

A Linear Algebra Framework for Static High Performance Fortran Code Distribution*

CORINNE ANCOURT, FABIEN COELHO, FRANÇOIS IRIGOIN, AND RONAN KERYELL

Centre de Recherche en Informatique, École Nationale Supérieure des Mines de Paris, 35, rue Saint-Honoré, F-77305 Fontainebleau cedex, France; e-mail: coelho@cri.ensmp.fr

ABSTRACT

High Performance Fortran (HPF) was developed to support data parallel programming for single-instruction multiple-data (SIMD) and multiple-instruction multiple-data (MIMD) machines with distributed memory. The programmer is provided a familiar uniform logical address space and specifies the data distribution by directives. The compiler then exploits these directives to allocate arrays in the local memories, to assign computations to elementary processors, and to migrate data between processors when required. We show here that linear algebra is a powerful framework to encode HPF directives and to synthesize distributed code with space-efficient array allocation, tight loop bounds, and vectorized communications for INDEPENDENT loops. The generated code includes traditional optimizations such as guard elimination, message vectorization and aggregation, and overlap analysis. The systematic use of an affine framework makes it possible to prove the compilation scheme correct. © 1997 John Wiley & Sons, Inc.

1 INTRODUCTION

Distributed memory multiprocessors can be used efficiently if each local memory contains the right pieces of data, if local computations use local data, and if missing pieces of data are quickly moved at the right time between processors. Macro packages and libraries are available to ease the programmer's burden but the

level of details still required transforms the simplest algorithm, e.g., a matrix multiply, into hundreds of lines of code. This fact decreases programmer productivity and jeopardizes portability as well as the economical survival of distributed memory parallel machines.

Manufacturers and research laboratories, led by Digital and Rice University, decided in 1991 to shift part of the burden to compilers by providing the programmer a uniform address space to allocate objects and a (mainly) implicit way to express parallelism. Numerous research projects [1–3] and a few commercial products had shown that this goal could be achieved and the High Performance Fortran (HPF) Forum was set up to select the most useful functionalities and to standardize the syntax. The initial definition of the new language, HPF, was frozen in May 1993 and corrections were added in November 1994 [4].

* An early version of this article was presented at the Fourth International Workshop on Compilers for Parallel Computers held in Delft, the Netherlands, December 1993.

Received May 1995

Revised November 1995

© 1997 by John Wiley & Sons, Inc.

Scientific Programming, Vol. 6, pp. 3–27 (1997)

CCC 1058-9244/97/010003-25

```

! non-independent
!   A in the rhs may
!   induce RW dependences...
FORALL (i=1:n, j=1:m, MASK(i,j))
  A(i,j) = f(A, ...)

! array TMP declared and mapped as A
! initial copy of A into TMP
!   because of potential RW dependences
INDEPENDENT(j,i)
do j=1, m
  do i=1, n
    TMP(i,j) = A(i,j)
  enddo
enddo
! possible synchro...
INDEPENDENT(j,i)
do j=1, m
  do i=1, n
    if (MASK(i,j)) A(i,j) = f(TMP, ...)
  enddo
enddo

```

FIGURE 1 Masked FORALL to INDEPENDENT loops.

Prototype compilers incorporating some HPF features are available [3, 5–9]. Commercial compilers from APR [10, 11], DEC [12, 13], IBM [14], and PGI [15, 16] are also being developed or are already available. These compilers implement part or all of the HPF subset, which only allows static distribution of data and prohibits dynamic redistributions.

This article deals with this HPF static subset and shows how changes of basis and affine constraints can be used to relate the global memory and computation spaces seen by the programmer to the local memory and computation spaces allocated to each elementary processor. These relations, which depend on HPF directives added by the programmer, are used to allocate local parts of global arrays and temporary copies which are necessary when nonlocal data is used by local computations. These constraints are also used in combination with the *owner-computes rule* to decide which computations are local to a processor, and to derive loop bounds. Finally, they are used to generate send and receive statements required to access nonlocal data.

These three steps, local memory allocation, local iteration enumeration, and data communication, are put together as a general compilation scheme for parallel loops, known as INDEPENDENT in HPF, with affine bounds and subscript expressions. HPF's FORALL statements or constructs, as well as a possible future ON extension to advise the compiler about the distribution of iterations onto the processors, can be translated into a set of independent loops by introducing a temporary array mapped as required to store the intermediate results. These translations are briefly outlined in Figures 1 and 2. The resulting code is a pair of loops which can be compiled by our scheme, following the *owner-computes rule*, if the ON clause is put into the affine framework. The FORALL transla-

tion requires a temporary array due to its single-instruction multiple-data (SIMD)-like semantics. However, if the assigned array is not referenced in the *rhs*, the FORALL loop is independent and should be tagged as such to fit directly our scheme. Such necessary temporary arrays are not expected to cost much, both from the compilation and execution point of view: The allocated memory is reusable (it may be allocated on the stack), and the copy assignment on local data should be quite fast.

This compilation scheme directly generates optimized code which includes techniques such as guard elimination [1], message vectorization, and aggregation [2, 3]. It is compatible with overlap analysis [1]. There are no restrictions either on the kind of distribution (general cyclic distributions are handled) or on the rank of array references (the dimension of the referenced space: for instance rank of $A(i, i)$ is 1). The memory allocation part, whether based on overlap extensions or dealing with temporary arrays needed to store both remote and local elements, is independent of parallel loops and can always be used. The relations between the global programmer space and the local processor spaces can also be used to translate sequential loops with a run-time resolution mechanism or with some optimizations. The reader is assumed knowledgeable in HPF directives [4] and optimization techniques [1, 3].

The article is organized as follow. Section 2 shows how HPF directives can be expressed as affine constraints and normalized to simplify the compilation process and its description. Section 3 presents an overview of the compilation scheme and introduces the basic sets *Own*, *Compute*, *Send*, and *Receive* that are used to allocate local parts of HPF arrays and temporaries, to enumerate local iterations, and to generate data exchanges between processors. Section 4 refines

```

INDEPENDENT(j,i), ON(...)
do j=1, m
  do i=1, n
    A(i,j) = f(...)
  enddo
enddo

! array TMP(iteration space)
! mapped as ON(iteration space)
INDEPENDENT(j,i)
do j=1, m
  do i=1, n
    TMP(i,j) = f(...)
  enddo
enddo
! possible synchro...
INDEPENDENT(j,i)
do j=1, m
  do i=1, n
    A(i,j) = TMP(i,j)
  enddo
enddo
    
```

FIGURE 2 ON to INDEPENDENT loops.

these sets to minimize the amount of memory space allocated, to reduce the number of loops whenever possible, and to improve the communication pattern. This is achieved by using different coordinates to enumerate the same sets. Examples are shown in Section 5 and the method is compared with previous work in Section 6.

2 HPF DIRECTIVES

The basic idea of this work was to show that the HPF compilation problem can be put into a linear form, including both equalities and inequalities, then to show how to use polyhedron manipulation and scanning techniques to compile an HPF program from this linear form. The linear representations of the array and HPF declarations, the data mapping, and the loop nest accesses are presented in this section.

HPF specifies data mappings in two steps. First, the array elements are aligned with a template, which is an abstract grid used as a convenient way to relate different arrays together. Each array element is assigned to at least one template cell thru the ALIGN directive. Second, the template is distributed onto the processors, which is an array of virtual processors. Each template cell is assigned to one and only one processor through the DISTRIBUTE directive. The template and processors are declared with the TEMPLATE and PROCESSORS directives, respectively.

Elements of arrays aligned with the same template cell are allocated on the same elementary processor. Expressions using these elements can be evaluated locally, without interprocessor communications. Thus, the alignment step mainly depends on the algorithm. The template elements are packed in blocks to reduce communication and scheduling overheads without in-

creasing load imbalance too much. The block sizes depend on the target machine, while the load imbalance stems from the algorithm. Templates can be bypassed by aligning an array on another array and by distributing arrays directly on processors. This does not increase the expressiveness of the language but implies additional check on HPF declarations. Templates are systematically used in this article to simplify algorithm descriptions. Our framework deals with both stages and could easily tackle direct alignment and direct distribution. The next sections show that any HPF directive can be expressed as a set of affine constraints.

2.1 Notations

Throughout this article, a lower case letter as v denotes a vector of integers, which may be variables and constants. v_i , ($i \geq 0$) is the i th element or variable of vector v . Subscript 0, as in v_0 , denotes a constant integer vector. As a convention, a denotes the variables which describe the elements of an array, t is used for templates, and p for processors. An upper case letter as A denotes a constant integer matrix. Constants are implicitly expanded to the required number of dimensions. For instance 1 may denote a vector of 1. $|A|$ denotes the determinant of matrix A .

2.2 Declarations

The data arrays, templates, and processors are declared as Cartesian grids in HPF. If a is the vector of variables describing the dimensions of array $A(3:8, -2:5, 7)$, then the following linear constraints are induced on a :

$$3 \leq a_1 \leq 8, \quad -2 \leq a_2 \leq 5, \quad 1 \leq a_3 \leq 7$$

These may be translated into the matrix form $D_A a \leq d$ where:

$$D_A = \begin{pmatrix} 1 & 0 & 0 \\ -1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & -1 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & -1 \end{pmatrix},$$

$$a = \begin{pmatrix} a_1 \\ a_2 \\ a_3 \end{pmatrix}, d = \begin{pmatrix} 8 \\ -3 \\ 5 \\ 2 \\ 7 \\ -1 \end{pmatrix}$$

Any valid array element must verify the linear constraints, i.e., $A(a_1, a_2, a_3)$ is a valid array element if Equation 1 is verified by vector a . In the remaining article it is assumed without loss of generality that the dimension lower bounds are equal to 0. This assumption simplifies the formula by deleting a constant offset at the origin. Thus, the declaration constraints for $A(0:5, 0:7, 0:6)$ can be written intuitively as

$$0 \leq a < D.1 \quad (1)$$

where

$$D.1 = \begin{pmatrix} 6 & 0 & 0 \\ 0 & 8 & 0 \\ 0 & 0 & 7 \end{pmatrix} \begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix} = \begin{pmatrix} 6 \\ 8 \\ 7 \end{pmatrix}$$

D is a diagonal matrix composed of the sizes of A . The rationale for storing the sizes in a diagonal matrix instead of a vector is that this form is needed for the distribution formula (see Section 2.4). Likewise, the template declaration constraints are $0 \leq t < T.1$ and the processors $0 \leq p < P.1$.

2.3 Alignment

The `ALIGN` directive is used to specify the alignment relation between one object and one or many other objects through an affine expression. The alignment of an array A on a template T is an affine mapping

from a to t . Alignments are specified dimension-wise with integer affine expressions as template subscript expressions. Each array index can be used at most once in a template subscript expression in any given alignment, and each subscript expression cannot contain more than one index [4]. Let us consider the following HPF alignment directive, where the first dimension of the array is collapsed and the most general alignment subscripts are used for the other dimensions:

```
align A(*, i, j) with T(2*j-1, -i+7)
```

it induces the following constraints between the dimensions of A and T represented, respectively, by vectors a and t :

$$t_1 = 2a_3 - 1, t_2 = -a_2 + 7$$

Thus, the relation between A and T is given by $t = Aa + s_0$ where

$$t = \begin{pmatrix} t_1 \\ t_2 \end{pmatrix}, A = \begin{pmatrix} 0 & 0 & 2 \\ 0 & -1 & 0 \end{pmatrix},$$

$$a = \begin{pmatrix} a_1 \\ a_2 \\ a_3 \end{pmatrix}, s_0 = \begin{pmatrix} -1 \\ 7 \end{pmatrix}$$

A (for alignment) is a matrix with at most one nonzero integer element per column (dummy variables must appear at most once in the alignment subscripts) and per row (no more than one dummy variable in an alignment subscript) and s_0 is a constant vector that denotes the constant shift associated to the alignment. Note that since the first dimension of the array is collapsed the first column of A is null, and that the remaining columns constitute an antidiagonal matrix since i and j are used in reverse order in the subscript expressions.

