 Introduction

The use of massively parallel machines to solve large scale computational problems in physics, chemistry, biology, engineering, medicine and other sciences has increased considerably in recent times. This is primarily due to the tremendous improvements in the computational speeds of parallel computers in the last few years. Many of these applications, also referred to as *Grand Challenge Applications* [CR93], have computational requirements which stretch the capabilities of even the fastest supercomputer available today.

In addition to requiring a great deal of computational power, these applications usually deal with large quantities of data. At present, a typical Grand Challenge Application could require 1Gbyte to 4Tbytes of data per run [dRC94]. Main memories are not large enough to hold this much amount of data; so data needs to be stored on disks and fetched during the execution of the program. Unfortunately, the performance of the I/O subsystems of massively parallel computers has not kept pace with their processing and communications capabilities [CFPB93]. Hence, the performance bottleneck is the time taken to perform disk I/O. The need for high performance I/O is so significant that almost all the present generation parallel computers provide some kind of hardware and software support for parallel I/O. An overview of the various issues involved in high performance I/O is given in [dRC94].

Data parallel languages like HPF [For93] and pC++ [BBG+93] have recently been developed to provide support for portable high performance programming on parallel machines. In order that these languages can be used for large scale scientific computations, support for performing large scale I/O from programs written in these languages is necessary. Issue of providing language support for high performance I/O has been addressed recently [CMZ92, Sni92]. It is, therefore, essential to provide compiler support for these languages so that the programs can be translated automatically and efficiently.

1.1 Contributions of the paper

In this paper we describe data access reorganization strategies for efficient compilation of out-of-core data parallel programs on distributed memory machines. This paper builds on our previous work on basic compilation techniques for out-of-core data parallel programs [TBC94a, TBC+94b]. In particular, this addresses the following issues, 1) how to estimate the I/O costs associated with different accesse patterns in out-of-core computations, 2) how to reorganize data storage on disks to reduce I/O costs, 3) how to reorganize computations based on the reorganized data, and 4) when multiple out-of-core arrays are involved in the computations, how to allocate memory to individual arrays to minimize I/O accesses. We demonstrate that these techniques can reduce I/O costs by as much as an order of magnitude compared to the costs of I/O when in-core compilation methods are straight forwardly extended for out-of-core computations. The compilation process is explained with the help of an out-of-core matrix multiplication program which uses a distributed GAXPY algorithm. This program is used only as an example to illustrate the various issues involved in compiling out-of-programs and optimizing the I/O requirements. The techniques described in this paper are applicable to other data parallel languages. Performance results on the Intel Touchstone Delta are presented and analyzed.

1.2 Organization

The rest of the paper is organized as follows. Section 2 describes the basic model used for out-of-core compilation. The compilation methodology is discussed in Section 3. Section 4 describes the I/O cost estimation and optimizations performed by the compiler, followed by conclusions in Section 5. In this paper, the term *in-core compiler* refers to a compiler for in-core programs and the term *out-of-core compiler* refers to a compiler for out-of-core programs.

2 Model for Out-of-Core Compilation

2.1 Programming Model

The most widely used programming model for large-scale scientific and engineering applications on distributed memory machines is the Single Program Multiple Data (SPMD) model. In this model, parallelism is achieved by partitioning data among processors which effectively represents parallelism in a class of applications called *loosely synchronous* applications [Fox91]. To achieve load-balance, express locality of access, reduce communication and other optimizations, several distribution and data alignment strategies are often used (eg., block, cyclic, along rows, columns, etc.). Many parallel programming languages or language extensions have been developed which support such distributions. These languages provide directives that enable the expression of mappings from the problem domain to the processing domain and allow the user to align and distribute arrays in the most appropriate fashion for the underlying computation. The compiler uses the information provided by these directives to compile global name space programs for distributed memory computers. Examples of parallel languages which support data distribution include Vienna Fortran [ZBC+92], Fortran D [FHK+90] and High Performance Fortran (HPF) [For93, KLS+94]. In this paper, we describe the compilation of out-of-core HPF programs, but the discussion is applicable to any other data parallel language in general.

The `DISTRIBUTE` directive in HPF specifies which elements of the array are mapped to each processor. This results in each processor having a *local array* associated with it. In an in-core program, the local array resides in a local memory of the processor. Our group at Syracuse University has developed a compiler for in-core HPF programs [BCF+93]. For large data sets, however, local arrays cannot entirely fit in main memory. In such cases, parts of the local array have to be stored on disk. We refer to such a local array as an **Out-of-core Local Array (OCLA)**. Parts of the OCLA need to be swapped between main memory and disk during the course of the computation.

2.2 Architectural Model

Figure 1 describes the architectural model used by the compiler. It assumes any general distributed memory computer in which the processors are connected together in some fashion. The system is provided with a set of disks. Each processor may either have its own local disk or all processors may share the set of disks. The system is provided with dedicated I/O processors which control the flow of data between the compute processors and the disks. The I/O subsystem may have a separate interconnection network or it can share the same network which connects the processors together.

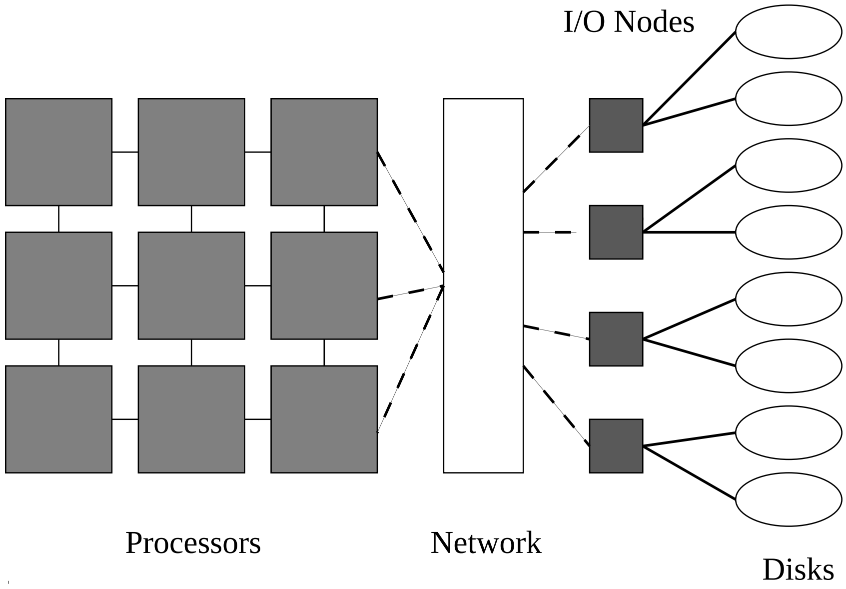

Figure 1: Architectural Model

2.3 Data Storage Model

The Data Storage Model shown in Figure 2 specifies how the out-of-core array is placed on disks and how it is accessed by the processors. The out-of-core local array of each processor is stored in a separate file called the **Local Array File (LAF)** of that processor. The LAF can be assumed to be *owned* by that processor. The node program explicitly reads from and writes into the LAF when required. If the I/O architecture of the system is such that each processor has its own disk, such as in the IBM SP-1, the LAF of each processor will be stored on the disk attached to that processor. If there is a common set of disks for all processors, such as on the Intel Paragon, the LAF will be distributed across one or more of these disks. In other words, we assume that each processor has its own *logical disk* with the LAF stored on that disk. The mapping of the logical disk to the physical disks is system dependent.

