Generating Communication for Array Statements: Design, Implementation, and Evaluation

James M. Stichnoth, David O'Hallaron, and Thomas R. Gross
School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213

Supported in part by the Defense Advanced Research Projects Agency, Information Science and Technology Office, under the title "Research on Parallel Computing," ARPA Order No. 7330. Work furnished in connection with this research is provided under prime contract MDA972-90-C-0035 issued by DARPA/CMO to Carnegie Mellon University. Supported also in part by the Air Force Office of Scientific Research under Contract F49620–92–J–0131. This material is based upon work supported in part under a National Science Foundation Graduate Research Fellowship.

Abstract

Array statements as included in Fortran 90 or High Performance Fortran (HPF) are a well-accepted way to specify data parallelism in programs. When generating code for such a data parallel program for a private memory parallel system, the compiler must determine when array elements must be moved from one processor to another. This paper describes a practical method to compute the set of array elements that are to be moved; it covers all the distributions that are included in HPF: block, cyclic, and block-cyclic. This method is the foundation for an efficient protocol for modern private memory parallel systems: for each block of data to be sent, the sender processor computes the local address in the receiver's address space, and the address is then transmitted together with the data. This strategy increases the communication load but reduces the overhead on the receiving processor. We implemented this optimization in an experimental Fortran compiler, and this paper reports an empirical evaluation on a 64-node private memory iWarp system, using a number of different distributions.

Running head:

Communication for Array Statements

Contact author:

Thomas Gross
School of Computer Science
Carnegie Mellon University
5000 Forbes Ave.
Pittsburgh, PA 15213-3891

412 268 7661

Symbols

$\mathcal{F}$

$\lfloor$

$\rfloor$

$\leq$

$\oplus$

$\cap$

∞

$\times$

$\emptyset$

$\equiv$

$\bigcup$

Ψ

$\in$

1 Introduction

One approach to programming parallel systems that has received considerable attention is to let the programmer explicitly expose the data parallelism in a program and to let the compiler deal with the details of mapping the program onto the target parallel system. Array statements are an attractive linguistic construct to capture the data parallelism in a program; they succinctly describe computations for a collection of data elements. Array statements have been added to earlier versions of Fortran (e.g., CM Fortran), are part of the standard Fortran (i.e., Fortran 90), and can be found in various extensions (e.g., Fortran D, High Performance Fortran (HPF), etc.). Conceptually similar constructs are also included in other languages [1], but we concentrate in this paper on the kind of array statement as defined in Fortran 90 or HPF.

Consider an array statement such as:

$$\mathtt{A}(f_1 : l_1 : s_1) = \mathcal{F}(\mathtt{B}(f_2 : l_2 : s_2)). \tag{1}$$

$\mathtt{A}(f : l : s)$ defines a sequence of array elements $\mathtt{A}(f + i \times s), 0 \le i \le \lfloor (l - f)/s \rfloor$. s is called the *access stride*, and the sequence of array elements is sometimes referred to as a *slice* of the array. $\mathcal{F}$ is an element-wise function; i.e., the result of $\mathcal{F}(\mathtt{B(i)})$ is stored in $\mathtt{A(j)}$ for corresponding values of `i` and `j`. The semantics of the assignment statement require that the operands on the right-hand side are read before any of the results are written. A compiler that is to translate the above array statement for execution on a private memory parallel system must map the global name space of the data parallel program into the memory organization provided by the parallel system. The compiler must determine how the elements of different arrays are stored relative to each other (alignment), and how they are distributed onto the parallel processors (distribution). For each array element `A(i)` of an array `A`, the compiler must define two mapping functions: an *ownership* function that produces the processor at which `A(i)` resides (the *owner*), and a *local memory* function that determines the actual address of `A(i)` in the local memory space of the owner. Furthermore, the compiler must determine on which processor the computation of $\mathcal{F}$(`B(i)`) is to take place, for each index `i` as defined by the array statement.

The alignment and distribution information can be specified with directives by the user [8, 14] (such as with `ALIGN` and `DISTRIBUTE` directives in HPF and Fortran D), or can be derived by a compiler (e.g., [5]). A number of compilers base the decision on where to perform a computation on the "owner computes rule" [14]: the processor that stores the array element on the left-hand side of an array statement computes the update. Those operands on the left-hand side that are non-local must be moved, either by calls to a communication library or by direct I/O operations.

The problem of mapping data and computation onto a parallel system is a key issue for any compiler, and a number of research projects have reported various methods (see [14] for a recent review). This paper concentrates on the phase of a compiler for private memory parallel systems that deals with the generation of communication operations. Given alignment and distribution, we describe a practical machine-independent method to determine the array elements that must be copied from one processor to another. This phase is independent of how the compiler determines the alignment or distribution but has a high impact on the quality of the code produced. The method described here is the foundation for an efficient protocol for a wide range of private memory parallel systems. Some aspects of our work have also been discovered independently by other researchers (e.g., [4]) or are included in other compilers (e.g., the CM Fortran compiler [11]). However, this is the first unified description and evaluation of a general algorithm for block, cyclic, and block-cyclic distributions. Section 2 states the problem; Section 3 describes our approach and presents the basic protocol for private memory parallel systems. Section 4 discusses an optimization that is attractive for parallel systems with ample communication bandwidth. This optimization is implemented in an experimental Fortran compiler for a private memory parallel system. In Section 5 we report on performance measurements obtained on a 64-node iWarp system that demonstrate the benefits of this optimization.

2 Compiling array statements

The discussion in the rest of the paper assumes that the compiler uses the owner computes rule and focuses on the issues of the code generation phase.

The earlier example of an array statement included several simplifications. An array statement is any statement of the source language that allows an array `A` or an array slice such as `A`$(f : l : s)$ as a source or destination operand. In general, arrays can have more than one dimension, and there can be multiple array operands in a single statement.