However, this representation cannot express replication of elements on a dimension of the template, as allowed by HPF. The former relation is generalized by adding a projection matrix R (for replication) which selects the dimensions of the template which are actually aligned. Thus the following example:

```
align A(i) with T(2*i-1, *)
```

is translated into

$$Rt = Aa + s_0 \quad (2)$$

```

template T(0:99), T'(0:99)
processors P(0:4)
distribute T(block(20)), T'(cyclic(1)) onto P
    
```

FIGURE 3 Example 1.

where $R = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$, $t = \begin{pmatrix} t_1 \\ t_2 \end{pmatrix}$, $A = (2)$, and $a = (a_1)$, $s_0 = (-1)$.

The same formalism can also deal with a chain of alignment relations (A is aligned to B which is aligned to C, etc.) by composing the alignments. To summarize, a “*” in an array reference induces a zero column in A and a “*” in a template reference a zero column in R . When no replication is specified, R is the identity matrix. The number of rows of R and A corresponds to the number of template dimensions that is actually aligned, thus leading to equations linking template and array dimensions. Template dimensions with no matching array dimensions are removed through R .

2.4 Distribution

Once optional alignments are defined in HPF, objects or aligned objects can be distributed on Cartesian processor arrangements, declared like arrays and templates, and represented by the inequality $0 \leq p < P$. Each dimension can be distributed by block, or in a cyclic way, on the processor set or even collapsed on the processors. Let us first consider Examples 1 and 2 in Figures 3 and 4:

Example 1

In Example 1, the distributions of the templates create a link between the template elements t and the processors p , so that each processor owns some of the template elements. These links can be translated into linear formulae with the addition of variables. For the `block` distribution of template T , assuming the local offset l within the size 20 block, $0 \leq l < 20$, then the formula simply is:

$$t = 20p + l$$

For a fixed processor p and the allowed range of offsets

```

template T(0:99,0:99,0:99)
processors P(0:9,0:9)
distribute T(*,cyclic(4),block(13)) onto P
    
```

FIGURE 4 Example 2.

within a block l , the formula gives the corresponding template elements that are mapped on that processor. There is no constant in the equality due to the assumption that template and processor dimensions start from 0. For the `cyclic` distribution of template T' , a cycle variable c counts the number of cycles over the processors for a given template element. Thus, the formula is:

$$t' = 5c + p$$

The cycle number c generates an initial offset on the template for each cycle over the processors. Then the p translation associates the processor's owned template element for that cycle. The general `cyclic`(n) multi-dimensional case necessitates both cycle c and local offset l variable vectors, and can be formulated with a matrix expression to deal with all dimensions at once.

General Case

HPF allows a different block size for each dimension. The extents (n in `BLOCK`(n) or `CYCLIC`(n)) are stored in a diagonal matrix, C . P is a square matrix with the size of the processor dimensions on the diagonal. Such a distribution is not linear according to its definition [4], but may be written as a linear relation between the processor coordinate p , the template coordinate t , and two additional variables, l and c :

$$\Pi t = CPc + Cp + l \quad (3)$$

Vector l represents the offset within one block in one processor and vector c represents the number of wrap-arounds (the sizes of which are CP) that must be performed to allocate blocks cyclically on processors as shown on Figure 5 for the example in Figure 7.

The projection matrix Π is needed when some dimensions are collapsed on processors, which means that all the elements on the dimension are on the same processor. These dimensions are thus orthogonal to the processor parallelism and can be eliminated. The usual modulo and integer division operators dealing with the block size are replaced by predefined constraints on p and additional constraints on l . Vector c is implicitly constrained by array declarations and/or (depending on the replication) by the `TEMPLATE` declaration.

Specifying `BLOCK`(n) in a distribution instead of `CYCLIC`(n) is equivalent but tells the compiler that the distribution will not wrap around and that Equation 3 can be simplified by zeroing the c coordinate for this dimension. This additional equation reduces the number of free variables and, hence, removes a loop in the generated code. `CYCLIC` is equivalent to `CYCLIC(1)`

$l \downarrow c \rightarrow$	$p = 0$				$p = 1$				$p = 2$				$p = 3$			
	0	1	2	3	0	1	2	3	0	1	2	3	0	1	2	3
0	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
1	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
2	32	33	34	35	36	37	38	39	40	41	42	43	44	45	46	47
3	48	49	50	51	52	53	54	55	56	57	58	59	60	61	62	63
4	64	65	66	67	68	69	70	71	72	73	74	75	76	77	78	79
5	80	81	82	83	84	85	86	87	88	89	90	91	92	93	94	95
6	96	97	98	99	100	101	102	103	104	105	106	107	108	109	110	111
7	112	113	114	115	116	117	118	119	120	121	122	123	124	125	126	127

```

chpf$ processors P(0:3)
chpf$ template T(0:127)
chpf$ distribute T(cyclic(4)) onto P

```

FIGURE 5 Example of a template distribution. (From Chatterjee et al. [17], with permission.)

which is translated into an additional equation: $l = 0$. BLOCK without a parameter is equivalent to $\text{BLOCK}(\lceil e_i/e_p \rceil)$ where e_i and e_p are, respectively, the template and processor extents in the matching dimensions. Since block distributions do not cycle around the processors, $c = 0$ can be added.

The variables are bounded by:

$$0 \leq p < P, 0 \leq t < T, 0 \leq l < C \quad (4)$$

Example 2

Let us consider the second, more general example, in Figure 4. These directives can be translated into the following matrix form:

$$\begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} t = \begin{pmatrix} 40 & 0 \\ 0 & 130 \end{pmatrix} c + \begin{pmatrix} 4 & 0 \\ 0 & 13 \end{pmatrix} p + l,$$

$$\begin{pmatrix} 0 \\ 0 \end{pmatrix} \leq l < \begin{pmatrix} 4 \\ 13 \end{pmatrix}$$

where

$$t = \begin{pmatrix} t_1 \\ t_2 \\ t_3 \end{pmatrix}, p = \begin{pmatrix} p_1 \\ p_2 \end{pmatrix}, c = \begin{pmatrix} c_1 \\ c_2 \end{pmatrix}, l = \begin{pmatrix} l_1 \\ l_2 \end{pmatrix}, c_2 = 0$$

and

$$\Pi = \begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}, C = \begin{pmatrix} 4 & 0 \\ 0 & 13 \end{pmatrix}, P = \begin{pmatrix} 10 & 0 \\ 0 & 10 \end{pmatrix}$$

In this example, the distribution is of type BLOCK in the last dimension, thus the added $c_2 = 0$. If a dimension of a template is collapsed on the processors, the corresponding coordinate is useless because even if it is used to distribute the dimension of an array, this dimension is collapsed on the processors. Thus, it can be discarded from the equations. It means that it can be assumed that $\Pi = I$ in Equation 3 if these useless dimensions are removed by a normalization process.

2.5 Iteration Domains and Subscript Functions

Although iteration domains and subscript functions do not concern directly the placement and the distribution of the data in HPF, they must be considered for the code generation. Since loops are assumed INDEPENDENT with affine bounds, they can be represented by a set of inequalities on the iteration vector i : The original enumeration order is not kept by this representation, but the independence of the loop means that the result of the computation does not depend on this very order. The array reference links the iteration vector to the array dimensions.

Let us consider for instance the example in Figure 6. The iteration domain can be translated into a parametric (n_1 and n_2 may not be known at compile-

```

INDEPENDENT(i2,i1)
do j=1, n1, 3
  do i1=j+2, n2+1
    A(2*i1+1, n1-j+i1) = ...
  enddo
enddo
! normalized (stride 1) version
do i2=0, (n1-1)/3
  do i1=3*i2+3, n2+1
    A(2*i1+1, n1-3*i2+i1-1) = ...
  enddo
enddo

```

FIGURE 6 Loop example.

time) form, where the constant vector is a function over the parameters. Thus, the set of constraints for the loop nest iterations, with the additional change of variable $j = 3i_2 + 1$ to normalize the loop [18], is:

$$\begin{pmatrix} 0 & -1 \\ 0 & 3 \\ 1 & 0 \\ -1 & 3 \end{pmatrix} \begin{pmatrix} i_1 \\ i_2 \end{pmatrix} \leq \begin{pmatrix} 0 \\ n_1 - 1 \\ n_2 + 1 \\ -3 \end{pmatrix}$$

Moreover, the array reference can also be translated into a set of linear equalities:

$$a = \begin{pmatrix} a_1 \\ a_2 \end{pmatrix} = \begin{pmatrix} 2 & 0 \\ 1 & -3 \end{pmatrix} \begin{pmatrix} i_1 \\ i_2 \end{pmatrix} + \begin{pmatrix} 1 \\ n_1 - 1 \end{pmatrix}$$

In the general case, a set of constraints is derived for the iteration domain of the loop nest (Equation 5, where n stands for a vector of parameters for the loop nest) and for the array references (Equation 6).

$$Li \leq b_0(n) \quad (5)$$

$$a = Si + a_0(n) \quad (6)$$

2.6 Putting It All Together

Let us now consider the example in Figure 7. According to the previous sections, all the declaration of distributions, alignments, arrays, processors and templates, iterations, and references can be rewritten as

$$0 \leq l < C.1,$$

$$0 \leq a < D.1,$$

$$0 \leq p < P.1,$$

$$0 \leq t < T.1,$$

$$Rt = Aa + s_0,$$

$$t = CPc + Cp + L,$$

$$Li \leq b_0(n),$$

$$a = Si + a_0(n)$$

where $D = (43)$, $T = (128)$, $A = (3)$, $s_0 = (0)$, $R = (1)$, $C = (4)$, $P = (4)$, $L = \begin{pmatrix} 3 \\ -3 \end{pmatrix}$, $b_0(n) = \begin{pmatrix} U \\ 0 \end{pmatrix}$, $S = (3)$, and $a_0(n) = (0)$.

```

real A(0:42)
!HPF$ template T(0:127)
!HPF$ processors P(0:3)
!HPF$ align A(i) with T(3*i)
!HPF$ distribute T(cyclic(4)) onto P
A(0:U:3) = ...
    
```

FIGURE 7 Example in Chatterjee et al. [17].

The array mapping of A is represented by the “○” or “●” on Figure 8 and the written elements (if $U = 42$) of A with “●”. As can be seen on this two-dimensional (2D) representation of a 3D space, the mapping and the access are done according to a regular pattern, namely a lattice. This is exploited to translate the compilation problem in a linear algebra framework.