A simple way to view this model is to think of each processor as having another level of memory (logical disk) which is much slower than the main memory. Both the main memory and this additional memory cannot be directly accessed by any other processor. Hence, a processor cannot directly access some other processor's LAF. If a processor needs data from the LAF of another processor, the required data will be first read by the owner processor and then communicated to the requesting processor.

In order to store data on the disks based on the distribution pattern specified in the program, redistribution of data may be needed in the beginning when data is first stored on disk. This is because the way data arrives (eg. from archival storage, satellite or over the network) may not conform to the distribution specified in the program. Redistribution requires reading data from disks, communicating data between processors and writing the data to the local array files. This involves some additional overhead which can be amortized if the array is used several times (eg. for many iterations).

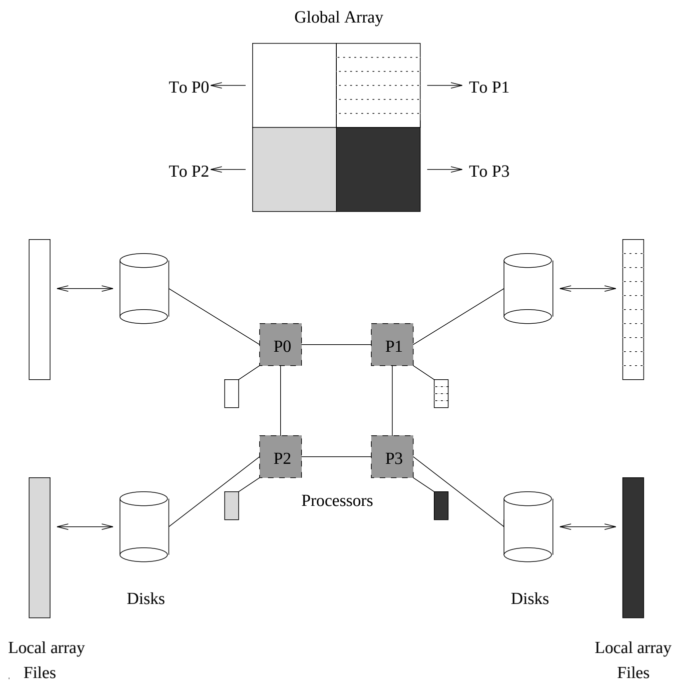

Figure 2: Data Storage Model

3 Compilation Methodology

This section describes the methodology used for compiling out-of-core HPF programs. We use an out-of-core matrix multiplication program example to illustrate the compilation and optimization process. We have chosen this example because it clearly brings out many of the important issues involved in compiling out-of-core programs. This example is used only for explanatory purposes and the methodology described in this paper is applicable to any other program in general.

We first describe the global name space matrix multiplication program and then explain how it is translated assuming that all the matrices are in-core. This is helpful in understanding how the program needs to be compiled when the arrays are out-of-core.

3.1 GAXPY Algorithm for Matrix Multiplication

Let A, B and C be $n \times n$ matrices such that $C = A \times B$. A, B and C can be represented in terms of their individual columns as

$$\begin{aligned} A &= [a_1, \cdots, a_n],\ a_j \in \mathcal{R}^n \\ B &= [b_1, \cdots, b_n],\ b_j \in \mathcal{R}^n \\ C &= [c_1, \cdots, c_n],\ c_j \in \mathcal{R}^n \end{aligned}$$

Then the GAXPY algorithm for computing $C = A \times B$ is

$$c_j = \sum_{k=1}^{n} b_{kj} a_k,\ \ j = 1:n \tag{1}$$

```
1            parameter (n=64, nprocs=4)
2            real a(n,n), b(n,n), c(n,n), temp(n,n)
3       !hpf$      processors Pr(nprocs)
4       !hpf$      template d(n)
5       !hpf$      distribute d(block) on Pr
6       !hpf$      align (*,:) with d :: a, c, temp
7       !hpf$      align (:,*) with d :: b
8            do j=1, n
9                FORALL (k=1:n)
10                    temp(1:n,k) = b(k,j)×a(1:n,k)
11                 end FORALL
12                 c(1:n,j) = SUM(temp,2)      ! Sum Intrinsic
13            end do
14            end
```

Figure 3: GAXPY Matrix Multiplication in HPF

We can see that in order to compute the j^{th} column of C, we need the j^{th} column of B and all columns of A. Figure 3 shows the HPF program for GAXPY matrix multiplication. Arrays A and C are distributed in column-block fashion whereas array B is distributed in row-block fashion over 4 processors. A temporary array is needed to store the products of element b_{kj} and column a_k, which can be computed for all k in parallel. The j^{th} column of the result is computed using the intrinsic function SUM.

3.2 In-core Compilation

Our research group at Syracuse University has developed a compiler to translate in-core HPF programs to message passing node programs for distributed memory machines [BCF+93]. We will focus on compilation of `FORALL` statements in the HPF program.[1] A `FORALL` statement is essentially a parallel loop with copy-in-copy-out semantics [For93]. According to the HPF specifications, the following steps describe a correct sequential implementation of a `FORALL` statement; 1) copy rhs, 2) synchronize, 3) evaluate expression, 4) synchronize, 5) assign to lhs. Note that synchronization and copying are only part of the specification and can be avoided in most cases with appropriate compiler analysis [BCF+93].