Array *assignment* statements exhibit all the interesting problems of mapping array statements onto private memory parallel systems. Any solution to this class of problems can easily be extended to cover the compilation of most other data parallel constructs, such as the `WHERE` statement, the `DOALL` statement, or building the parameter list for a subroutine call. For this reason, we concentrate on the array assignment statement in this paper. The reader is referred to [12] for a full treatment of the compilation of this statement and other classes of array statements.

In this paper we describe only how to compile the simple array statement given by equation (1). If there are multiple terms on the right-hand side, we repeat the communication step for each term. If the terms contain multiple array slices, we treat each pair of corresponding slices orthogonally, nesting the resulting loops. Then the basic method to execute this statement conceptually consists of the following steps for each processor P_i [10]:

1. Send those elements of `B` that are owned by P_i but required by another processor P_j to this P_j, $P_j \in \{P_0, \ldots, P_{n-1}\}$.

2. Receive non-local elements of `B` required to update a local element of `A`$(f_1 : l_1 : s_1)$, and store these elements locally.

3. Execute the updates of the local elements of `A`$(f_1 : l_1 : s_1)$, using local elements `B`(j).

If $P_i = P_j$, no inter-processor communication is required. A compiler has the option to use the original elements of `B` in the computation in step 3, or to copy the local elements of `B` to the same space that holds the values received in step 2. In our experience, this latter option simplifies step 3 (and the resulting code) and is therefore implemented in our compiler. However, this decision is not central to the topic of this paper. There exist also other optimizations (e.g., detecting communication patterns [3, 12]), however; we concentrate here on the more fundamental task of producing efficient communication code when compiling the array statement, i.e. to determine which elements of B to send (and where to send them) and which elements to receive (and where in memory to place the received values).

3 Communication for the basic assignment statement

When executing an array statement, the actual communication pattern can be arbitrarily complex: anywhere from no communication at all, to a simple shift operation, to an all-to-all communication. For each processor, the compiler must determine the set of processors to which this processor must send data, and must correspondingly determine the set of processors from which to receive data. Regardless of the sophistication of the compiler in detecting communication patterns, the compiler must be able to handle the general problem of determining communication between an arbitrary pair of processors.

Given the core array assignment statement discussed previously, our goal is to calculate the *exact* set of array elements that need to be sent between any two processors. An alternative approach is to calculate a superset that includes all array elements required and leave it to the receiving processor to determine which elements are actually needed. The advantages and tradeoffs of such an overestimate are not clear, so we choose to perform an exact calculation.

In the most general case, the arrays A and B of $\mathtt{A}(f_1 : l_1 : s_1) = \mathtt{B}(f_2 : l_2 : s_2)$ are distributed across the processors in block-cyclic fashion.[1] We concentrate in this paper on the general block-cyclic distributions, since a code generator phase that can handle block-cyclic distributions of arrays can also handle the simpler cases of block and cyclic distributions. Figure 1 depicts a block-cyclic distribution; elements of the array that reside on one particular processor are shown shaded. This processor owns indices beginning with f, in blocks of size n, with *distribution stride* s, which is usually the product of the block size n and the number of processors over which the array dimension is distributed. (When we refer to *indices* we actually mean the corresponding *array elements*.) Because of the regular nature of the block-cyclic distribution, this index set, which describes the elements owned by a processor, can be specified using only four parameters: f, l, s, and n (l is the last element of the array).

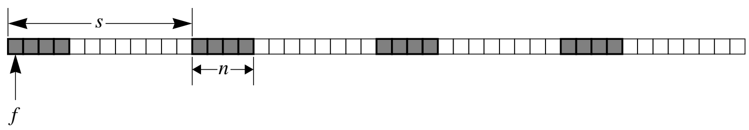

Figure 1: Block-cyclic distribution.

[1]For the code generation phase, it is not important that alignment and distribution of an array remain unchanged during the execution of a program. Re-alignment or re-distribution can be implemented by array assignment statements, where the right-hand side operand has the old alignment or distribution, and the left-hand side operand has the new one.

The key parts of the compilation of the array assignment are to determine which array elements to send, and to determine where in memory to store received elements. The basis of our solution is a practical way to compute the intersections of slices. Slices can then be used to determine the array elements involved in an array expression as well as to characterize block-cyclic ownership sets. Since slices can be characterized by three parameters (first element, last element, access stride), and four parameters can be used to characterize ownership sets in terms of slices (as shown in Section 3.2), the resulting intersections can also be specified using a small number of parameters.

3.1 Slice intersection

A slice parameterizes an arithmetic sequence of integers, and our goal is to find the intersection of two or more such slices, while preserving the ordering. As we show below, such an intersection either is empty or is another slice; thus slices are closed under intersection. Our goal is to find the *first*, *last*, and *stride* components of the intersection. Figure 2 shows two examples of slice intersections.

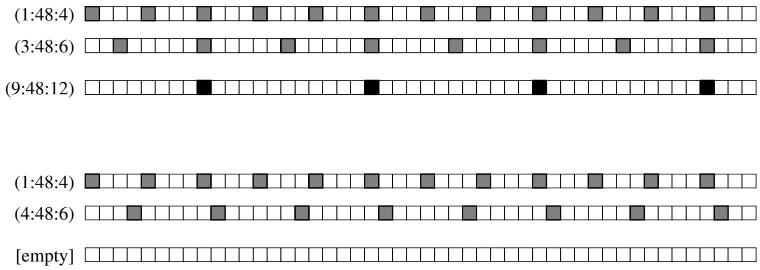

Figure 2: Two examples of slice intersection: $(1 : 48 : 4) \cap (3 : 48 : 6) = (9 : 48 : 12)$ and $(1 : 48 : 4) \cap (4 : 48 : 6) = \emptyset$.