3 OVERVIEW OF THE COMPILATION SCHEME

The compilation scheme is based on the HPF declarations, the owner computes rule, and the single-program multiple data (SPMD) paradigm, and deals with INDEPENDENT loops. Loop-bound and array subscript expressions are assumed affine. Such a parallel loop independently assigns elements of some left-hand side x with some right-hand side expressions f on variables Y, Z , etc. (Fig. 9). The loop is normalized to use unit steps as explained in the previous section. Since the loop is independent, S_x must be injective on the iteration domain for an independent parallel loop to be deterministic. The n models external run-time parameters that may be involved in some subexpressions, namely the loop bounds and the reference shifts.

The HPF standard does not specify the relationship between the HPF processors and the actual physical processors. Thus, efficiently compiling loop nests involving arrays distributed onto different processor sets without further specification is beyond the scope of this article. But if the processor virtualization process can be represented by affine functions, which would probably be the case, we can encompass these different processor sets in our framework by remaining on the physical processors. In the following, we assume for clarity that the interfering distributed arrays in a loop nest are distributed on the same processor set.

3.1 Own Set

The subset of x that is allocated on processor p according to HPF directives must be mapped onto an

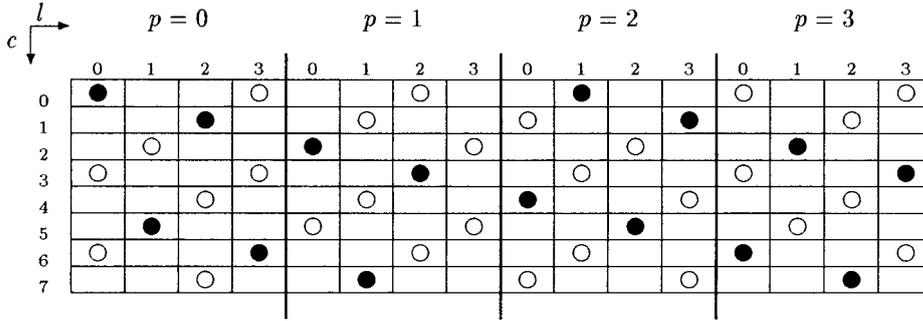


FIGURE 8 Example of an array mapping and writing. (From Chatterjee et al. [17], with permission.)

array of the output program. Let $Own_x(p)$ be this subset of X and X' the local array mapping this subset. Each element of $Own_x(p)$ must be mapped to one element of X' but some elements of X' may not be used. Finding a good mapping involves a tradeoff between the memory space usage and the access function complexity. This is studied in Section 4.3. Subset $Own_x(p)$ is derived from HPF declarations, expressed in an affine formalism (see Section 2), as:

$$\begin{aligned}
 Own_x(p) &= \{a | \exists t, \exists c, \exists l, s.t. \quad R_x t = A_x a + s_{x0} \\
 &\quad \wedge t = C_x P c + C_x p + l_x \\
 &\quad \wedge 0 \leq a < D_x . 1 \\
 &\quad \wedge 0 \leq p < P . 1 \\
 &\quad \wedge 0 \leq l < C_x . 1 \\
 &\quad \wedge 0 \leq t < T_x . 1\}
 \end{aligned}$$

Subset $Own_x(p)$ can be mapped onto a Fortran array by projection on c and l , or on any equivalent set of variables, i.e., up to an injective mapping. Although Own is not a dense polyhedron as defined here, it can be seen as such in the higher dimensional (a, c, l) space. Thus, this view is used in the following, although our interest is just to represent the array element a . Note that for a given distributed dimension a_i there is one and only one corresponding (p_j, c_j, l_j) triplet.

3.2 Compute Set

Using the owner computes rule, the set of iterations local to p , $Compute(p)$, is directly derived from the

```

INDEPENDENT(i)
forall(Li ≤ b0(n))
  X(Sxi + ax0(n)) = f(Y(Syi + ay0(n)), Z(Szi + az0(n)))

```

FIGURE 9 Generic loop.

previous set, $Own_x(p)$. The iterations to be computed by a given processor are those of the loop nest for which the lhs are owned by the processor:

$$\begin{aligned}
 Compute(p) &= \{i | S_x i + a_{x0}(n) \in Own_x(p) \\
 &\quad \wedge Li \leq b_0(n)\}
 \end{aligned}$$

Note that $Compute(p)$ and $Own_x(p)$ are equivalent sets if the reference is direct ($S_x = I$, $a_{x0} = 0$) and if the iteration and array spaces conform. This is used in Section 4.3 to study the allocation of local arrays as a special case of local iterations. According to this definition of $Compute$, when X is replicated, its new values are computed on all processors having a copy. Depending on a tradeoff between communication and computation speed, the optimal choice may be to broadcast the data computed once in parallel rather than to compute each value locally.

3.3 View Set

The subset of Y (or Z , etc.) used by the loop that computes X is induced by the set of local iterations:

$$View_y(p) = \{a | \exists i \in Compute(p) s.t. a = S_y i + a_{y0}(n)\}$$

Note that unlike S_x , matrix S_y is not constrained and cannot always be inverted.

If the intersection of this set with $Own_y(p)$ is non-empty, some elements needed by processor p for the computation are fortunately on the same processor p . Then, in order to simplify the access to $View_y(p)$ without having to care about dealing differently with remote and local elements in the computation part, the local copy Y' may be enlarged so as to include its neighborhood, including $View_y(p)$. The neighborhood is usually considered not too large when it is bounded by a numerical small constant, which is typically the case if X and Y are identically aligned and accessed at least on one dimension up to a translation, such as

encountered in numerical finite difference schemes. This optimization is known as overlap analysis [1]. Once remote values are copied into the overlapping \mathbb{Y}' , all elements of $View_{\mathbb{Y}}(p)$ can be accessed uniformly in \mathbb{Y} with no overhead. However, such an allocation extension may be very rough and expensive in some cases (e.g., a matrix multiplication) because of the induced memory consumption. A reusable temporary array might be preferred, and locally available array elements must be copied. Such tradeoffs to be considered in the decision process are not discussed in this article, but we present the techniques for implementing both solutions.

When \mathbb{Y} (or \mathbb{X} , or \mathbb{Z} , etc.) is referenced many times in the input loop or in the input program, these references must be clustered according to their connection in the dependence graph [19]. Input dependencies are taken into account as well as usual ones (flow, anti, and output dependencies). If two references are independent, they access two distant areas in the array and two different local copies should be allocated to reduce the total amount of memory allocated: a unique copy would be as large as the convex hull of these distant regions. If the two references are dependent, only one local copy should be allocated to avoid any consistency problem between copies. If \mathbb{Y} is a distributed array, its local elements must be taken into account as a special reference and be accessed with (p, c, l) instead of i .

The definition of *View* is thus altered to take into account array regions. These regions are the result of a precise program analysis which is presented in [20–27]. An array region is a set of array elements described by equalities and inequalities defining a convex polyhedron. This polyhedron may be parameterized by program variables. Each array dimension is described by a variable. The equations due to subscript expression $S_{\mathbb{Y}}$ are replaced by the array region, a parametric polyhedral subset of \mathbb{Y} which can be automatically computed. For instance, the set of references to \mathbb{Y} performed in the loop body of:

```
do i2 = 1, N
  do i1 = 1, N
    X(i1, i2) = Y(i1, i2) + Y(i1-1, i2)
              + Y(i1, i2-1)
  enddo
enddo
```

is automatically summarized by the parametric region on $a_{\mathbb{Y}} = (y_1, y_2)$ for the \mathbb{Y} reference in the loop body of the example above as:

$$\mathbb{Y}(y_1, y_2) : \{(y_1, y_2) | y_1 \leq i_1 \wedge y_2 \leq i_2 \wedge i_1 + i_2 \leq y_1 + y_2 + 1\}$$

Subset $View_{\mathbb{Y}}$ is still polyhedral and array bounds can be derived by projection, up to a change of basis. If regions are not precise enough because convex hulls are used to summarize multiple references, it is possible to use additional parameters to exactly express a set of references with regular holes [28]. This might be useful for red-black SOR.

As mentioned before, if \mathbb{Y} is a distributed array and its region includes local elements, it might be desired to simply extend the local allocation so as to simplify the addressing and to allow a uniform way of accessing these array elements for a given processor $p_{\mathbb{Y}}$ and cycle $c_{\mathbb{Y}}$. The required extension is computed simply by mixing the distribution and alignment equations to the $View_{\mathbb{Y}}$ description. For the example above, assuming that \mathbb{X} and \mathbb{Y} are aligned and thus distributed the same way, the overlap is expressed on the block offset $l_{\mathbb{Y}}$ with:

$$\begin{aligned} R_{\mathbb{Y}} t &= A_{\mathbb{Y}} a_{\mathbb{Y}} + s_{\mathbb{Y}0} \\ t &= CPc_{\mathbb{Y}} + Cp_{\mathbb{Y}} + l_{\mathbb{Y}} \\ a_{\mathbb{Y}} &\in View_{\mathbb{Y}}(p) \end{aligned}$$

Vectors $c_{\mathbb{Y}}$ and $p_{\mathbb{Y}}$ are constrained as usual by the loop, the processor declaration, and the template declaration of \mathbb{X} , but vector $l_{\mathbb{Y}}$ is no more constrained in the block size (C). It is indirectly constrained by $a_{\mathbb{Y}}$ and the \mathbb{Y} region. This leads to a $-1 \leq l_{\mathbb{Y}}$ lower bound instead of the initial $0 \leq l_{\mathbb{Y}}$, expressing the size one overlap on the left.

3.4 Send and Receive Sets

The set of elements that are owned by p and used by other processors and the set of elements that are used by p but owned by other processors are both intersections of previously defined sets:

$$\begin{aligned} Send_{\mathbb{Y}}(p, p') &= Own_{\mathbb{Y}}(p) \cap View_{\mathbb{Y}}(p') \\ Receive_{\mathbb{Y}}(p, p') &= View_{\mathbb{Y}}(p) \cap Own_{\mathbb{Y}}(p') \end{aligned}$$

These two sets cannot always be used to derive the data exchange code. First, a processor p does not need to exchange data with itself. This cannot be expressed directly as an affine constraint and must be added as a guard in the output code.* Second, replicated arrays

* This could be avoided by exchanging data first with processors p' such that $p' < p$ and then with processors such that $p' > p$, using the lexicographic order. But this additional optimization is not likely to decrease the execution time much, because the loop over p' is an outermost loop.

would lead to useless communication: Each owner would try to send data to each viewer. An additional constraint should be added to restrict the sets of communication to needed ones. Different techniques can be used to address this issue: (1) Replication allows broadcasts and/or load balance, which is simply translated into linear constraints as described in [29]. (2) The assignment of owners to viewers can also be optimized in order to reduce the distance between communicating processors. For instance, the cost function could be the minimal Manhattan distance* between p and p' or the lexicographically minimal vector† $p' - p$ if the interconnection network is a $|p|$ -dimension grid.‡ A combined cost function might even be better by taking advantage of the Manhattan distance to minimize the number of hops and of the lexicographical minimum to ensure uniqueness. These problems can be cast as linear parametric problems and solved [30].