The compiler uses distribution directives (Figure 3, lines 3–7) in the source program to find the distribution pattern of the arrays. Using the data distribution information, the arrays are partitioned into local arrays. After data distribution, the compiler analyzes the array operations (Figure 3, lines 8–13). The compiler checks that the outer loop (lines 8–13) is a sequential DO loop whereas the inner loop (lines 9–11) is a `FORALL` construct. The inner `FORALL` loop (indexed by variable k) is sequentialized into local DO loops. After the local computation is done, the temporary results are added to give the j^{th} column of the resultant C. This operation is performed using a global sum reduction routine. Using the knowledge that the index j is in global name space and

[1] Any array assignment statement can be converted into a corresponding `FORALL` statement, so we will use them interchangeably [FHK+90, For93].

```
        parameter (n=64, nproc=4, local_n=16)
C       Partition the arrays using the distribution information.
        real a(n,local_n), b(local_n,n), c(n,local_n)
        do j=1, n
             Initialize the temporary array.
             do i= 1, local_n
                  do k=1, n
                       temp(k,i) = a(k,i)×b(i,j)
                  end do
             end do
C       Perform Global Sum of the temporary arrays along dimension 2.
             result = global_sum(temp, 2)
C       Find the owner of the j^th column and store the column.
             owner = global_to_processor(j)
             local_index = global_to_local(j)
             if (mynode = owner) then
                  store the result as (local_index)^th column of C
             end if
        end do
        end
```

Figure 4: Translated code for in-core matrix multiplication

that C is distributed in column-block fashion, the compiler computes the owner of the resultant column which stores the result in the appropriate location in the local C array. Figure 4 shows the resultant node plus message passing program.

3.2.1 Comparison with a Hand-coded Program

Equation 1 can be rewritten as a sum of p partial sums as follows

$$c_j = \underbrace{\sum_{k=1}^{\lfloor \frac{n}{p} \rfloor} b_{kj} a_k + \sum_{k=\lfloor \frac{n}{p} \rfloor + 1}^{2\lfloor \frac{n}{p} \rfloor} b_{kj} a_k + \cdots + \sum_{k=((p-1)\times\lfloor \frac{n}{p} \rfloor)+1}^{n} b_{kj} a_k}_{p\ Sums}, \ j = 1:n, \ a_k \in \mathcal{R}^n \tag{2}$$

Each of these partial sums can be obtained on individual processors. Consider the partial sum $\sum_k b_{kj} a_k$. Each partial sum returns an intermediate vector in $\mathcal{R}^n$. Each vector is a linear combination of $\lfloor \frac{n}{p} \rfloor$ columns of A and $\lfloor \frac{n}{p} \rfloor$ elements of a column j of B. These intermediate vectors are then added to give the j^{th} column of matrix C. This process is repeated n times. It can be observed that to obtain the intermediate vectors, the best way to distribute A is in column-block form and B in row-block form. For this distribution, the number of rows of B in each processor is equal to the number of columns of A in that processor. Moreover, since in each step j of the summing process column c_j of array C is computed, the natural distribution for C is the same as that for A, namely column-block.

```
1 do j=1, columns_b (n)
2      do k=1, rows_b (n/p)
3           do i=1, rows_a (n)
4                temp(i) = b(k,j)×a(i,k) + temp(i)        ! Find Partial Sum
5           end do
6      end do
7      temp_sum = global sum of temp.
8      if (mynode is owner of column j) then
9           store temp_sum as column c(j'), where j' = global_to_local(j)
10     end if
11 end do
```

Figure 5: Hand-coded Distributed GAXPY Program

Figure 5 shows a hand-coded distributed memory GAXPY matrix multiplication program. The outer-most loop (j) varies from 1 to columns_b (n). In each iteration (j), the column j of array B is used for computation. Two inner loops multiply the k^{th} column of A by the k^{th} element of column j of B (lines 2-6). The intermediate vector *temp* is then added by all processors to give the global sum (temp_sum in line 7). Using the global index j, the owner of column $c(j)$ is calculated. This processor stores temp_sum as the j^{th} column of array C in the corresponding local array position. Note that the two inner loops operate in the local index space whereas the outer loop operates in the global index space. Figure 6 illustrates the computation in the j^{th} iteration of the algorithm. The elements of array B and the corresponding columns of array A are shown using the same shade.

A comparison of the programs in Figures 4 and 5 shows that the code generated by the in-core compiler is similar to the hand-coded version. That is, in-core compilation produces a good code in comparison with a hand-coded program.

3.3 Out-of-core Compilation

The out-of-core HPF compiler follows an approach similar to the in-core HPF compiler. In order to translate out-of-core programs, in addition to following the steps used for in-core compilation, the compiler also has to schedule explicit I/O accesses to fetch/store appropriate data from/to disks. The compiler has to take into account the data distribution on disks, the number of disks used for storing data and the prefetching/caching strategies used. As stated earlier, the local array of each processor is stored in a local array file (LAF). The portion of the local array currently required for computation is fetched from disk into the in-core local array (ICLA). The size of the ICLA is specified at compile time and usually depends on the amount of memory available. The larger the ICLA the better, as it reduces the number of disk accesses. Each processor performs computation on the data in its ICLA.

Some of the issues in out-of-core compilation are similar to compiler optimizations carried out to gain advantage of processor caches or pipelines. This optimization, commonly known as *stripmining* [Wol89a, ZC91], sections the loop iterations so that data of a fixed size (equal to cache size or pipeline stages) could be operated on in each iteration. In the case of out-of-core programs, the computation involving the entire local array is performed in stages, where each stage operates