Our approach for finding the intersection of two slices has three conceptual steps. The first step is to extend the lower bound and the upper bound of each sequence to $-\infty$ and $+\infty$, respectively, while remembering the original bounds. The second step is to find the intersection of these infinite sequences, which, as shown below, is either empty or another infinite sequence. The final step is to find the true lower and upper bounds of the intersection by using the remembered original bounds. This approach works well for finding the intersections of more than two slices, since we can avoid the steps of converting intermediate intersections from finite sequences to infinite sequences, and vice versa.

Consider the slice $(f : l : s)$. Recall that f represents the first element (with the smallest index) of the slice, l is an upper bound on elements in the set, and s is the stride. Assume that s is positive; if s is negative, then f is the element with the largest index, and l is a lower bound.

To capture an infinite sequence that contains $(f : l : s)$ as a subsequence, we notice that there exist an infinite number of integers r such that each element of $(f : l : s)$ can be described as $r + q \times s$ for *some* integer q. We therefore introduce a modified representation of a slice, $(f : l : s : r)$. This represents a slice with infinite bounds, as described above, where f and l are the original remembered bounds. The s parameter serves the same purpose as before. The r parameter is a "representative" of the set, in the sense that any member of the set is equal to $r + qs$ for some integer q. For example, $(0 : 10 : 2 : 37)$ represents the same sequence as $(1 : 10 : 2)$, which is also the same as $(1 : 9 : 2)$. In this representation, the f parameter need not be a representative of the set, in the same sense that the l parameter need not be a representative of the set. The only requirement is that for any integer q, $r + qs$ is a member of the original set if and only if $f \leq r + qs \leq l$ in the new representation.

The conversion between representations is simple. The sequence $(f : l : s)$ is the same as $(f : l : s : f)$, and the sequence $(f : l : s : r)$ is the same as $(f + ((r - f) \bmod s) : l : s)$.

Consider the intersection of the slices $S_i = (f_i : l_i : s_i : r_i)$ and $S_j = (f_j : l_j : s_j : r_j)$. A representative of the intersection is some integer r for which there exist integers m and n such that $r = r_i + ms_i = r_j + ns_j$. By solving this Diophantine equation ([6], Chapter 33), we find that a solution exists if and only if g evenly divides $(r_j - r_i)$, and that solution is given by $r = r_i + \frac{(r_j - r_i)x}{g} s_i$, where x and g are integers such that $xs_i + ys_j = g = \gcd(s_i, s_j)$. These integers can be computed using Euclid's extended GCD algorithm [9]. It is clear that the resulting stride of the intersection is $\mathrm{lcm}(s_i, s_j) = s_i s_j / g$. As our new upper and lower bounds, we must take the tightest of the two original bounds. Therefore:

$$(f_i : l_i : s_i : r_i) \cap (f_j : l_j : s_j : r_j) = (\max(f_i, f_j) : \min(l_i, l_j) : s_i s_j / g : r_i + (r_j - r_i) x s_i / g).$$

Note that the representative is not necessarily the same as either of the lower bounds. This freedom provides the motivation for using this modified slice representation. If we are taking the intersections of more than two slices, as is necessary when deriving the communication sets, it is more efficient to minimize the number of conversions between representations.

3.2 Block-cyclic sets

As illustrated by Figure 1, the set of array elements stored on a processor when the array is distributed in a block-cyclic fashion can be represented using four parameters: f, l, s, and n. The f parameter is the lowest-numbered array index owned; the l parameter is an upper bound on the highest-numbered array index owned; the n parameter is the block size; and the s parameter is the distance between the start of one block and the start of the next. We call this set of indices that is induced by the block-cyclic distribution of the array a *block-cyclic set*.

We showed above that slices are closed under intersection. Unfortunately, the same does not hold for block-cyclic sets, as illustrated in Figure 3. Thus to take intersections of block-cyclic sets, we need to develop a different representation. We choose to represent a block-cyclic set as a union of disjoint slices $\bigcup_{k=0}^{n-1} (f + k : l : s)$, as pictured in Figure 4.

With the above method to compute the intersection of slices, we are now ready to introduce the algorithm to determine the array elements that must be sent and received by a given processor. Our goal is to devise a code template that can be included in the code for each processor to determine the communication sets.

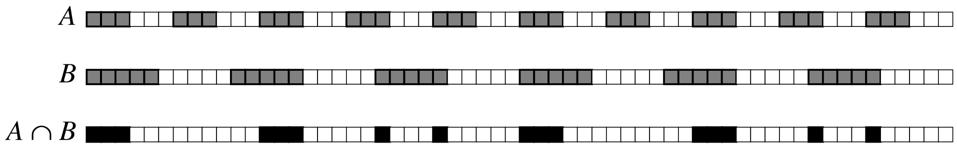

Figure 3: Block-cyclic sets are not closed under intersection.

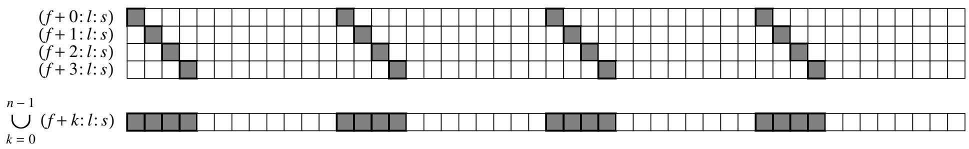

Figure 4: A block-cyclic distribution as a union of disjoint slices.

3.3 Communication: sender

The compiler must produce code for a given sending processor S and a receiving processor D such that S computes the elements of B that must be sent to D. Since A and B both have block-cyclic distributions, there exist f_3, l_3, n_3, s_3, f_4, l_4, n_4, and s_4 such that the portion of A that processor D owns, and the portion of B that processor S owns, are correspondingly given by:

$$Own_D(\mathtt{A}) = \bigcup_{j=0}^{n_3-1} (f_3 + j : l_3 : s_3), \qquad Own_S(\mathtt{B}) = \bigcup_{i=0}^{n_4-1} (f_4 + i : l_4 : s_4). \tag{2}$$

Note that the names of processors S and D never appear in the set definitions; they are implicitly encoded in the parameters f_3 and f_4, which can be computed at runtime. Also note that each ownership set is the union of disjoint sets, provided that $n_3 \leq s_3$ and $n_4 \leq s_4$.