When no replication occurs, elementary data communications implied by $Send_v$ and $Receive_v$ can be parametrically enumerated in basis (p, u) , where u is a basis for Y' , the local part of Y , the allocation and addressing of which are discussed in Section 4.3. $Send$ and $Receive$ are polyhedral sets and algorithms in [28] can be used. If the last component of u is allocated contiguously, vector messages can be generated.

3.5 Output SPMD Code

The generic output SPMD code, parametric on the local processor p , is shown in Figure 10. U represents any local array generated according to the dependence graph. Communications between processors can be executed in parallel when asynchronous send/receive calls and/or parallel hardware is available as in transputer-based machines or in the Paragon. The guards $Send_0$, etc. can be omitted if unused HPF processors are not able to free physical processors for other useful tasks. The loop-bound correctness ensures that no spurious iterations or communications occur. The parallel loop on U can be exchanged with the parallel inner loop on p' included in the `forall`. This other ordering makes message aggregation possible.

This code presented in Figure 10 is quite different from the one suggested in [31, 32]. First, the set of local iterations $Compute(p)$ is no longer expressed in

```

real X'((c, l) ∈ OwnX(p)),
      Y'((c, l) ∈ ViewY(p)),
      Z'((c, l) ∈ ViewZ(p))
forall(U ∈ {Y', Z', ...})
  forall((p, p'), p ≠ p', SendU(p, p') ≠ ∅ )
    forall((l, c) ∈ SendU(p, p'))
      send(p', U(l, c))
forall(U ∈ {Y', Z', ...})
  forall((p, p'), p ≠ p', ReceiveU(p, p') ≠ ∅ )
    forall((l, c) ∈ ReceiveU(p, p'))
      U(l, c) = receive(p')
if Compute(p) ≠ ∅
  forall((l, c) ∈ Compute(p))
    X'(SX(l, c)) = f(Y'(SY(l, c)),
                    Z'(SZ(l, c)), ...)

```

FIGURE 10 The output SPMD code.

the i frame but in the (c, l) frame: There is no actual need to know about the initial index values in the original code. Second, no difference is made between local and nonlocal iteration computations, at the price of possible communication and computation overlaps, but at the benefice of a homogeneous addressing: Only one reference and addressing should be generated for each original reference in the computation expression, thus avoiding costly code duplications or run-time guards to deal with each reference of an expression that may or may not be locally available at every iteration. The temporary array management and addressing issues are discussed further in the next sections. Third, nonlocal iterations computed as a set difference are not likely to produce an easily manageable set for code generation, since they should generally be nonconvex. Fourth, the parametric description of the communications allows enumeration of a subset of active processors.

Additional changes of basis or changes of coordinates must be performed to reduce the space allocated in local memory and to generate a minimal number of more efficient scanning loops. Each change must exactly preserve integer points, whether they represent array elements or iterations. This is detailed in the next section.

4 REFINEMENT

The pseudocode shown in Figure 10 is still far from Fortran. Additional changes of coordinates are needed to generate proper Fortran declarations and loops.

* The Manhattan norm of a vector is the sum of the absolute values of its coordinates, i.e., the l_1 norm.

† That is, the closest on the first dimension, and then the closest on the second dimension, and so on, dimension per dimension, till one processor is determined.

‡ For a vector v , let $|v|$ denotes its dimension.

4.1 Enumeration of Iterations

A general method to enumerate local iterations is described below. It is based on (1) solving HPF and loop equations, both equalities (Equations 2, 3, 6) and inequalities (Equations 4, 5), on (2) searching a good lattice basis to scan the local iterations in an appropriate order using p as a parameter since each processor knows its own identity, and (3) on using linear transformations to switch from the user-visible frame to the local frame.

In this section, a change of frame is computed based on the available equalities to find a dense polyhedron that can be scanned efficiently with standard techniques. The change-of-frame computation uses two Hermite forms to preserve the order of the (l, c) variables which are related to the allocation scheme, for benefiting from cache effects.

Simplifying Formalism

The subscript function S and loop bounds are affine. All object declarations are normalized with 0 as the lower bound. The HPF array mapping is also normalized: $\Pi = I$. Under these assumptions and according to Section 2, HPF declarations and Fortran references are represented by the following sets of equations:

$$\left. \begin{array}{l} \text{alignment (2)} \quad Rt = \quad Aa + s_0 \\ \text{distribution (3)} \quad t = \quad CPc + Cp + l \\ \text{affine reference (6)} \quad a = \quad Si + a_0(n) \end{array} \right\} \quad (7)$$

where p is the processor vector, l the local offset vector, c the cycle vector, a the accessed element, i the iteration, and n the parameters.

Problem Description

Let x be the following set of variables:

$$x = (l, c, a, t, i) \quad (8)$$

They are needed to describe array references and iterations. The order of the l and c variables chosen here should be reflected by the local allocation in order to benefit from cache effects. This order best suits cyclic distributions, as discussed in [33], because there should be more cycles (along c) than block elements (along l) in such cases. Putting the block offset after the cycle number for a given dimension would lead to larger inner loops for more block distributions. Note that the derivations discussed in this section are independent of this order. Equation 7 can be rewritten as

$$Fx = f_0(n, p) \quad (9)$$

with

$$F = \begin{pmatrix} 0 & 0 & A & -R & 0 \\ I & CP & 0 & -I & 0 \\ 0 & 0 & -I & 0 & S \end{pmatrix} \text{ and} \quad (10)$$

$$f_0(n, p) = \begin{pmatrix} s_0 \\ -Cp \\ a_0(n) \end{pmatrix}$$

For instance, in the Chatterjee et al. [17] example in Figure 7, Equation 10 is:

$$\begin{pmatrix} 0 & 0 & 1 & -1 & 0 \\ 1 & 16 & 0 & -1 & 0 \\ 0 & 0 & -1 & 0 & 3 \end{pmatrix} \begin{pmatrix} l \\ c \\ a \\ t \\ i \end{pmatrix} = \begin{pmatrix} 0 \\ -4p \\ 0 \end{pmatrix}$$

The set of interest is the lattice $\mathcal{F} = \{x | Fx = f_0(n, p)\}$. The goal of this section is to find a parametric solution for each component of x . This allows each processor p to scan its part of \mathcal{F} with minimal control overhead.

It is important to note that F is of full row-rank. The rows contain distinct variables that ensure their independence one from the other. First, the alignment equations are composed of independent equalities (they differ from a variables); this is also the case for the distribution and reference equations because of the l and t variables. Second, the equalities composing the distribution equations are the only to contain l variables, thus they are independent from the alignment and reference. Third, the alignment and reference equations are separated by the template variables through R . Since all equalities are independent, F is a full row-rank. Note that HPF restrictions about alignment are not exploited.

If additional equalities are taken into account, such as those arising from simple block ($c_i = 0$) or cyclic ($l_i = 0$) distributions, they can be used to remove the variables from the equation and do not actually change this property. Other degenerated cases may arise (for instance, there is only one processor on a given dimension), but they are not actually discussed here. Our general scheme can directly take advantage of such extra information to optimize the generated code by including it as additional equalities and inequalities.

Parametric Solution of Equations

Lattice \mathcal{F} is implicitly defined but a parametric definition is needed to enumerate its elements for a given processor p . There are two kinds of parameters that must be set apart. First, the constant unknown parameters n and local processor id p , the values of which are known at run-time on each processor. Second, the parameters of interest to us, i.e., those that have to be enumerated or instantiated on each processor to scan the integer solutions to the HPF equations, are the variables in vector x .

A Hermite form [34] of integer matrix F is used to find the parameters. This form associates to F (an $n \times m$ matrix with $m \geq n$) and three matrices H , P , and Q , such that $H = PFQ$. P is a permutation (a square $n \times n$ matrix), H an $n \times m$ lower triangular integer matrix, and Q an $m \times m$ unimodular change of basis. Since F is of full row-rank, no permutation is needed: $P = I$ and $H = FQ$ (a). By definition, H is a lower triangular matrix, and thus can be decomposed as $H = (H_L 0)$, where H_L is an $n \times n$ integer triangular square matrix. We know that $|H_L| \in \{-1, 1\}$. Indeed, H_L is of full rank (as F) and the column combinations performed by the Hermite form computation put unit coefficients on the H_L diagonal. This is ensured since independent unit coefficients appear in each row of F . Thus H_L is an integer triangular unimodular matrix, and has an integral inverse.

Now we can use Q as a change of basis between new variables v and x , with $v = Q^{-1}x$ (b). Vector v can also be decomposed like H in two components: $v = (v_0, v')$, where $|v_0|$ is the rank of H . Using (a) and (b), Equation 9 can be rewritten as:

$$Fx = FQQ^{-1}x = Hv = (H_L 0)(v_0, v') = H_L v_0 = f_0(n, p)$$

where v_0 is a parametric solution at the origin which depends on the run-time value of the n and p parameters. Thus we have $v_0(n, p) = H_L^{-1}f_0(n, p)$. By construction, H does not constrain v' , and Q can also be decomposed like v as $Q = (Q_0 F')$. Lattice \mathcal{F} can be expressed in a parametric linear way:

$$x = Qv = (Q_0 F')(v_0(n, p), v') = Q_0 v_0(n, p) + F'v'$$

and with $x_0(n, p) = Q_0 v_0(n, p) = Q_0 H_L^{-1} f_0(n, p)$:

$$\mathcal{F} = \{x | \exists v' \text{ s.t. } x = x_0(n, p) + F'v'\} \quad (11)$$

We have switched from an implicit (Equation 9) description of a lattice on x to an explicit (Equation 11) one through F' and v' . Note that F' is of full column-rank.

Cache-Friendly Order

As mentioned earlier, the allocation is based somehow on the (l, c) variables. The order used for the enumeration should reflect as much as possible the one used for the allocation, so as to benefit from memory locality. Equation 11 is a parametric definition of x . However, the new variables v' are not necessarily ordered as the variables of interest to us. The aim of this paragraph is to reorder the variables as desired, by computing a new transformation based on the Hermite form of F' . Let $H' = P'F'Q'$ be this form. Let Q' define a new basis:

$$u = Q'^{-1}v' \quad (12)$$

$$x - x_0(n, p) = F'v' = P'^{-1}H'Q'^{-1}v' = P'^{-1}H'u \quad (13)$$

If $P' = I$, the new generating system of \mathcal{F} is based on a triangular transformation between x and u (Equation 13). Since H' is a lower triangular integer matrix, the variables of u and of x simply correspond one to the other. Knowing this correspondence allows us to order the variable u components to preserve locality of accesses. If $P' \neq I$, the variables are shuffled and some cache effect may be lost. However, we have never encountered such an example.

Code Generation

Variable u defining x in lattice \mathcal{F} and polyhedron \mathcal{K} is easy to scan with DO loops. The constraints defining polyhedron \mathcal{K} are coming from declarations, HPF directives, and normalized loop bounds. They are:

$$0 \leq l < C.l, \quad 0 \leq t < T.l, \quad 0 \leq a < D.l, \quad Li \leq b_0(n)$$

Let $Kx \leq k_0(n)$ be these constraints on x . Using Equation 13, the constraints on u can be written $K(x_0(n, p) + P'^{-1}H'u) \leq k_0(n)$, i.e., $K'u \leq k'_0(n, p)$, where $K' = KP'^{-1}H'$ and $k'_0(n, p) = k_0(n) - Kx_0(n, p)$.