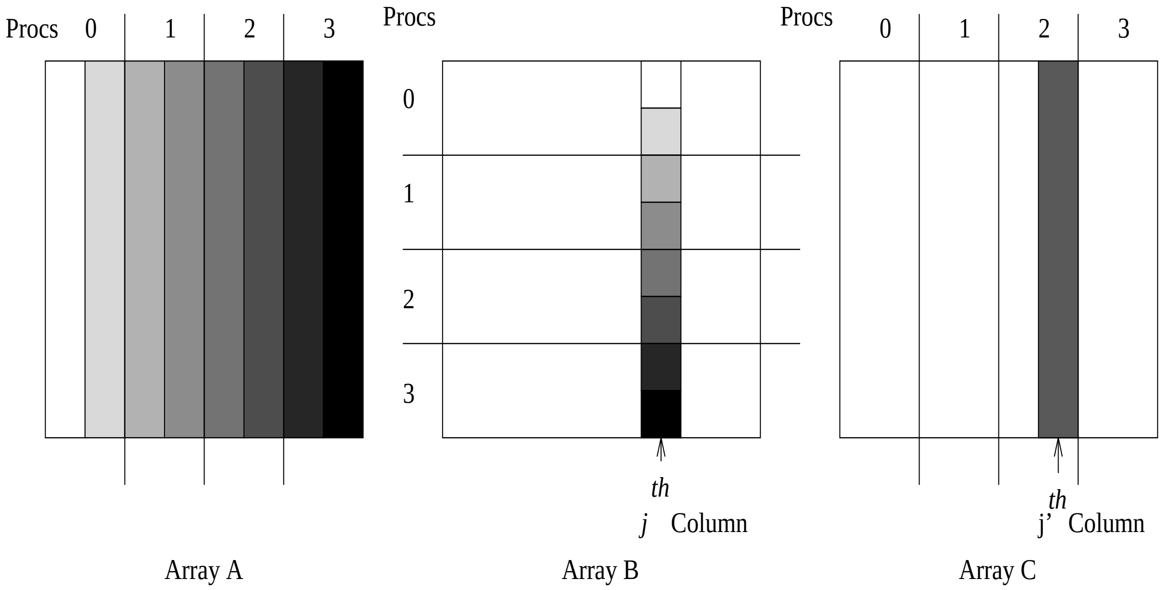

Figure 6: Distributed GAXPY for in-core matrices

on a different part of the array called a *slab*. The size of each slab is equal to the size of the ICLA. As a result, during the compilation, the iteration space of a `FORALL` statement is sectioned (*stripmined*) so that each iteration operates on the data that can fit in the processor's memory (ie. the size of ICLA).

Figure 7 shows the various steps involved in translating an out-of-core HPF program. The compilation consists of two phases. In the first phase, called the *in-core phase*, the arrays in the source HPF program are partitioned according to the distribution information (provided by the HPF directives) and local lower/upper bounds for each local array are calculated. Array expressions are then analyzed for detecting communication. In other words, the compilation in this phase proceeds in the manner compilation is done for in-core programs. The second phase, called the *out-of-core phase*, involves adding appropriate statements to perform I/O and communication. The local array is first stripmined according to the memory available in each processor. The resulting *slabs* are analyzed for communication. The local `FORALL` loop is then sequentialized and the loops are modified to insert necessary I/O calls. Note that I/O is performed in the local name space.

3.3.1 Compiling the Out-of-core Matrix Multiplication Program

We illustrate the out-of-core compilation process using the HPF matrix multiplication example given in Figure 3 and assume that the arrays A, B and C are out-of-core. Arrays A and C are distributed in column-block fashion over p processors, whereas array B is distributed in row-block fashion. Figure 8 shows the global arrays and their local array files. Figure 8(A) shows array A, Figure 8(B) shows array B and Figure 8(C) shows array C. Consider processor 3 in Figures 8(A) and 8(C). Figure 8(D) shows the local array corresponding to either array A or C and the corresponding OCLA. The OCLA of processor 3 is divided into *slabs*, each of which is equal to the size of the in-core local array (ICLA). The slabs are shown with different shades. Figure 8(E) shows the local array corresponding to array B for processor 3. The OCLA is divided into two slabs where each slab contains four subcolumns.

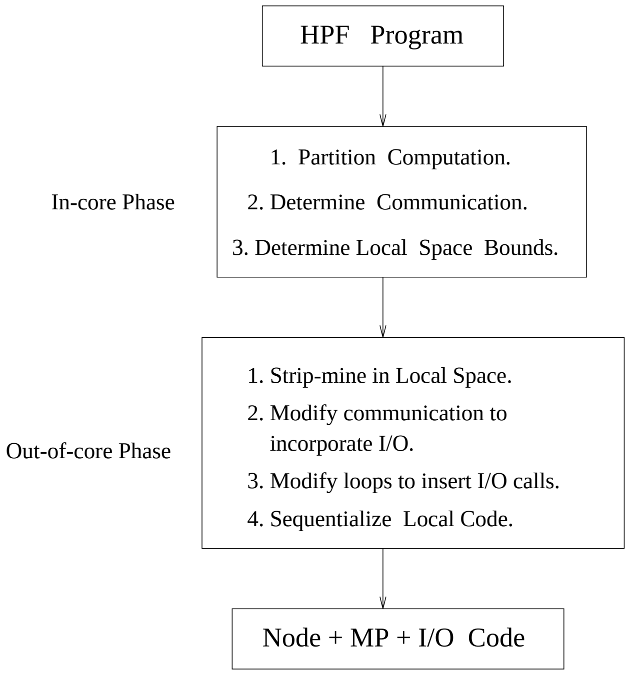

Figure 7: Flow chart for out-of-core compilation

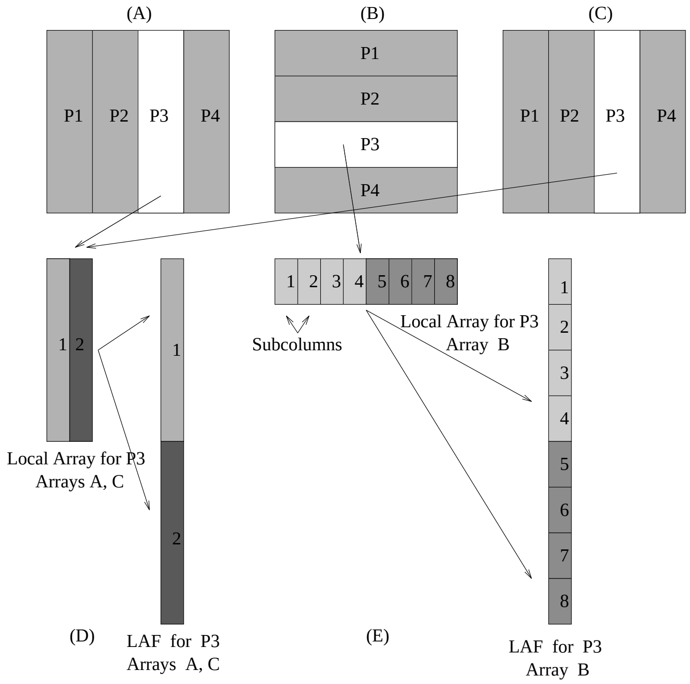

Figure 8: Compiling the Out-of-core Matrix Multiplication Program

The compilation is performed in two phases. In the in-core phase, the compiler obtains the necessary information about the arrays from the HPF directives (lines 3–7, Figure 3). Using this information, the compiler analyzes the array operations (lines 8–13, Figure 3). The compiler analyzes the outer loop (line 8, Figure 3), finds that the loop is a sequential DO loop and hence does not partition it. The inner loop consists of a `FORALL` construct which is parallelized. Using the array distribution information, the compiler computes the local array bounds and partitions the computation. The compiler analyzes the array assignment statement in line 10 to determine if communication is required. In this case, no communication is required in the innermost loop, but the outer loop requires a global sum operation.