The basic method of determining the communication is to take intersections of sets. All these sets can be described using either 3 or 4 parameters (3 in the case of a single slice, and 4 in the case of a block-cyclic set), although the resulting intersections are not nearly as simple to represent. Figure 5 gives an example of the communication that might arise between processors during the evaluation of the assignment statement $\mathtt{A}(1 : n : 1) = \mathtt{B}(1 : n : 1)$. Each array is distributed over two processors, where A has a block size of 3 and B has a block size of 5. The shaded boxes represent the indices of each array that are owned by the corresponding processor. The arrows represent the array elements that processor S must send and that processors d_1 and d_2 must receive; this corresponds to the intersection of the block-cyclic ownership sets. Even though the ownership sets for the sender and receiver are simple to represent in isolation, their intersection, although somewhat regular, is more complicated. The pattern becomes even more complicated when we use non-unit access strides in the assignment statement. This example also illustrates that there can be little similarity between the two intersections, even though $Own_{d_1}(\mathtt{A})$ is simply the complement of $Own_{d_2}(\mathtt{A})$.

The set of indices of A owned by processor D that are stored in the computation is $S_1 = Own_D(\mathtt{A}) \cap (f_1 : l_1 : s_1)$, and the set of indices of B owned by processor S that are read is

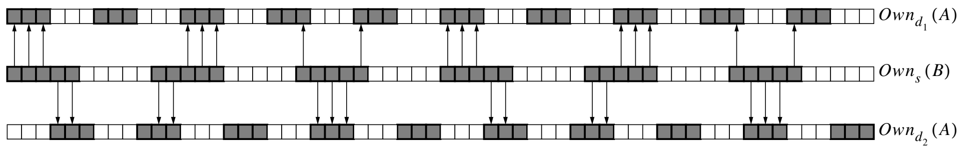

Figure 5: Communication sets that are difficult to parameterize.

$S_2 = Own_S(\mathtt{B}) \cap (f_2 : l_2 : s_2)$. For any integer i, if $f_1 + is_1 \in S_1$ and $f_2 + is_2 \in S_2$, then processor S sends $\mathtt{B}[f_2 + is_2]$ to D, which stores it in $\mathtt{A}[f_1 + is_1]$. To express this in terms of set intersections, we define a Map function, which associates $\mathtt{A}(f_1 + is_1)$ with $\mathtt{B}(f_2 + is_2) = \mathtt{B}(Map(f_1 + is_1))$, on an index y as $Map(y) = \frac{y - f_1}{s_1} \cdot s_2 + f_2$. The Map function is easily extended to a set or sequence in the obvious fashion. Hence for any integer i, if $Map(f_1 + is_1) = f_2 + is_2 \in Map(S_1)$ and $f_2 + is_2 \in S_2$, then processor S sends $\mathtt{B}[f_2 + is_2]$ to D. Thus the set of indices of $\mathtt{B}$ that processor S sends to processor D is the following set:

$$\Psi_S = S_2 \cap Map(S_1) = (Own_S(\mathtt{B}) \cap (f_2 : l_2 : s_2)) \cap Map(Own_D(\mathtt{A}) \cap (f_1 : l_1 : s_1)).$$

But since $Map(f_1 : l_1 : s_1) = (f_2 : l_2 : s_2)$, we can remove the $(f_2 : l_2 : s_2)$ term. Thus the intersection to compute is the following:

$$\Psi_S = Own_S(\mathtt{B}) \cap Map(Own_D(\mathtt{A}) \cap (f_1 : l_1 : s_1)). \tag{3}$$

The basic algorithm for computing the indices of the array elements to send is then the following, which can be translated directly into Fortran code:

```
DO j = 0, n3 - 1
   DO i = 0, n4 - 1
      (first, last, stride) = (f4 + i : l4 : s4) ∩ Map((f3 + j : l3 : s3) ∩ (f1 : l1 : s1))
```

This algorithm computes at most $2n_3n_4$ slice intersections. With a small change to the i and j loop bounds, fewer intersections might be required, since it is possible for i and j to loop over only those combinations for which the corresponding intersections are non-empty. In addition, since the strides and upper bounds of the slices are the same for all iterations of the loop, we can compute the resulting $last$ and $stride$ parameters outside the loop.

Local addresses

While the analysis above determines which global indices of the array must be sent, we also need to map each global index to the processor's local memory. This is the *local memory* function described previously.

To minimize the amount of wasted memory space, we want to compact the array (i.e., put the blocks contiguous in memory). [7] argues that this compaction is actually necessary, and should not be viewed merely as an optimization. Given a global index x and a block-cyclic distribution characterized as above in terms of f_i, l_i, n_i, and s_i, we find that the local index $LM_i(x)$ is defined as:

$$LM_i(x) = \overbrace{\left\lfloor \frac{x - f_i}{s_i} \right\rfloor}^{block} \cdot n_i + \overbrace{((x - f_i) \text{ mod } s_i)}^{offset} + \texttt{K}$$

where K is the index of the first array element. In the C programming language, K is 0; in Fortran K is usually 1. This mapping is illustrated in Figure 6.

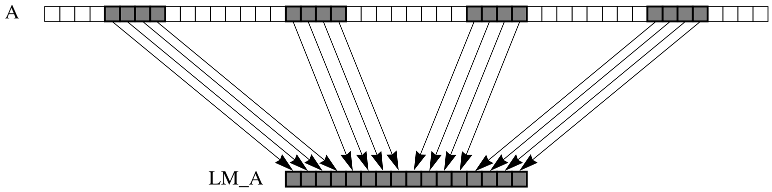

Figure 6: Local memory mapping function for a processor.