Algorithms presented in [28] or others [30, 35–43] can be used to generate the loop nest enumerating the local iterations. When S is of rank $|a|$, optimal code is generated because no projections are required. Otherwise, the quality of the control overhead depends on the accuracy of integer projections [44] but the correctness does not.

Correctness

The correctness of this enumeration scheme stems from (1) the exact solution of integer equations using the Hermite form to generate a parametric enumeration, from (2) the unimodularity of the transformation used to obtain a triangular enumeration, and from

(3) the independent parallelism of the loop which allows any enumeration order.

4.2 Symbolic Resolution

The previous method can be applied in a symbolic way, if the dimensions are not coupled, and thus can be dealt with independently, as array sections in [17, 45, 46]. Equation 7 then becomes for a given dimension:

$$\left. \begin{array}{l} \text{alignment} \quad t = \alpha a + t_0 \\ \text{distribution} \quad t = \pi\gamma c + \gamma p + l \\ \text{affine reference} \quad a = \sigma i + a_0 \end{array} \right\} \quad (14)$$

where π is the number of processors (a diagonal coefficient of P) and γ the block size (i.e., a diagonal coefficient in C). In order to simplify the symbolic resolution, variables a and t are eliminated. The matrix form of the system is then $f_0 = (\alpha a_0 + t_0 - \gamma p)$, $x = (l, c, i)$, and $F = (1 \quad \pi\gamma \quad -\alpha\sigma)$.

The Hermite form is $H = (1 \quad 0 \quad 0) = PFQ$, with $P = I$ and:

$$Q = \begin{pmatrix} 1 & -\pi\gamma & \alpha\sigma \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}$$

Let g , μ , and ω be such that g is $\gcd(\pi\gamma, \alpha\sigma)$ and $g = \pi\gamma\mu - \alpha\sigma\omega$ is the Bezout identity. The Hermite form H' of the two rightmost columns of Q noted $F'(H' = P'F'Q')$ is such that $x - x_0 = H'u$ with:

$$H' = \begin{pmatrix} -g & 0 \\ \mu & \frac{\alpha\sigma}{g} \\ \omega & \frac{\pi\gamma}{g} \end{pmatrix}, x_0 = \begin{pmatrix} \alpha a_0 + t_0 - \gamma p \\ 0 \\ 0 \end{pmatrix},$$

$$P' = I \text{ and } Q' = \begin{pmatrix} \mu & \frac{\alpha\sigma}{g} \\ \omega & \frac{\pi\gamma}{g} \end{pmatrix}$$

This links the two unconstrained variables u to the elements x of the original lattice \mathcal{F} . Variables a and t can be retrieved using Equation 14.

The translation of constraints on x to u gives a way to generate a loop nest to scan the polyhedron. Under the assumption $\alpha > 0$ and $\sigma > 0$, assuming that

$$\begin{aligned} &\forall p \in [0 \dots \pi - 1] \\ \text{do } u_1 &= \frac{\alpha a_0 + t_0 - \gamma p + \gamma + g - 2}{g}, \frac{\alpha a_0 + t_0 - \gamma p}{g} \\ \text{do } u_2 &= \frac{-\omega u_1 + \frac{\pi\gamma}{g} - 1}{\frac{\pi\gamma}{g}}, \frac{\beta - 1 - \omega u_1}{\frac{\pi\gamma}{g}} \\ &x = H'u + x_0(p) \end{aligned}$$

FIGURE 11 Generated code.

loop bounds are rectangular* and using the constraints in \mathcal{X} :

$$0 \leq l < \gamma, 0 \leq a < s, 0 \leq i < \beta$$

the constraints on x (the a one is redundant if the code is correct) are:

$$\begin{pmatrix} -1 & 0 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & -1 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} l \\ c \\ i \end{pmatrix} \leq \begin{pmatrix} 0 \\ \gamma - 1 \\ 0 \\ \beta - 1 \end{pmatrix}$$

and can be translated as constraints on u :

$$\begin{pmatrix} g & 0 \\ -g & 0 \\ -\omega & -\frac{\pi\gamma}{g} \\ \omega & \frac{\pi\gamma}{g} \end{pmatrix} \begin{pmatrix} u_1 \\ u_2 \end{pmatrix} \leq \begin{pmatrix} \alpha a_0 + t_0 - \gamma p \\ -\alpha a_0 - t_0 + \gamma p + \gamma - 1 \\ 0 \\ \beta - 1 \end{pmatrix}$$

The resulting generic SPMD code for an array section is shown in Figure 11. As expected, the loop nest is parameterized by p , the processor identity. Integer divisions with positive remainders are used in the loop bounds.

4.3 HPF Array Allocation

The previous two sections can also be used to allocate local parts of HPF distributed arrays. A loop nest referencing a whole array through an identity subscript function ($S = I, a_0 = 0$) serves as a basis for the allocation. The dense polyhedron obtained by the changes of bases for the enumeration purpose can be used to store the required elements, since local

* This assumption is, of course, not necessary for the general algorithm described in Section 4.1, but met by array sections.

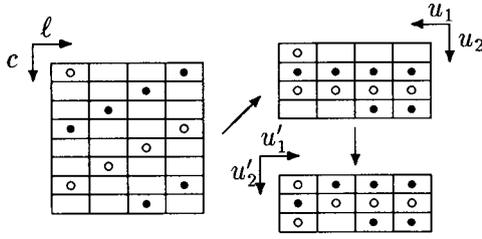


FIGURE 12 Packing of elements.

iterations and local elements are strictly equivalent. Thus, the constraints on local iterations can be reused as constraints on local elements.

However, Fortran array declarations are based on Cartesian sets and are not as general as Fortran loop bounds. General methods have been proposed to allocate polyhedra [47]. For our particular case, it is possible to add another change of frame to better fit Fortran array declaration constraints and to reduce the amount of allocated memory at the expense of the access function complexity.

The local array elements are packed onto local memories dimension by dimension. The geometrical intuition of the packing scheme for the [17] example is shown in Figure 12. The basic idea is to remove the regular holes due to the alignment stride by allocating the dense u space on each processor. The “○” and “●” are used to support the geometrical intuition of the change of frame. The first grid is the part of the template local to the processor, in the cycle and offset coordinate system. The second grid shows the same grid through transformation from x to u . The third one is the final packed form, which could be used if space must not be wasted, but at the expense of a complex access function.

An array dimension can be collapsed or distributed. If the dimension is collapsed, no packing is needed, so the initial declaration is preserved. If the dimension is aligned, there are three corresponding coordinates (p, l, c) . For every processor p , local (l, c) pairs have to be packed onto a smaller area. This is done by first packing up the elements along the columns, then removing the empty ones. Of course, a dual method is to pack first along the rows, then remove the empty ones. This last method is less efficient for the example shown in Figure 12 since it would require 16 (8×2) elements instead of 12 (3×4). The two packing schemes can be chosen according to this criterion. Other issues of interest are the induced effects for the cache behavior and the enumeration costs. Formulae are derived below to perform these packing schemes.

Packing of the Symbolic Resolution

Let us consider the result of the above symbolic resolution, when the subscript expression is the identity ($\sigma = 1, a_0 = 0$). The equation between u and the free variables of x is obtained by selecting the triangular part of H' , i.e., its first rows. If $H'' = (I0)H'$ is the selected submatrix, we have $x' - x'_0 = H''u$, i.e.,

$$\begin{pmatrix} l - t_0 + \gamma p \\ c \end{pmatrix} = \begin{pmatrix} -g & 0 \\ \mu & \frac{\alpha}{g} \end{pmatrix} \begin{pmatrix} u_1 \\ u_2 \end{pmatrix}$$

Variables σ and a_0 were substituted by their values in the initial definition of H' .

Variables (u_1, u_2) could be used to define an efficient packing, since holes are removed by the first change of basis (Fig. 12). In order to match simply and closely the l, c space, the sign of u_1 (linked to l) can be changed, and the vector should be shifted so that the first local array element is mapped near $(0, 0)$.

Some space may be wasted at the beginning and end of the allocated space. For a contrived example (with very few cycles), the wasted space can represent an arbitrary amount on each dimension. Let us assume that two thirds of the space is wasted for a given dimension. Thus, the memory actually used for an array with three of these dimensions is $1/27$, and $26/27$ of the allocated memory is wasted. If such a case occurs, the allocation may be skewed to match a rectangle as closely as possible. This may be done if space is at a premium and if more complex, nonaffine access functions are acceptable. The improved and more costly scheme is described in the next section.

Allocation Basis

Let M be the positive diagonal integer matrix composed of the absolute value of the diagonal coefficients of H'' .

$$u' = \text{alloc}(u) = \lfloor M^{-1}(x' - x'_0) \rfloor = \lfloor M^{-1}H''u \rfloor \quad (15)$$

M provides the right parameters to perform the proposed packing scheme. To every u a vector u' is associated through Equation 15. This formula introduces an integer division. Let's show why u' is correct and induces a better mapping of array elements on local memories than u . Since H'' is lower triangular, Equation 15 can be rewritten

$$\forall i \in [1 \dots |u|], u'_i = \frac{h_{i,i}}{|h_{i,i}|} u_i + \frac{\sum_{j=1}^{i-1} h_{i,j} u_j}{|h_{i,i}|} \quad (16)$$

$$\begin{aligned}
 &0 \leq p \leq \pi - 1, \quad 0 \leq \ell \leq \gamma - 1 \\
 &\left\lfloor \frac{t_0}{\gamma\pi} \right\rfloor \leq c \leq \left\lfloor \frac{\alpha(s-1) + t_0}{\gamma\pi} \right\rfloor \\
 \text{array A' (0:} &\left\lfloor \frac{\max(c) - \min(c) + 1}{\frac{\alpha}{g}} \right\rfloor, 0: \left\lfloor \frac{\gamma}{g} \right\rfloor)
 \end{aligned}$$

FIGURE 13 Local new declaration.

Function $\text{alloc}(u)$ is bijective. $\text{Alloc}(u)$ is injective: If u^a and u^b are different vectors, and i is the first dimension for which they differ, Equation 16 shows that u_i^a and u_i^b will also differ. The function is also surjective, since the property allows construction of a vector that matches any u' by induction on i .

Array Declaration

Two of the three components of x' , namely p and l , are explicitly bounded in \mathcal{K} . Implicit bounds for the third component, c , are obtained by projecting \mathcal{K} on c . These three pairs of bounds, divided by M , are used to declare the local part of the HPF array. Figure 13 shows the resulting declaration for the local part of the array, in the general symbolic case, for one dimension of the array. Bounds $\min(c)$ and $\max(c)$ are computed by Fourier projection.

The packing scheme induced by u' is better than the one induced by u because there are no nondiagonal coefficients between u' and x' that would introduce a waste of space, and u' is as packed as u because contiguity is preserved by Equation 16. The row packing scheme would have been obtained by choosing (c, l) instead of (l, c) for x' . Different choices can be made for each dimension. The access function requires an integer division for the reduced memory allocation. Techniques have been suggested to handle divisions by invariant integers efficiently [48] that could help reduce this cost. Also, because contiguity is preserved, only one division per column is required to compute the base location. These packing schemes define two parameters (u'_1, u'_2) to map one element of one dimension of the HPF array to the processor's local part. The local array declaration can be linearized with the u'_3 dimension first, if the Fortran limit of seven dimensions is exceeded.