In the second phase, stripmining of the index spaces is carried out using the memory (ICLA) size. Since the outer loop successively fetches elements of B, an I/O routine for fetching the slabs of array B is inserted. The inner loop is also stripmined and another I/O routine is inserted for fetching the slabs of array A. After the execution of the inner loop, all processors add their temporary results to obtain the corresponding columns of C. Using the distribution information of array C and the value of the outer loop index (j), the index of the processor that owns these columns is computed. This processor computes the local indices of these columns and stores the columns in the local array file. The resulting node program with the communication and I/O calls is shown in Figure 9.

3.4 Experimental Results

Figure 10 shows the performance for multiplication of 1K$\times$1K real arrays on 4, 16, 32 and 64 processors. The slab ratio, which is the ratio of the slab size to the out-of-core local array size is varied from (1/8) to 1. The slab-size for array A is chosen to be equal to the slab-size for array B. Note that the case when the slab-size is equal to the OCLA size (slab ratio = 1) is different from the case when the entire data is stored in main memory. When the slab-size equals the size of the OCLA, the slab-ratio (k in Figure 9) is 1. Even so, data is still accessed from disk, but only once for each column of C. We observe that as the slab ratio is decreased, the time taken increases. This is because a lower slab ratio means a smaller slab size and more number of slabs. This increases the number of I/O requests, though the total amount of data fetched from disk remains the same. The larger number of I/O requests increases the time taken for I/O which results in higher overall execution time. In the next section, we present optimizations which reduce the I/O time significantly.

4 Data Access Reorganization

For in-core programs, interprocessor communication is often the bottleneck which can degrade the overall performance considerably. Hence, an important optimization to be performed by any compiler for in-core programs is to minimize the communication overhead. This is usually done by aggregating small messages into a single long message so as to reduce communication latency, using collective communication routines etc. For an out-of-core compiler, it is very important to minimize the I/O cost because the time required to fetch data from disk is at least an order of magnitude more than the time required to communicate data between processors.

We measure I/O cost in terms of two metrics, namely the number of I/O requests per processor and the total amount of data fetched from disk per processor. Since the cost associated with physically accessing data (e.g. seek time, latency time etc.) is dictated by the hardware and to a

```
C       (N,N) Arrays distributed over p processors.
C       Stripmine code based on the slab size M.
C       Repeat operation k times, k=(no. of cols. in OCLA of A/no. of cols. in slab) = N^2/(M P)
C       Initialize global index.
          global_index=0
          do l=1, k
               Call I/O routine to read the ICLA of array B.
               do m=1, no_columns_in_icla_of_B
                    global_index=global_index+1
                    column_count = 0
                    do n=1, k
                         Call I/O routine to read the ICLA of array A.
                         do i=1, no_columns_in_icla_of_A {M/N}
                              column_count = column_count + 1
                              do j=1, no_elements_per_column {N}
                                   temp(j,i) = temp(j,i) + A(j,i)×B(column_count,m)
                              end do
                         end do
                    end do
                    Call Global Sum routine to obtain the (global_index)^th column of C
                    if (mynode is owner of this column) then
                         Store the column in the corresponding ICLA.
                              if ICLA is full then
                                   Call I/O routine to write the ICLA of array C.
                              end if
                    end if
               end do
          end do
```

Figure 9: Node+MP+I/O Pseudo-Code for the Matrix Multiplication Program

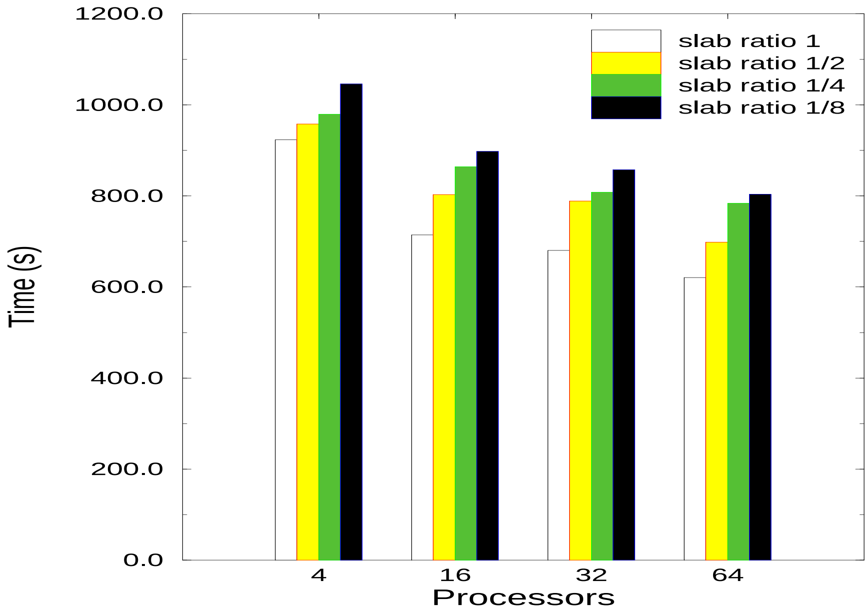

Figure 10: Effect of slab size variation

certain extent by the parallel file system, these two metrics can be used to effectively analyze the I/O costs associated with a user program.

4.1 I/O Cost Estimation

The previous section presented a simple extension of the in-core compilation method to compile out-of-core programs. Specifically, the extension of the in-core compilation technique did not explicitly consider the I/O costs associated with array assignment statements involving out-of-core arrays. In this section, we describe a framework for estimating the I/O costs in such statements and using this estimate to determine better access patterns which reduce the I/O cost.

In order to estimate the I/O cost associated with the compiled code, we analyze the node+MP+I/O program generated by the out-of-core compiler, which was described in the previous section (Figure 9). Using the local loop bounds, slab sizes and index variables for each out-of-core array, we compute the number of I/O accesses and the total amount of data accessed. We illustrate this by computing the dominant I/O costs in the out-of-core program in Figure 9. We call this version of the translated program as the *column slab version* because the out-of-core local array is divided into slabs along columns as shown in Figure 11(I). Note that for each column of B, the entire local array of A is required. Thus the dominant I/O cost is given by I/O accesses associated with the array A. Further, note that arrays B and C are accessed once during the entire computation, one slab at a time.