To calculate the local memory indices of the array elements to send, we simply apply the LM function to Ψ_S. Although the LM function seems expensive to apply to every element in Ψ_S, with a division, modulo, and multiplication operation, we can actually easily incorporate the LM function into the algorithm that computes Ψ_S, with little additional cost.

3.4 Communication: receiver

There are two problems that the compiler must solve for the receiving processor R. First it must determine which array elements R is supposed to receive. Second, it must anticipate the order in which the sending processor sends the array elements and associate each sent element with the desired receiving index. The first problem can be addressed by using similar analysis for computing Ψ_S, which yields:

$$\Psi_D = Map^{-1}(S_2) \cap S_1 = Map^{-1}(Own_S(\texttt{B}) \cap (f_2 : l_2 : s_2)) \cap (Own_D(\texttt{A}) \cap (f_1 : l_1 : s_1)).$$

But since $Map^{-1}(f_2 : l_2 : s_2) = (f_1 : l_1 : s_1)$, we can remove the $(f_1 : l_1 : s_1)$ term. Thus the intersection to compute is the following:

$$\Psi_D = Map^{-1}(Own_S(\texttt{B}) \cap (f_2 : l_2 : s_2)) \cap Own_D(\texttt{A}).$$

To address the second problem, we simply ensure that we use the same looping structure as in the computation of Ψ_S. The corresponding code template is:

```
DO j = 0, n_3 - 1
   DO i = 0, n_4 - 1
      (first, last, stride) = Map^-1((f_4 + i : l_4 : s_4) ∩ (f_2 : l_2 : s_2)) ∩ (f_3 + j : l_3 : s_3)
```

```
C Fill buffer, send to _dest_          | C Receive from processor _from_, unpack
  bufptr = 1                           |   CALL receive(buffer,from)
  DO loop0 = MAX(first_A(dest,1),1),   |   bufptr = 1
$        MIN(first_A(dest,1)+15,1024)  |   DO loop0 = 1, 16
    DO loop1 = 1, 16                   |     DO loop1 = MAX(first_B(from,1),1),
      buffer(bufptr) = B(loop1,loop0)  | $         MIN(first_B(from,1)+15,1024)
      bufptr = bufptr + 1              |       A(loop0,loop1) = buffer(bufptr)
    ENDDO                              |       bufptr = bufptr + 1
  ENDDO                                |     ENDDO
  IF (bufptr.GT.1)                     |   ENDDO
$   CALL send(dest,buffer,(bufptr-1))  |
```

Figure 7: Fortran code generated for $\mathtt{A}(1:n, 1:n) = \mathtt{TRANSPOSE}(\mathtt{B}(1:n, 1:n))$ (variable names edited for clarity).

3.5 Optimizations for special distributions

If either or both of the arrays are known by the compiler to be block or cyclic, further optimization is possible. Since the ownership set is a single slice, rather than a union of slices, one or both of the outer loops can be eliminated. When f is the first index that a processor owns, l is the largest array index, and the dimension is distributed across v processors, a cyclic distribution can be described as $(f : l : v)$, and a block distribution can be described as $\left(f : f + \left\lceil \frac{E}{v} \right\rceil - 1 : 1\right)$, where E is the number of elements in the array dimension. Note, however, that there are nine sets of communication algorithms to derive, since each of the two arrays in the assignment statement can have one of three types of distributions.

Figure 7 shows the Fortran code generated by the compiler for a simple example, $\mathtt{A}(1 : n, 1 : n) = \mathtt{TRANSPOSE}(\mathtt{B}(1 : n, 1 : n))$, for $n = 1024$ and 64 processors. Both arrays have a block distribution in the first dimension and are not distributed in the second dimension; hence there is only a single loop for each array dimension. Only the send and receive code fragments are shown; we omit the initialization and the details of the loop concerning the interleaving of the sends and receives. Timings for this code are given in Section 5.3.

4 Deposit model communication

The above algorithms result in efficient code but there is still room for improvement. The sending and receiving algorithms are quite similar and share many common subexpressions; thus a significant fraction of the combined computation can be avoided if the sender computes the ownership information for both itself and the receiver, avoiding the redundancy, and then passes the result on to the receiver. Although this requires more work for the sender, it usually pays off overall, since each processor tends to send to as many processors as it receives from.

One way to achieve this goal is to have the sender compute the array *indices* for the receiver. However, this approach requires that k words of address data must be sent for a k-dimensional array, and the receiver must perform $k - 1$ multiplications and additions to compute the final destination address. A much better approach is for the sender to calculate the destination *memory addresses* for the receiver, which can then plug the data directly into memory as it arrives. We call this the *deposit model* of communication, since the sender deposits the data directly into the final destination memory locations on the receiver. This is one example of the "door-to-door delivery" communication style [2], which moves data from the user's address space on one processor directly into the user's address space on another processor, without any intermediate copies to or from system space. Sending the addresses to the receiver minimizes the amount of work that is done by the receiver; it does not have to compute the index sets of the data items to be received. Since we are compiling to a single program which is run on all processors in SPMD fashion, the sending processor has complete knowledge of the receiver's memory space and can thus easily determine the corresponding destination memory locations. If the sender does not know the base address of the receiver's array, it can simply send the offset, and the receiver can add in the base.

Note that sending along address information for the deposit model increases the overall communication load in the system. The usefulness of the deposit model depends on the ratio of communication/computation speed. For example, if the computation speed is much faster than the communication speed, the doubled cost of the communication may outweigh the time saved by the receiver not computing the destination addresses, and thus it would not be worthwhile to use the deposit model. However, the link bandwidths in parallel systems continue to increase (i.e., an iWarp node has a peak performance of 20 MFLOPS and a single iWarp link has a bandwidth of 40 MByte/sec; a Paragon XP/S node has a peak performance of 75 MFLOPS and a link bandwidth of 200 MByte/sec), so there exists a class of machines and systems that are an attractive target for this optimization. We measured the effects of the deposit model on one such system (the iWarp), and our experiments demonstrate that the deposit model pays off for a large class of communication patterns. Furthermore, it is simple for a compiler to detect those combinations of right-hand side distribution and left-hand side distribution that are not improved by the deposit model.