4.4 Properties

The proposed iteration enumeration and packing scheme has several interesting properties. It is compatible with efficient cache exploitation and overlap

analysis. Moreover, some improvements can statically enhance the generated code.

According to the access function, the iteration enumeration order and the packing scheme in Figure 11 can be reversed via loop u_2 direction in order that accesses to the local array are contiguous. Thus, the local cache and/or prefetch mechanisms, if any, are efficiently used.

The packing scheme is also compatible with overlap analysis techniques [1]. Local array declarations are extended to provide space for border elements that are owned by neighbor processors, and to simplify accesses to nonlocal elements. The overlap is induced by relaxing constraints on l , which is transformed through the scheme as relaxed constraints on u'_2 . This allows overlaps to be simply considered by the scheme. Moreover, a translated access in the original loop leading to an overlap is transformed into a translated access in the local SPMD-generated code.

For example, using the HPF array mapping of Figure 7:

```
align B with A
A(1:42) = B(0:41)
```

has a $View_B$ area represented in grey on Figure 14 in the unpacked template space and in the local packed array space. The local array B' can be extended by overlap to contain the grey area. Thus, constraint $0 \leq l \leq 3$ in *Own* becomes $-3 \leq l \leq 3$, expressing the size 3 overlap on the left.

The generic code proposed in Figure 11 can be greatly improved in many cases. Integer division may be simplified, performed efficiently with shifts, or even removed by strength reduction. Node splitting and loop invariant code motion should be used to reduce the control overhead. Constraints may also be simplified, for instance, if the concerned elements just match a cycle. Moreover, it is possible to generate the loop nest directly on u' , when u is not used in the loop body.

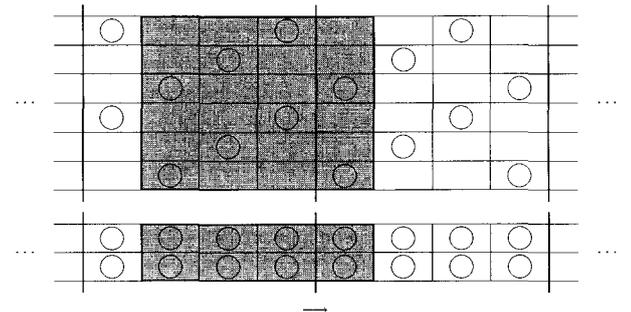


FIGURE 14 Overlap before and after packing.

$$\alpha = 3, t_0 = 0, \sigma = 3, a_0 = 0, \gamma = 4, \pi = 4, \beta = 15$$

```

 $\forall p \in [0 \dots 3]$ 
do  $u'_1 = 0, 3$ 
   $lb_2 = \frac{16p+4u'_1+8}{9}, ub_2 = \frac{16p+4u'_1+7}{9}$ 
   $lb'_2 = \frac{9lb_2-4u'_1-16p}{3}$ 
  do  $u'_2 = lb'_2, lb'_2 + (ub_2 - lb_2)$ 
  ...

```

FIGURE 15 Optimized code.

For the main example in [17], such transformations produce the code shown in Figure 15.

In the general resolution (Section 4.1), the cycle variables c were put after the local offsets l . The induced inner loop nest is then on c . It may be interesting to exchange l and c in x' when the block size is larger than the number of cycles: The loop with the larger range would then be the inner one. This may be useful when elementary processors have some vector capability.

4.5 Allocation of Temporaries

Temporary space must be allocated to hold nonlocal array elements accessed by local iterations. For each loop L and array X , this set is inferred from $Compute_L(p)$ and from subscript function S_X . For instance, references to array Y in Figure 16 require local copies.

```

implicit integer (a-z)
real X(0:24,0:18), Y(0:24,0:18)
!HPF$ template T(0:80,0:18)
!HPF$ align X(I,J) with T(3*I,J)
!HPF$ align Y with X
!HPF$ processors P(0:3,0:2)
!HPF$ distribute T(cyclic(4),block) onto P

read *, m, n

!HPF$ independent(I,J)
do I = m, 3*m
  do J = 0, 2*I
    X(2*I,J) = Y(2*I,J) + Y(2*I+1,J) + I - J
  enddo
enddo

!HPF$ independent(I)
do I = 0, n
  Y(I,I) = I
enddo

```

FIGURE 16 Code example.

Because of Y 's distribution, overlap analysis is fine but another algorithm is necessary for other cases.

Consider the generic loop in Figure 9 and assume that local elements of Y cannot efficiently be stored using overlap analysis techniques. First, if all or most of local elements of the lhs reference are defined by the loop nest, and if S_Y has the same rank as S_X , temporaries can be stored as X . If, furthermore, the access function S_X uses one and only one index in each dimension, the resulting access pattern is still HPF like, so the result of Section 4.3 can be used to allocate the temporary array. Otherwise, another multistage change of coordinates is required to allocate a minimal area.

The general idea of the temporary allocation scheme is to reduce the number of necessary dimensions for the temporary array via a Hermite transformation, then to use this new basis for declaration, or the compressed form to reduce further the allocated space.

The set of local iterations, $Compute(p)$, is now defined by a new basis and new constraints, such that $x - x_0 = P'^{-1}H'u$ (Equation 13) and $K'u \leq k'$. Some rows of $P'^{-1}H'$ define iteration vector i , which is part of x (Equation 8). Let H'_i be this submatrix: $i = H'_i u$.

Let $H_Y = P_Y S_Y Q_Y$ be S_Y 's Hermite form. Let $v = Q_Y^{-1}i$ be the new parameters, then $S_Y i = P_Y^{-1}H_Y v$. Vector v can be decomposed as (v', v'') where $v' = \Pi v$ contributes to the computation of i and v'' belongs to H_Y 's kernel. If H_{Yl} is a selection of the nonzero columns of $H_Y = (H_{Yl}, 0)$, then we have:

$$a_Y - a_{Y0} = S_Y i = P_Y^{-1} H_{Yl} v'$$

$$v' = \Pi Q_Y^{-1} i$$

and by substituting i :

$$v' = \Pi Q_Y^{-1} H'_i u$$

It is sometimes possible to use x' instead of v' . For instance, if the alignments and the subscript expressions are translations, i is an affine function of x' and v' simply depends on x' . The allocation algorithm is based on u but can also be applied when x' is used instead.

Since the local allocation does not depend on the processor identity, the first $|p|$ components of u should not appear in the access function. This is achieved by decomposing u as (u', u'') with $|u'| = |p|$ and $\Pi Q_{\bar{y}}^{-1} H_i$ in (F_p, F_y) such that:

$$v' = F_p u' + F_y u''$$

and the local part v'' is introduced as:

$$v'' = F_y u''$$

Then the first solution is to compute the amount of memory to allocate by projecting constraints \mathcal{K} onto v'' . This is always correct but may lead to a waste of space because periodic patterns of accesses are not recognized and holes are allocated. In order to reduce the amount of allocated memory, a heuristic, based on large coefficients in F_y and constraints \mathcal{K} , is suggested.

Each dimension, i.e., component of v'' , is treated independently. Let F_y^i be a line of F_y and y the corresponding component of v'' : $y = f v''$. A change of basis G (not related to the previous allocation operation) is applied to v'' to reduce f to a simpler equivalent linear form f' where each coefficient appears only once and where coefficients are sorted by increasing order. For instance:

$$f = (16 \quad 3 \quad -4 \quad -16 \quad 0)$$

is replaced by

$$f = f' G$$

$$G = \begin{pmatrix} 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & -1 & 0 & 0 \\ 1 & 0 & 0 & -1 & 0 \end{pmatrix}$$

$$f' = (3 \quad 4 \quad 16)$$

and v'' is replaced by $w = G v''$. Constraints \mathcal{K} on u are rewritten as constraints \mathcal{W} on w by removing the constraints due to processors and by using projections and G .

input: $\{y = f.w, W w \leq 0\}$

output: $\{y' = f'.w\}$

initial system: $f^0 = f$

for $i = 1 \dots (|w| - 1)$

$g_i = \gcd(\{f_j^{i-1}, j > i\})$

$T_i = \{t | t = \sum_{j \leq i} f_j^{i-1} w_j \wedge W w_j \leq 0\}$

$s_i = \max_{t \in T_i} t - \min_{t \in T_i} t$

if $(s_i \geq g_i)$

then $f^i = f^{i-1}$

else $\forall j \leq i, f_j^i = f_j^{i-1}$ and $\forall j > i, f_j^i = \frac{f_j^{i-1}}{g_i} s_i$

end for

$f' = f^{|w|-1}$

FIGURE 17 Heuristic to reduce allocated space.

Linear form f is then processed by the heuristic shown in Figure 17. It reduces the extent by $\frac{s_i}{g_i}$ at each step, if the constraints W show that there is a hole. A hole exists when g_i is larger than the extent of the partial linear form being built, under constraints W . Linearity of access to temporary elements is preserved.

This scheme is correct if and only if a unique location y' is associated to each y . Further insight on this problem, the minimal covering of a set by interval congruences, can be found in [49, 50].

4.6 Data Movements

The relationships between the bases and frames defined in the previous sections are shown in Figure 18. Three areas are distinguished. The top area contains user-level bases for iterations, i , and array elements,

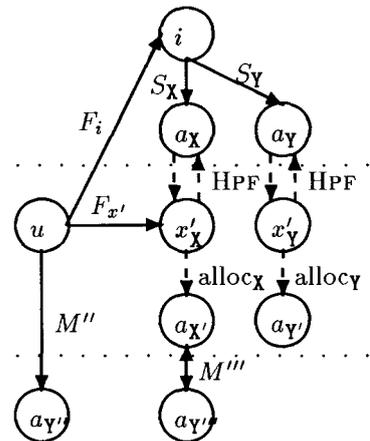


FIGURE 18 Relationships between frames.

a_x , a_y , and so on. The middle area contains the bases and frames used by the compiler to enumerate local iterations, u , and to allocate local parts of HPF arrays, a'_x , a'_y , and so on as well as the universal bases, x'_x and x'_y , used to define the lattices of interest, \mathcal{F} and \mathcal{F}' . The bottom area shows new bases used to allocate temporary arrays, a''_y and a'''_y , as defined in Section 4.5.

Solid arrows denote affine functions between different spaces and dashed arrows denote possible non-affine functions built with integer divide and modulo operators. A quick look at these arrows show that u , which was defined in Section 4.1, is the central basis of the scheme, because every other coordinate can be derived from u along one or more paths.

Two kinds of data exchanges must be generated: updates of overlap areas and initializations of temporary copies. Overlap areas are easier to handle because the local parts are mapped in the same way on each processor, using the same alloc function. Consider array X in Figure 18. The send and receive statements are enumerated from the same polyhedron, up to a permutation of p and p' , with the same basis u . To each u corresponds only one element in the user space, a_x , by construction* of u . To each a_x corresponds only one a'_x on each processor.† As a result, data exchanges controlled by loop on u enumerate the same element at the same iteration.

Because the allocation scheme is the same on each processor, the inner loop may be transformed into a vector message if the corresponding dimension is contiguous and/or the send/receive library supports constant but nonunit strides. Block copies of larger areas also are possible when alloc is affine, which is not the general case from a theoretical point of view but should very often happen in real applications.