The I/O cost for accessing array A in the column slab version can be calculated as follows. Let

N = number of rows and columns in the global array A,

P = number of processors,

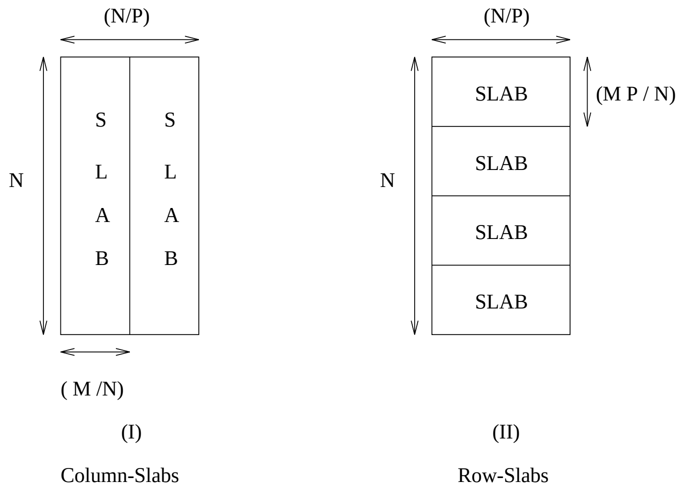

Figure 11: Column slabs versus row slabs

M = number of array elements in a slab (This depends on the available node memory).

To calculate one column of C, all columns of A are required. Therefore, the number of I/O requests per processor for one column of C is $(N/P)/(M/N) = N^2/(P\,M)$. The number of I/O requests per processor to calculate all columns of C is given by

$$T_{fetch}(A) = N^3/(M\,P) \tag{3}$$

The number of elements of array A that are fetched by each processor to compute one column of C is $N(N/P) = N^2/P$. Hence, the total number of elements of A that are accessed by each processor to compute all columns of C is given by

$$T_{data}(A) = N^3/P \tag{4}$$

Clearly, the I/O cost associated with this code is as large as the amount of computation. As explained in the previous section, it is important to note that the in-core version of this compiled code is as optimized as the hand-coded version.

In this example, another way of accessing the array A is to create slabs along rows as shown in Figure 11(II). This would require reordering the loops as illustrated in Figure 12. We call this version of the translated program as the *row slab version*. The I/O cost associated with this version can be calculated as follows. We note that a row slab of array A consists of all the columns of A within a particular set of rows. Thus when a processor fetches a slab, it has all the subcolumns of the out-of-core local array and the size of the subcolumns is the same as the number of rows in the slab. Since each slab has all the subcolumns necessary to calculate one subcolumn of C, each slab needs to be fetched only once to compute all the columns of C. Hence, in the row slab version, the number of I/O requests per processor is equal to the number of slabs which can be calculated as

$$T_{fetch}(A) = N/((M\,P)/N) = N^2/(M\,P) \tag{5}$$

Since the array A is accessed only once, the number of data elements of A fetched per processor to compute all columns of C is given by

$$T_{data}(A) = N(N/P) = N^2/P \tag{6}$$

```
C       (N,N) Arrays distributed over p processors.
C       Stripmine code based on the slab size M.
C       Repeat Operation k times, k=(no. of rows in OCLA of A/no. of rows in slab)=N^2/(M P)
          do l=1, k
              Call I/O routine to read the ICLA of array A.
              global_index=0
              do n=1, k
                  Call I/O routine to read the ICLA of array B.
                  do m=1, no_columns_in_icla_of_B
                      global_index=global_index+1
                      do i=1, no_columns_in_icla_of_A {N/P}
                          do j=1, no_elements_per_column {(M P)/N}
                              temp(j,i) = temp(j,i) + A(j,i)×B(i,m)
                          end do
                      end do
                      Call Global Sum intrinsic to obtain the (global_index)^th subcolumn of C
                      if (mynode is owner of this subcolumn) then
                          Store the subcolumn in the corresponding ICLA.
                          if ICLA is full then
                              Call I/O routine to write the ICLA of array C.
                          end if
                      end if
                  end do
              end do
          end do
```

Figure 12: Row Slab Version of the translated code

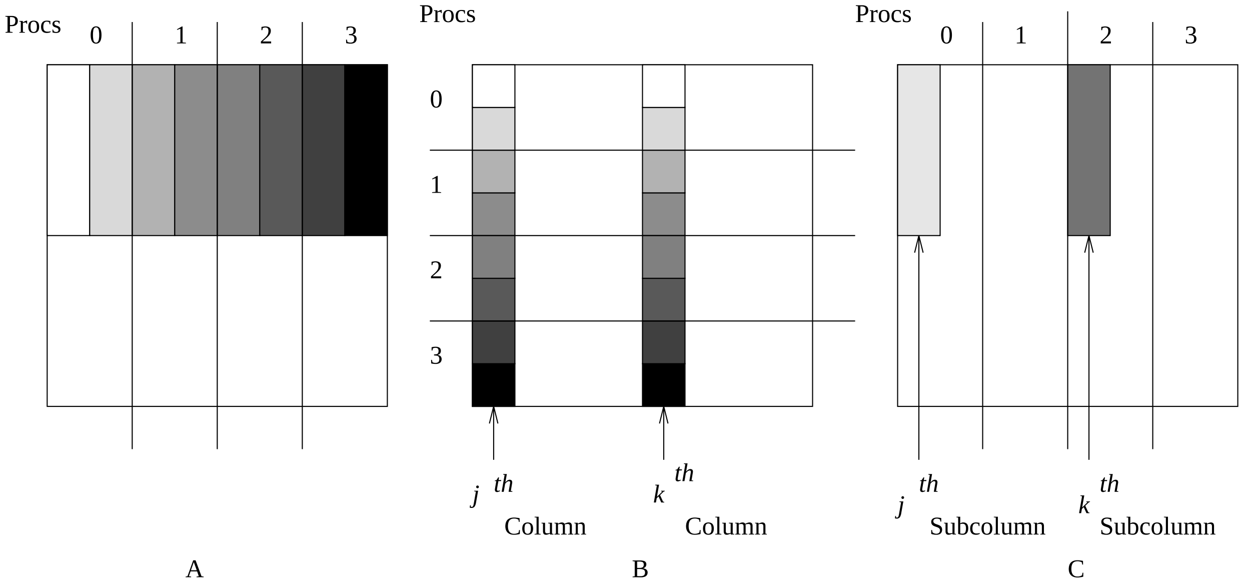

Figure 13: The figure shows the row slab version of the translated code. The subcolumns of array A and the corresponding elements of array B are shown using same shade.