5 Performance results

In this section, we provide an empirical evaluation of the deposit model for compiler generated code (see [13] for more information about the compiler). The performance was measured in an implementation on a 64-node iWarp array, using a general-purpose message passing library for interprocessor communication. We base our discussion on three primitives (redistributions of one-dimensional arrays, one-dimensional shift, and two-dimensional transpose), which are compiled in isolation, and the execution of a complete program (2DFFT).

5.1 One-dimensional redistribution

Figures 8, 9, and 11 show the number of clocks for the execution of the assignment statement $\mathtt{A}(1 : n) = \mathtt{B}(1 : n)$, with various distributions for two single-precision arrays A and B, and different values of n. The communication statements for this array assignment are the same as the communication patterns occurring in any redistribution of any one-dimensional array. Three different distributions were chosen: block, cyclic, and block-cyclic with a block size of 2, and we measured all combinations of distributions for `A` and `B`. We do not report the cases in which `A` and `B` have the same distribution, because no communication is necessary, thus yielding identical code from the compiler for both situations.

The results show that the deposit model always outperforms the standard model for one-dimensional redistributions, usually by a significant amount. However, if one array has a cyclic distribution and the other has a block-cyclic distribution, each processor receives messages from k other processors, where k the block size. In our case, $k = 2$, so the communication pattern is sparse, and the deposit provides little benefit over the standard method.

These experiments lead to another important conclusion: for dense communication patterns, it is extremely important for the receiver to process incoming messages as quickly as possible. Consider an all-to-all communication pattern. With $O(p^2)$ messages travelling through the network, at any given time many messages are likely to be blocked waiting on the completion of other messages. In addition, a receiver is able to process only a fixed number of messages at a time (one message for most systems). This means that the longer the receiver spends processing an incoming message, the longer other messages must wait, even messages destined for other receivers. The deposit model allows messages to be processed as quickly as possible, subject to the extra time it takes to transmit and receive address information in addition to data. A possible solution for the standard model is for the receiver to buffer all messages and only process them after all messages have been received. However, this method requires $O(p)$ times more memory for buffers, and most large programs simply cannot afford to give up this amount of memory for communication.

5.2 One-dimensional shift

Figure 10 shows the time required for the execution of the assignment statement $\mathtt{A}(1 : n-1) = \mathtt{B}(2 : n)$, using the same distributions and problem sizes as for the one-dimensional redistribution case. We always chose the same distribution for both arrays, so that the communication follows a shift pattern. With this shift pattern, communication occurs only at block boundaries. While this minimizes the amount of communication for block distributions, it maximizes the communication for cyclic distributions, since every element is on a block boundary.

The communication time increases sharply as n crosses some threshold value. This effect is caused by an interaction between the message routing method, the hardware buffers, and the message sizes. When message sizes are sufficiently small, an entire message can fit into the hardware buffers while waiting for the receiver to become ready. This means that the sender can finish sending its message(s) and does not block on the receiver, and can thus finish any receive operations sooner. This effect propagates throughout the processors, allowing them all to complete quickly in the case of short messages. As expected, this effect shows up for smaller values of n if the deposit model is used, since both addresses and data values are transmitted.

We notice that the standard method usually outperforms the deposit model, especially in the cyclic case when the amount of data to transfer is large. In a shift operation, regardless of the distribution, the communication pattern is always sparse, thus minimizing the penalty paid for delaying the communication. The results for the block distribution are nearly identical, as expected, since only a single value is sent regardless of n.

5.3 Two-dimensional transpose

The assignment statement $\mathtt{A}(1 : n, 1 : n) = \mathtt{TRANSPOSE}(\mathtt{B}(1 : n, 1 : n))$ results in a more complicated communication pattern. Each array is of size $n \times n$ for $n \in \{128, 256, 512, 1024\}$. Both arrays are distributed in their first dimension, and the second dimensions are not distributed. For each case, the distributed dimensions have the same distribution: block, cyclic, or block-cyclic with a block size of 2. If the distributed dimensions are distributed over all the processors, this assignment statement results in an all-to-all communication pattern.

Our measurements, depicted in Figure 12, show that the deposit model performs significantly better than the standard method for this example. Since the deposit model method processes the data more quickly as it arrives, the deposit model cuts down on the network congestion and speeds up the overall communication.

5.4 2DFFT

This two-dimensional FFT program works by performing a one-dimensional FFT (1DFFT) independently on each column of the $n \times n$ matrix, performing a transpose operation, and then performing another one-dimensional FFT independently on each column. We achieve the best parallelism by distributing the matrix by columns, and performing the 1DFFT computations on the columns with a `DOALL` loop. This means that each processor can independently perform a 1DFFT on the data it owns, without any communication except for the transpose operation. The two-dimensional transpose operation is implemented with the `TRANSPOSE` array assignment statement described above.

Figure 13 shows a scatter plot of the absolute performance for each communication method. An entry above the diagonal indicates that the deposit model performed better. We see that the deposit model provides a significant improvement over the standard communication model in all cases.

5.5 Discussion

Our tests have shown that the deposit model usually executes more quickly than the standard method, due both to the complexity in determining the array indices to send and receive, and to the increased speed with which the receiver can process incoming messages. But we also note that we can relatively easily generalize the types of array statements that perform better with the particular communication methods, simply by anticipating whether the resulting communication pattern is sparse or dense. For example, we might instruct the compiler to use the standard communication method in the case of one-dimensional assignment statements where both arrays have the same distribution and access stride (since it invokes a sparse shift communication pattern), and to use the deposit model in the other cases. This choice also depends somewhat on the underlying computation and communication architecture of the system. For example, as we increase the computation power of the processor relative to the communication bandwidth, eventually the deposit model will never pay off. However, most current private memory systems provide ample communication bandwidth and can easily tolerate an increase in the message size without a comparable increase in communication time. Once the startup overhead has been paid, sending extra words in a message is cheap. For such systems, the deposit model provides a low-cost optimization.