Temporary arrays like Y'' and Y''' are more difficult to initialize because there is no such identity between their allocation function as local part of an HPF array, alloc_y , and their allocation functions as temporaries, $\text{t_alloc}_{y''}$ and $\text{t_alloc}_{y'''}$. However, basis u can still be used.

When the temporary is allocated exactly like the master array, e.g., Y''' and X' , any enumeration of elements u in \mathcal{U} enumerates every element a'''_y once and only once because the input loop is assumed independent. On the receiver side, a'''_y is directly derived from u . On the sender side, a_y is computed as $S_y F_y u$ and a nonaffine function is used to obtain x'_y and

a'_y . Vector messages can be used when every function is affine, because a constant stride is transformed into another constant stride, and when such transfers are supported by the target machine. Otherwise, calls to packing and unpacking routines can be generated.

When the temporary is allocated with its own allocation function M , it is more difficult to find a set of u enumerating exactly once its elements. This is the case for copy Y'' in Figure 18. Elements of Y'' , a''_y , must be used to compute one related u among many possible ones. This can be achieved by using a pseudo-inverse of the access matrix M :

$$u = M'(MM')^{-1}a''_y \quad (17)$$

Vector u is the rational solution of equation $a''_y = Mu$ which has a minimum norm. It may well not belong to \mathcal{U} , the polyhedron of meaningful u which is linked to a user iteration. However, M was built to have the same kernel as $S_y F_y$. Thus, the same element a_y is accessed as it would be by a regular u . Since only linear transformations are applied, these steps can be performed with integer computations by multiplying Equation 17 with the determinant of MM' and by dividing the results by the same quantity. The local address, a'_y , is then derived as above. This proves that sends and receives can be enumerated by scanning elements a''_y .

5 EXAMPLES

The different algorithms presented in the previous section were used to distribute the contrived piece of code of Figure 16, using functions of a linear algebra library developed at École des mines de Paris.

This is an extension of the example in [17] showing that allocation of HPF arrays may be nontrivial. The reference to X in the first loop requires an allocation of X' and the computation of new loop bounds. It is not a simple case because the subscript function is not the identity and because a `cyclic(4)` and `block` distribution is specified. The two references to Y in the first loop are similar but they imply some data exchange between processors and the allocation of an overlap area in Y' . Furthermore, the values of I and J are used in the computation, the iteration domain is nonrectangular, and is parameterized with m .

The second loop shows that the `Compute(p)` set may have fewer dimensions than the array referenced and that fewer loops have to be generated. The output code is too long to be printed here. Interesting excerpts are shown in Figures 19 and 20 and commented on below.

* As explained in Section 4.3, allocation is handled as a special case by choosing the identity for S_x .

† Elements in overlap area are allocated more than once, but on different processors.

5.1 HPF Declarations

HPF declarations are already normalized. Templates are all used, thus $\Pi = I$. Other key parameters are:

$$P = \begin{pmatrix} 4 & 0 \\ 0 & 3 \end{pmatrix}, T = \begin{pmatrix} 80 & 0 \\ 0 & 18 \end{pmatrix}, C = \begin{pmatrix} 4 & 0 \\ 0 & 7 \end{pmatrix}$$

$$A_X = A_Y = \begin{pmatrix} 3 & 0 \\ 0 & 1 \end{pmatrix}, R_X = R_Y = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix},$$

$$t_{0X} = t_{0Y} = \begin{pmatrix} 0 \\ 0 \end{pmatrix}$$

5.2 Allocation of X'

Array X is referenced only once, as an lhs. Thus, its local part is directly derived from the set $Own_x(p)$ using the algorithm described in Section 4.3.

$$Own_x(p) = \{(x_1, x_2) | t_1 = 3x_1, t_2 = x_2$$

$$t_1 = 16c_1 + 4p_1 + l_1, t_2 = 21c_2 + 7p_2 + l_2$$

$$0 \leq x_1 \leq 24, 0 \leq x_2 \leq 18,$$

$$0 \leq p_1 \leq 3, 0 \leq p_2 \leq 2,$$

$$0 \leq l_1 \leq 3, 0 \leq l_2 \leq 6,$$

$$0 \leq t_1 \leq 80, 0 \leq t_2 \leq 18\}$$

The equations are parametrically solved and x is replaced by u in the constraints according to the following change of basis:

$$u_1 = p_1, u_2 = p_2, u_3 = c_1, 3u_5 - 16u_3 - 4u_1 = l_1$$

$$u_4 = c_2, u_5 = x_1, u_6 = x_2, u_6 - 21u_4 - 7u_2 = l_2$$

Then, Frame u is replaced by Frame u' (represented with the w variables in the Fortran code) to avoid nondiagonal terms in the constraints. Frame u' is linked to the user basis x by:

$$u'_1 = p_1, u'_2 = p_2, u'_3 = c_1, u'_4 = c_2, 3u'_5 = l_1, u'_6 = l_2$$

and constraints \mathcal{K} are projected on each component of u' to derive the declaration bounds of X' :

$$0 \leq u'_1 \leq 3, 0 \leq u'_2 \leq 2, 0 \leq u'_3 \leq 4$$

$$0 \leq u'_4 \leq 0, 0 \leq u'_5 \leq 1, 0 \leq u'_6 \leq 6$$

Coordinates u'_1 and u'_2 are the processor coordinates. The other coordinates, u'_3 to u'_6 , are the local array coordinates. Note that since we have a block distribution in the second dimension, $u'_4 = 0$. The Fortran declaration of X' is shown in Figure 19 as `x_hpf`. 10×7 elements are allocated while 7×7 would be enough. The efficiency factor may be better in the (p, l, c) basis. It would also be better with less contrived alignments and subscript expressions and/or larger blocks.

5.3 Sending Y'

Before the first loop nest is executed, each processor (p_1, p_2) must send some elements of Y to its right neighbor $(p_1 + 1, p_2)$, with a wraparound for the rightmost processors, $(3, p_2)$. The pairs of communicating processors are not accurately represented by convex constraints and the `dest_p1` and `dest_p2` loop is not as tight as it could be. This suggests that central and edge processors should be handled differently when overlap is observed. A linearized view of the processors could also help to fix this issue.

The bounds for the inner loops are complicated but they could mostly be evaluated iteratively if node splitting and loop invariant code motion are used. Also, the test on `u7` can be simplified as a modulo operation. Finally, `w5` and `w6` can also be evaluated iteratively.

5.4 Local Iterations for the Second Loop

As for the first loop, the second loop cannot be represented as an array section assignment, which makes it harder to handle than the first one. The set of local iteration is defined by:

$$Compute_Y(p) = \{(y_1, y_2) | y_1 = i, y_2 = i,$$

$$t_1 = 3y_1, t_2 = y_2$$

$$t_1 = 16c_1 + 4p_1 + l_1,$$

$$t_2 = 21c_2 + 7p_2 + l_2$$

$$0 \leq i \leq n, 0 \leq y_1 \leq 24, 0 \leq y_2 \leq 18,$$

$$0 \leq p_1 \leq 3, 0 \leq p_2 \leq 2$$

$$0 \leq l_1 \leq 3, 0 \leq l_2 \leq 6, 0 \leq t_1 \leq 80,$$

$$0 \leq t_2 \leq 18\}$$

The loop bounds were retrieved by projection on u . It is probably useless to generate a guard on the processor identity because noninvolved processors have nothing to do and because this does not delay

```

implicit integer (a-z)
c Local declaration of X and Y(0:42,0:42):
real X_hpf(0:6, 0:1, 0:4), Y_hpf(0:6, 0:2, 0:4)

read *, m, n

c For each processor P(p1,p2)

c Send view_Y to compute_X:
do dest_p1 = 0, 3
do dest_p2 = 0, 2
if (dest_p1.ne.p1.and.dest_p2.ne.p2) then
do u6 = max(-15, -3 - 4*p1), min(0, -4*p1)
do u7 = max(0, divide(15 - 5*u6, 16)),
& min(6, divide(24 - 5*u6, 16))
do u8 = 7*p2, min(20, 7*p2)
do u9 = 0, 4
do u11 = max(0, divide(16*u7 + 5*u6,2), m, (2 + 2*dest_p1 + 8*u9)/3),
& min(12, divide(16*u7 + 5*u6,2), 3*m, (2*dest_p1 + 8*u9)/3)
do u12 = max(0, divide(2*u11 - u8, 21), divide(7*dest_p2 - u8, 21)),
& min(0, divide(20 + 7*dest_p2 - u8, 21))
w3 = 3*u7 + u6
w5 = (-u6 - 4*p1)/3
w6 = u8 - 7*p2
send(P(dest_p1,dest_p2), Y_hpf(w6,w5,w3))
enddo
enddo
enddo
enddo
enddo
endif
enddo
enddo

c receive_Y
c ...

c compute_X:
do u4 = 0, 4
do u6 = max(0, m, (2 + 2*p1 + 8*u4)/3),
& min(12, 3*m, (1 + 2*p1 + 8*u4)/3)
do u7 = max(0, 7*p2), min(18, 2*u6, 6 + 7*p2)
w3 = u4
w5 = (6*u6 - 16*u4 - 4*p1)/3
w6 = u7 - 7*p2
i = u6
j = u7
X_hpf(w6, w5, w3) = Y_hpf(w6, w5, w3)
& + i - j + Y_hpf(w6, w5 + 1, w3)
enddo
enddo
enddo

```

FIGURE 19 Output code (excerpts, part 1).

the other ones. The guard may as well be hidden in the inner loop bounds. Experiments are needed to find out the best approach.

Note that an extra loop is generated. The Y diagonal can be enumerated with only two loops and three are generated. This is due to the use of an imprecise projection algorithm, but does not endanger correctness [28]. Further work is needed in this area.

Integer Divide

One implementation of the integer divide is finally shown. The divider is assumed strictly positive, as is the case in all call sites. It is necessary because the Fortran remainder is not positive for negative numbers. It was added to ensure the semantics of the output code. Nevertheless, if we can prove the dividend is positive, we can use the Fortran division.

```

c   compute_Y:
do u4 = 0, 3
do u5 = max(0, 7*p2, (2 + 4*p1 + 16*u4)/3),
&   min(18, 6 + 7*p2, n, (3 + 4*p1 + 16*u4)/3)
w3 = u4
w5 = (3*u5 - 16*u4 - 4*p1)/3
w6 = u5 - 7*p2
i = u5
Y_hpf(w6, w5, w3) = i
enddo
enddo

end

c   Integer divide with positive remainder. Assumes j>0.
integer function divide(i,j)
integer i,j
if(i.ge.0) then
divide = i/j
else
divide = -(i+j-1)/j
endif
end
    
```

FIGURE 20 Output code (excerpts, part 2).

6 RELATED WORK

Techniques to generate distributed code from sequential or parallel code using a uniform memory space have been extensively studied since 1988 [51–53]. Techniques and prototypes have been developed based on Fortran [1, 2, 5, 6, 8, 54–56], C [57–62], or other languages [31, 32, 63–65].

The most obvious, most general, and safest technique is called run-time resolution [51, 52, 63]. Each instruction is guarded by a condition which is only true for processors that must execute it. Each memory address is checked before it is referenced to decide whether the address is local and the reference is executed, whether it is remote and a receive is executed, or whether it is remotely accessed and a send is executed. This rewriting scheme is easy to implement [7], but is very inefficient at run-time because guards, tests, sends, and receives are pure overhead. Moreover, every processor has to execute the whole control flow of the program, and even for a parallel loop, communications may sequentialize the program at run-time [51].