Comparing equations 3 and 4 with equations 5 and 6, we observe that the row slab version requires an order of magnitude less number of I/O requests and an order of magnitude less amount of data to be fetched from disk than the column slab version. The I/O costs associated with arrays B and C are the same in both versions. So, the row slab version is clearly the method of choice for translating the out-of-core GAXPY matrix multiplication program. This data access pattern and the corresponding computation reorganization is further illustrated using Figure 13. The elements of array B that are multiplied to the corresponding columns of array A are shown using the same shade. The same slab of array A is multiplied with the j^{th} column to k^{th} column of B to produce the j^{th} to k^{th} subcolumns of C. Thus, the repeated I/O accesses to array A in the column slab version are eliminated.

In general, this approach for estimating the I/O cost requires analyzing the storage and access patterns along each dimension of the distributed out-of-core array. Based on this analysis, the loops are reorganized and the corresponding I/O costs are computed. The version with the minimum I/O cost is selected. This is summarized in the algorithm given in Figure 14.

4.2 Performance Results

Table 1 compares the performance of the row and column slab versions of the out-of-core matrix multiplication program. Two arrays of 1K×1K real numbers are multiplied on 4, 16, 32 and 64 processors. The slab ratio, which is the ratio of the slab size to the OCLA size, is varied from (1/8) to 1. We have also measured the time for an in-core version of the program which requires only an initial read of the arrays from disk. As explained earlier, this is different from the case when the slab size is 1 because in the latter case, the array is assumed to be out-of-core even though the entire out-of-core local array is stored in one slab. This slab is fetched from disk whenever necessary.

We observe that the row slab version performs considerably better than the column slab version

```
Determine the amount of available memory.
For each array used in the array assignment statement do
    For each dimension of the out-of-core array do
        Use index variables to analyze access patterns.
        Compute the I/O costs for stripmining using slabs along this dimension.
    end for
end for
Determine which array requires the largest amount of I/O.
Select the stripmining strategy which results in lowest I/O cost for that array.
```

Figure 14: General Algorithm for I/O Cost Estimation

Table 1: Performance of matrix multiplication for various slab sizes, time in seconds

Slab Ratio	4 Procs		16 Procs		32 Procs		64 Procs	
	Col. slab	Row slab	Col. slab	Row slab	Col. slab	Row slab	Col. slab	Row slab
1/8	1045.84	239.97	897.59	161.02	857.62	97.08	803.57	90.29
1/4	979.20	226.08	864.08	118.20	807.99	92.43	783.79	75.56
1/2	958.17	205.91	802.69	96.79	788.47	80.45	698.29	66.70
1	923.11	194.15	714.15	84.77	680.40	66.94	620.70	60.11
In-core	140.91		40.40		20.14		9.58	

Table 2: Performance of the row slab version for different slab sizes of arrays A and B

2K × 2K arrays, 16 processors, time in seconds

Slab_B size	Slab_A=256 Time (s)	Slab_A size	slab_B=256 Time (s)	Total Memory (Slab_A + Slab_B)
256	826.94	256	826.94	512
512	548.13	512	510.02	768
1024	507.01	1024	492.87	1280
2048	493.04	2048	452.29	2304

for any number of processors and any slab size. This is because it requires an order of magnitude less amount of I/O, as proved earlier. In both versions, the time taken increases as the slab ratio (or slab size) is decreased. A smaller slab size results in higher number of I/O requests, which increases the I/O cost. The difference between the in-core version and any of the out-of-core versions shows the corresponding time spent in performing I/O.

4.2.1 Selecting Slab Sizes for Multiple Arrays

The compiler has to choose the slab sizes to be used for all arrays in the program depending on the amount of available memory. One approach is to distribute the available memory equally among all the arrays, so that they all have the same slab size. Another approach is to analyze the I/O access patterns of the arrays and assign a larger slab size to the array with more frequent accesses. We have studied the effect of different slab sizes on the overall performance. Table 2 shows the performance of the matrix multiplication program for different slab sizes for arrays A and B. The arrays are chosen to be of size 2K×2K. In the first experiment, the slab size for array A is fixed and the slab size for array B is varied. The second experiment is performed by keeping the slab size for array B fixed. While in the first experiment, the execution time improves from 826.94 seconds to 493.04 seconds, in the second experiment the best performance observed is 452.29 seconds. Hence, instead of equally dividing the available memory between the slabs of A and B, if a larger portion is allocated to the slab of A, better performance is obtained. This is because A is accessed more often than B or C. Hence the compiler should allocate more memory to array A. As explained earlier, using the loop bounds and index variables, the compiler can determine which array requires more I/O accesses and accordingly allocate the available memory.

5 Related Work

Abu-Sufah first investigated strategies for improving performance of fortran programs in virtual memory environment [ASKL81]. Compiler transformations such as tiling, strip-mining, loop interchange, loop skewing are proposed by Wolfe [Wol89b]. Transformations like Unroll-and-jam and Scalar replacement are proposed by Carr [Car93]. Callahan studies the problem of register allocation [CCK90]. The notion *reference window* is proposed by Gannon et al. [GJG88]. The reference window is used by the compiler to study reuse in the program and perform corresponding transformations. Irigoin and Triolet also propose transformations to improve locality [IT88].

Schriber and Dongarra describe strategies to perform automatic tiling. They propose a linear algebraic formulation to find optimal size and shape of the data tile [SD90]. Ramanujam and Sadayappan use locality transformations to minimize interprocessor communication. They use a linear programming formulation to obtain optimal shape of the tile [RS90]. Further studies in *optimal* tiling are done by Boulet et al. [BDRR93]. Wolf and Lam propose an elegant loop transformation theory to improve locality and parallelism [WL91]. Blockability in numerical algorithms has also been studied extensively [CK92, GJMS88].

An excellent description of the compiler transformations is given in [BGS93].

Most of this work is targeted for locality optimizations for sequential programs. Our work deals with locality optimizations for out-of-core programs running on distributed memory machines. We also assume that the out-of-core programs access the data from files which are distributed over multiple disks.

6 Conclusions

We have described how an out-of-core program written in a data parallel language like HPF can be translated into a message passing node program with explicit communication and parallel I/O. Such a compiler is necessary for compiling large scale scientific applications written in a data parallel language. These applications typically handle large quantities of data which results in the program being out-of-core.

We have discussed the basic compilation model and various steps involved in the compilation. The compilation process was illustrated using an out-of-core matrix multiplication example. We described how the basic in-core compilation method can be extended to compile out-of-core programs. However, the code generated this way may not give good performance. We have proposed an optimization by which the compiler can improve the code generated by the above method. The compiler estimates the I/O costs associated with different array access patterns and selects the method with the least I/O cost. This can reduce the amount of I/O by as much as an order of magnitude. We also discussed how the performance of the program varies with slab size. Instead of dividing the available memory equally among all arrays, the best performance is obtained when the most frequently accessed array is allocated a larger slab size.