As with all measurements on a real system, there are several features of the system that influence these measurements. The message passing system used for these measurements is based on synchronous sends and receives (i.e., both processors must be explicitly calling the send and receive routines for communication to proceed). To avoid deadlock, this method requires a choice of routing and scheduling that results in extra congestion in the system. This choice of routing and scheduling influences the overall execution time, but we feel that these perturbations do not impact our assessment of the deposit model. (In fact, the longer messages of the deposit model may even be penalized more than the shorter messages of the default model.)

6 Concluding remarks

Array statements are an attractive and powerful way for a user to specify the data parallelism in a program. When mapping an array statement onto a target private memory parallel system for a given distribution and alignment, the key task of the compiler is to generate the communication statements that provide local copies of non-local data as needed. The array assignment statement is the core statement, and solutions for this statement type can easily be extended to other array statements.

This paper introduced a novel representation of slices that does justice to the fact that real compilers do not translate array expressions in isolation but must deal with statements involving multiple operands. Similar code patterns are used to determine which elements to send or receive, but any such computation is overhead that reduces the speedup obtained from a parallel system. To keep the amount of overhead computation small and to reduce the overhead on the receiver, a compiler can exploit the similarity of determining sender and receiver information and compute this information together on the sender processor. Then the sender transmits (address, data) tagged blocks to the receiving processor. Using this method, our measurements on a 64-node iWarp system indicate savings up to 50% to 75%. Of course, the impact of this optimization depends on the specific data distribution and on the size of the data set. This optimization increases the amount of communication, but since most modern private memory parallel systems provide ample communication bandwidth, this optimization is quite worthwhile.

References

[1] G. Blelloch and J. Sipelstein. Collection-oriented languages. *Proc. IEEE*, 79(4):504–523, Apr 1991.

[2] S. Borkar, R. Cohn, G. Cox, S. Gleason, T. Gross, H. T. Kung, M. Lam, B. Moore, C. Peterson, J. Pieper, L. Rankin, P. S. Tseng, J. Sutton, J. Urbanski, and J. Webb. iWarp: An integrated solution to high-speed parallel computing. In *Proceedings of Supercomputing '88*, pages 330–339, Orlando, Florida, November 1988. IEEE Computer Society and ACM SIGARCH.

[3] M. Bromley, S. Heller, T. McNerney, and G. Steele. Fortran at ten gigaflops: The Connection Machine convolution compiler. In *Proc. ACM SIGPLAN'91 Conf. on PLDI*, pages 145–156, Toronto, June 1991. ACM.

[4] S. Chatterjee, J. Gilbert, F. Long, R. Schreiber, and S. Teng. Generating local addresses and communication sets for data-parallel programs. In *Proc. of the ACM Symposium on Principles and Practice of Parallel Programming (PPoPP)*, pages 149–158, May 1993.

[5] S. Chatterjee, J. Gilbert, R. Schreiber, and S. Teng. Automatic array alignment in data-parallel programs. In *Proc. 20th POPL Conf.*, pages 16–28. ACM, January 1993.

[6] T. Cormen, C. Leiserson, and R. Rivest. *Introduction to Algorithms*. The MIT Press, 1990.

[7] S. Hiranandani, K. Kennedy, and C. W. Tseng. Evaluation of compiler optimizations for Fortran D on MIMD distributed-memory machines. Technical Report COMP TR91-170, Dept. of Computer Science, Rice University, November 1991.

[8] S. Hiranandani, K. Kennedy, and C. W. Tseng. Compiling Fortran D for MIMD distributed-memory machines. *CACM*, 35(8):66–80, Aug 1992.

[9] D. E. Knuth. *The Art of Computer Programming*, volume 2. Addison-Wesley, 1981.

[10] C. Koelbel and P. Mehrotra. Compiling global name-space parallel loops for distributed execution. *IEEE TPDS*, 2(4):440–451, Oct 1991.

[11] G. Sabot. Internals of the Connection Machine Fortran 1.0 compiler. Talk at Carnegie Mellon, Apr 1991.

[12] James M. Stichnoth. Efficient compilation of array statements for private memory multicomputers. Technical Report CMU-CS-93-109, School of Computer Science, Carnegie Mellon University, February 1993.

[13] J. Subhlok, J. Stichnoth, D. O'Hallaron, and T. Gross. Programming task and data parallelism on a multicomputer. In *Proc. of the ACM Symposium on Principles and Practice of Parallel Programming (PPoPP)*, pages 13–22, May 1993.

[14] H. Zima and B. Chapman. Compiling for distributed-memory systems. *Proc. IEEE*, 81(2):264–287, February 1993.

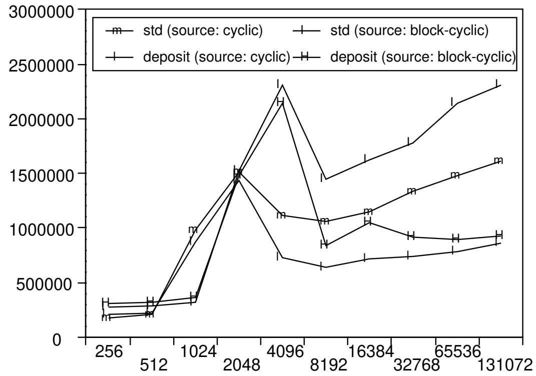

Figure 8: Execution time for $\mathtt{A}(1:n) = \mathtt{B}(1:n)$ with a *block* distribution for destination array $\mathtt{A}$.