Some of the compilation techniques described in this paper are currently being done by hand. We are in the process of implementing them in the compiler. Due to stability problems with the hardware, we have done experiments with relatively small data sets.

Acknowledgments

We would like to thank Ken Kennedy and Chuck Koelbel for many enlightening discussions. We would also like to thank our compiler group at Syracuse University for their help with the basic infrastructure of the in-core HPF compiler.

References

[ASKL81] W. Abu-Sufah, David Kuck, and D. Lawrie. On the performance enhancement of paging systems through program analysis and program transformation. *IEEE Transactions on Computers*, C-30(5):341–356, May 1981.

[BBG+93] F. Bodin, P. Beckman, D. Gannon, S. Narayana, and S. Yang. Distributed pC++: Basic Ideas for an Object Parallel Language. In *Proceedings of the First Annual Object-Oriented Numerics Conference*, pages 1–24, April 1993.

[BCF+93] Z. Bozkus, A. Choudhary, G. Fox, T. Haupt, and S. Ranka. Fortran 90D/HPF Compiler for Distributed Memory MIMD Computers: Design, Implementation, and Performance Results. In *Proceedings of Supercomputing '93*, pages 351–360, November 1993.

[BDRR93] Pierre Boulet, Alain Darte, Tanguy Risset, and Yves Robert. (Pen)-ultimate tiling? Technical Report 93-36, Ecole Normale Superieure de Lyon, 46, Alle'e d'Italie, 69364 Lyon Cedex 07, France, November 1993.

[BGS93] David F. Bacon, Susan Graham, and Oliver J. Sharp. Compiler Transformations for High-Performance Computing. Technical Report UCB/CSD-93-781, Computer Science Division, University of California, Berkeley, Computer Science Division, University of California, Berkeleyley, California 94720, 1993.

[Car93] Steve Carr. *Memory-Hierarchy Management.* PhD thesis, Rice University, February 1993.

[CCK90] David Callahan, Steve Carr, and Ken Kennedy. Improving register allocation for subscripted variables. *Proc. of SIGPLAN'90 Conference on Program Language Design and Implementation*, June 1990.

[CFPB93] P. Corbett, D. Feitelson, J. Prost, and S. Baylor. Parallel Access to Files in the Vesta File System. In *Proceedings of Supercomputing '93*, pages 472–481, November 1993.

[CK92] Steve Carr and Ken Kennedy. Compiler Blockabilty of Numerical Algorithms. *Proc. of Supercomputing'92*, November 1992.

[CMZ92] B. Chapman, P. Mehrotra, and H. Zima. Programming in Vienna Fortran. *Scientific Programming, Vol.1, No.1*, 1992.

[CR93] High Performance Computing and Communications: Grand Challenges 1993 Report. A Report by the Committee on Physical, Mathematical and Engineering Sciences, Federal Coordinating Council for Science, Engineering and Technology, 1993.

[dRC94] J. del Rosario and A. Choudhary. High Performance I/O for Parallel Computers: Problems and Prospects. *IEEE Computer*, March 1994.

[FHK+90] G. Fox, S. Hiranandani, K. Kennedy, C. Koelbel, U. Kremer, and C. Tseng. Fortran D Language Specifications. Technical Report COMP TR90-141, Rice University, 1990.

[For93] High Performance Fortran Forum. *High Performance Fortran Language Specification.* Version 1.0, May 1993.

[Fox91] Geoffrey Fox. The architecture of problems and portable parallel software systems. Technical Report SCCS-78b, Northeast Parallel Architectures Center, Syracuse University, Syracuse, NY 13244, 1991.

[GJG88] D. Gannon, W. Jalby, and K. Gallivan. Strategies for the Cache and Local Memory Management by global program transformation. *Journal of Parallel and Distributed Computing*, 5(5):587–616, October 1988.

[GJMS88] K. Gallivan, W. Jalby, U. Meier, and A. Sameh. The impact of hierarchical memory systems on linear algebra algorithm design. *Intl. Journal of Supercomputing Applications*, 2(1):14–48, 1988.

[IT88] Francois Irigoin and Remi Triolet. Supernode Partitioning. *POPL'88, Fifteenth Annual ACM Symposium on Principles of Programming Languages*, January 1988.

[KLS+94] C. Koelbel, D. Loveman, R. Schreiber, G. Steele, and M. Zosel. *High Performance Fortran Handbook.* The MIT Press, 1994.

[RS90] J. Ramanujam and P. Sadayappan. Tiling of Iteration Spaces for Multicomputers. *International Conference on Parallel Processing*, pages II179–II186, August 1990.

[SD90] R. Schriber and J. Dongarra. Automatic Blocking of Nested Loops. Technical report, Research Institute for Advanced Computer Science, May 1990.

[Sni92] Marc Snir. Proposal for IO. Posted to HPFF I/O Forum by Marc Snir, July 7 1992.

[TBC94a] R. Thakur, R. Bordawekar, and A. Choudhary. Compiler and Runtime Support for Out-of-Core HPF Programs. In *Proceedings of the* 8^{th} *ACM International Conference on Supercomputing*, pages 382–391, July 1994.

[TBC+94b] R. Thakur, R. Bordawekar, A. Choudhary, R. Ponnusamy, and T. Singh. PASSION Runtime Library for Parallel I/O. In *Proceedings of the Scalable Parallel Libraries Conference*, October 1994.

[WL91] Michael E. Wolf and Monica S. Lam. A Loop Transformation Theory and An Algorithm to Maximize Parallelism. *IEEE Transactions on Parallel and Distributed Systems*, 2(4):452–471, October 1991.

[Wol89a] M. Wolfe. *Optimizing Supercompilers for Supercomputers.* The MIT Press, Cambridge, MA, 1989.

[Wol89b] M. J. Wolfe. More iteration space tiling. *Proceedings of Supercomputing'89*, November 1989.

[ZBC+92] H. Zima, P. Brezany, B. Chapman, P. Mehrotra, and A. Schwald. Vienna Fortran - a language specification. Technical Report ICASE Interim Report 21, MS 132c, ICASE, NASA, Hampton VA 23681, 1992.

[ZC91] H. Zima and B. Chapman. *Supercompilers for Parallel and Vector Computers.* ACM Press, New York, NY, 1991.