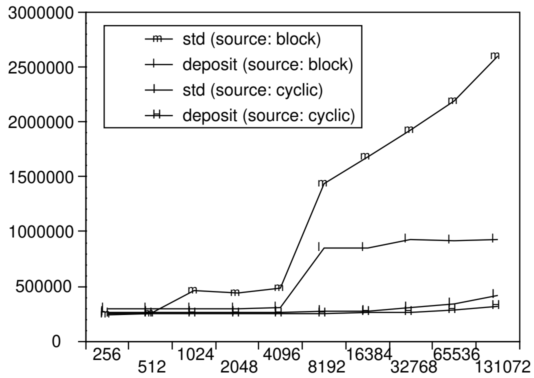

Figure 11: Execution time for $\mathtt{A}(1:n) = \mathtt{B}(1:n)$ with a *block-cyclic* distribution for destination array $\mathtt{A}$.

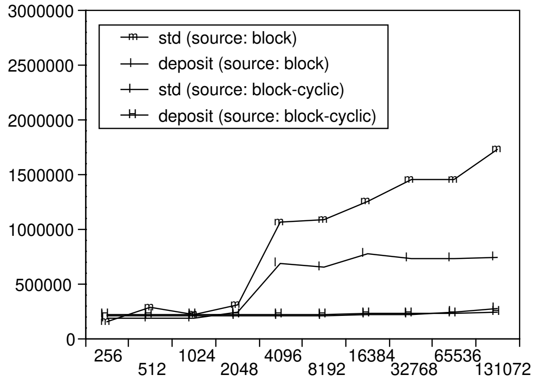

Figure 9: Execution time for $\mathtt{A}(1:n) = \mathtt{B}(1:n)$ with a *cyclic* distribution for destination array $\mathtt{A}$.

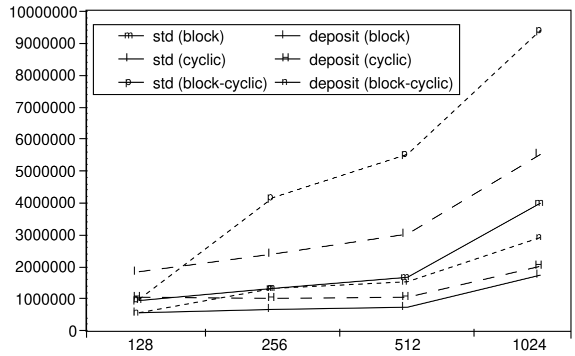

Figure 12: Execution time for $\mathtt{A}(1:n, 1:n) = \mathtt{TRANSPOSE}(\mathtt{B}(1:n, 1:n))$.

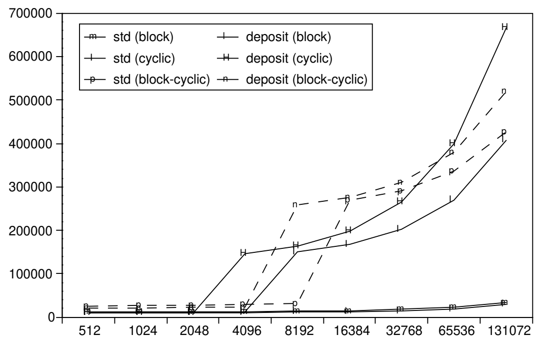

Figure 10: Execution time for $\mathtt{A}(1:n-1) = \mathtt{B}(2:n)$.

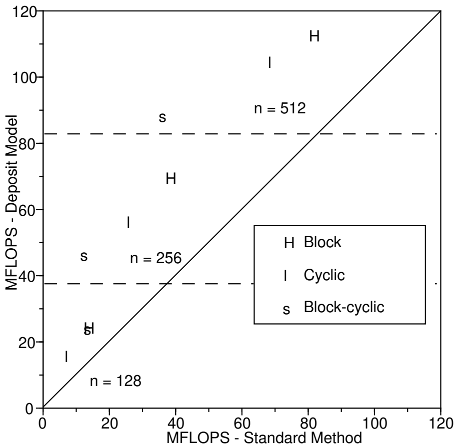

Figure 13: Execution rate comparison in MFLOPS for the 2DFFT.

Biographies

Thomas Gross is a faculty member in the School of Computer Science at Carnegie Mellon University. In 1983, he received his Ph. D. from Stanford University, where he worked on the MIPS project. Since joining Carnegie Mellon, he has been involved in the design and implementation of parallel systems. He is still interested in compilers.

Captions for Figures

Block-cyclic distribution.

Two examples of slice intersection: $(1 : 48 : 4) \cap (3 : 48 : 6) = (9 : 48 : 12)$ and $(1 : 48 : 4) \cap (4 : 48 : 6) = \emptyset$.

Block-cyclic sets are not closed under intersection.

A block-cyclic distribution as a union of disjoint slices.

Communication sets that are difficult to parameterize.

Local memory mapping function for a processor.

Fortran code generated for $\mathtt{A}(1 : n, 1 : n) = \mathtt{TRANSPOSE}(\mathtt{B}(1 : n, 1 : n))$ (variable names edited for clarity).

Execution time for $\mathtt{A}(1 : n) = \mathtt{B}(1 : n)$ with a *block* distribution for destination array `A`.

Execution time for $\mathtt{A}(1 : n) = \mathtt{B}(1 : n)$ with a *cyclic* distribution for destination array `A`.

Execution time for $\mathtt{A}(1 : n - 1) = \mathtt{B}(2 : n)$.

Execution time for $\mathtt{A}(1 : n) = \mathtt{B}(1 : n)$ with a *block-cyclic* distribution for destination array `A`.

Execution time for $\mathtt{A}(1 : n, 1 : n) = \mathtt{TRANSPOSE}(\mathtt{B}(1 : n, 1 : n))$.

Execution rate comparison in MFLOPS for the 2DFFT.