Many optimization techniques have been introduced to handle specific cases. Gerndt introduced overlap analysis in [1] for block distributions. When local array parts are allocated with the necessary overlap and when parallel loops are translated, the instruction guards can very often be moved in the loop bounds and the send/receive statements are globally executed before the local iterations, i.e., the loop nest with run-time resolution is distributed into two loop nests, one

for communications and one for computations. The communication loops can be rearranged to generate vector messages.

Tseng [3] presents lots of additional techniques (message aggregation and coalescing, message and vector message pipelining, computation replication, and collective communication). He assumes affine loop bounds and array subscripts to perform most optimizations. He only handles block and `cyclic(1)` distributions and the alignment coefficient must be 1.

Recent publications tackle any alignment and distribution but often restrict references to array sections. Each dimension is independent of the others as was assumed in Section 4.2.

In Paalvast et al. [33, 66] a technique based on the resolution of the associated Diophantine equations is presented. Row- and column-wise allocation and addressing schemes are discussed. Benkner et al. [67, 68] present similar techniques.

Chatterjee et al. [17] developed a finite state machine (FSM) approach to enumerate local elements. No memory space is wasted and local array elements are ordered by Fortran lexicographic order exactly like user array elements. They are sequentially accessed by `while` loops executing the FSM's, which may be a problem if vector units are available. Moreover, accesses to an auxiliary data structure, the FSM transition map, add to the overhead. Note that the code generated in Figure 11 may be used to compute the FSM. In fact, the lower iteration of the innermost loop is computed by the algorithm that builds the FSM. Kennedy et al. [69–72] and others [17] have suggested improvements to this technique, essentially to compute faster at run-time the automaton transition map. Also multidimensional cases need many transition maps to be handled.

Articles by Stichnoth et al. [46, 74] on the one hand and Gupta et al. [45, 75] and Kaushik et al. [76] on the other hand present two similar methods based on closed forms for this problem. They use array sections but compute some of the coefficients at run-time. Gupta et al. [45, 75] and Kaushik et al. [76] solve the block distribution case and use processor virtualization to handle cyclic distributions. Arrays are densely allocated as in [17] and the initial order is preserved but no formulae are given. Stichnoth uses the dual method for array allocation as in [7], i.e., blocks are first compressed, and the cycle number is used as a second argument.

In [19, 77] polyhedron-based techniques are presented to generate transfer code for machines with a distributed memory. In [37, 78] advanced analyses are used as an input to a code generation phase for

distributed memory machines. Polyhedron scanning techniques are used for generating the code. Two families of techniques have been suggested for that purpose. First, Fourier elimination-based techniques [28, 36, 37, 40–43, 79], and second, parametric integer programming-based methods [30, 35, 38, 39]. In [80], a twofold Hermite transformation is also used to remove modulo indexing from a loop nest. First, variables are added to explicit the modulo computation, then the Hermite computations are used to regenerate simply new loop bounds. While the aim is different, the transformations are very similar to those presented here.

7 CONCLUSION

The generation of efficient SPMD code from an HPF program is not a simple task and, up to now, many attempts have provided partial solutions and many techniques. A translation of HPF directives in affine constraints, avoiding integer division and modulo, was presented in Section 2 to provide a unique and powerful framework for HPF optimizations. A normalization step is applied to reduce the number of objects used in the compilation scheme overviewed in Section 3. New algorithms are presented to (1) enumerate local iterations, (2) allocate local parts of distributed arrays, (3) generate send and receive blocks, and (4) allocate temporaries implied by rhs references. They are based on changes of basis and enumeration schemes. Section 4 shows that problem (2) can be casted as a special case of problem (1) by using an identity subscript function. Problem (3) is an extension of problem (1): A unique reference to an array is replaced by a set of references and the equation used to express the reference is replaced by a set of inequalities. Problem (4) is an extension of problem (2). The set of elements to allocate is no longer the image of a rectangle but the image of an iteration set which can have any polyhedral shape. This shows that all these problems are closely linked.

Although the usual affine assumptions are made for loop bounds and subscript expressions, our compilation scheme simultaneously lifts several restrictions: Array references are not restricted to array sections. General HPF alignments and distributions are supported, and the same algorithms also generate efficient codes for classical block distributions, similar to the ones produced by classical techniques. Memory allocation is almost 100% efficient on large blocks and performs quite well on small ones when strange alignments are used. We believe that this slight memory

waste is more than compensated by the stride-1 vector load, store, send and receive which can be performed on the copies and which are necessary for machines including vector units. These contiguous accesses also perform well with a cache. The overlap analysis and allocation are integrated to the basic allocation scheme. Finally, most computations are performed at compile-time and no auxiliary data structures are used.

Our scheme can also be extended to cope with processor virtualization if the virtualization scheme is expressed with affine constraints. Such a scheme could reuse HPF distribution to map HPF processors on physical processors.

Many partial optimization techniques are integrated in our direct synthesis approach: message vectorization, aggregation [2], and overlap analysis [1]. A new storage management scheme is also proposed. Moreover, other optimization techniques may be applied to the generated code such as vectorization [18], loop invariant code motion [81], and software pipelining [82, 83].

This technique uses algorithms, directly or indirectly, that may be costly, such as Fourier elimination or the simplex algorithm, which have exponential worst-case behaviors. They are used for array region analysis, in the set manipulations, and in the code generation for polyhedron scanning. However, our experience with such algorithms is that they remain practical for our purpose: Polyhedron-based techniques are widely implemented in the PIPS framework [84] where HPFC, our prototype HPF compiler, is developed. First, for a given loop nest, the number of equalities and inequalities is quite low, typically a dozen or less. Moreover, these systems tend to be composed of independent subsystems on a dimension per dimension basis, resulting in a more efficient practical behavior. Second, efficient and highly tuned versions of such algorithms are available, for instance, in the Omega library. Third, potentially less precise but faster program analysis [85–87] can also be used in place of the region analysis.

Polyhedron-based techniques are already implemented in HPFC, our prototype HPF compiler [88], to deal with I/O communications in a host/nodes model [89] and also to deal with dynamic remappings [29] (`realign` and `redistribute` directives). For instance, the code generation times for arbitrary remappings are in the 0.1–2 s range. Future work includes the implementation of our scheme in HPFC, experiments, extensions to optimize sequential loops, to overlap communication and computation, and to handle indirections.

ACKNOWLEDGMENTS

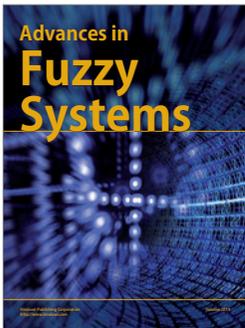
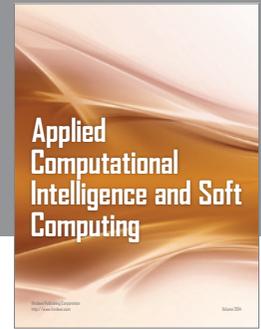
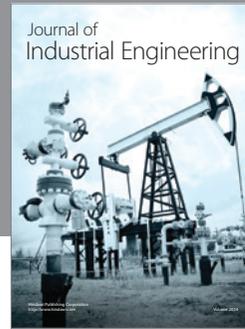
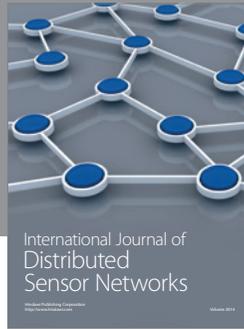
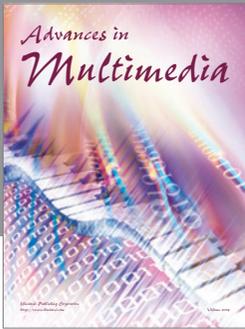
We thank Béatrice Creusillet for her many questions, Pierre Jouvelot for his careful reading and many suggestions, William Pugh for helpful comments and debugging, William (Jingling) Xue for the many remarks which helped improve the article, and the referees for their precise and constructive comments.

REFERENCES

- [1] H. M. Gerndt, "Automatic parallelization for distributed-memory multiprocessing systems," PhD thesis, University of Vienna, 1989.
- [2] S. Hiranandani, K. Kennedy, and C.-W. Tseng, "Compiling Fortran D for MIMD distributed-memory machines," *Commun. ACM*, vol. 35, pp. 66–80, Aug. 1992.
- [3] C.-W. Tseng, "An optimising Fortran D compiler for MIMD distributed memory machines," PhD thesis, Rice University, Houston, TX, Jan. 1993.
- [4] High Performance Fortran Forum, *High Performance Fortran Language Specification (Version 1.1)*. Houston: Rice University, 1994.
- [5] T. Brandes, "Efficient data-parallel programming without explicit message passing for distributed memory multiprocessors," High Performance Computing Center, German National Research Institute for Computer Science, Sankt Augustin, Internal Rep. AHR-92 4, Sept. 1992.
- [6] T. Brandes, "Adaptor: A compilation system for data parallel Fortran programs," High Performance Computing Center, German National Research Institute for Computer Science, Sankt Augustin, Tech. Rep. Aug. 1993.
- [7] F. Coelho, "Étude de la compilation du high performance Fortran," Master's thesis, Université Paris VI, Sept. 1993. Rapport de DEA Systèmes Informatiques, TR EMP E/178/CRI.
- [8] H. Zima and B. Chapman, "Compiling for distributed-memory systems," *Proc. IEEE*, Feb. 1993.
- [9] P. Banerjee, J. A. Chandy, M. Gupta, E. W. Hodges IV, J. G. Holm, A. Lain, D. J. Palermo, S. Ramaswamy, and E. Su, "An overview of the PARADIGME compiler for distributed-memory multicomputers," *IEEE Comput.*, vol. 28, pp. 37–47, Oct. 1995.
- [10] J. M. Levesque, *FORGE 90 and High Performance Fortran*. Applied Parallel Research, Inc., 1992. xHPF77 presentation.
- [11] J. M. Levesque, "Applied parallel research's xHPF system," *IEEE Parallel Distrib. Technol.*, p. 71, Fall 1994.
- [12] J. A. Bircsak, M. R. Bolduc, J. A. Diewald, I. Gale, J. Harris, N. W. Johnson, S. Lee, C. A. Nelson, and C. D. Offner, "Compiling high performance Fortran for distributed-memory systems," Report, Digital Equipment Corp., Oct. 1995. (To be published in the *Digital Techn. J.*)
- [13] C. D. Offner, "Digital's HPF compiler: Meeting the challenge of generating efficient code on a workstation farm," presented at the *NASA Ames Workshop*, 1993.
- [14] M. Gupta, S. Midkiff, E. Schonberg, V. Seshadri, D. Shields, K. Y. Wang, W.-M. Ching, and T. Ngo, "An HPF compiler for the IBM SP2," presented at the Workshop on Compilers for Parallel Computers, Malaga, Spain, 1995.
- [15] L. Dagum, L. Meadows, and D. Miles, "Data parallel direct simulation Monte Carlo in high performance Fortran," *Sci. Prog.* (in press).
- [16] L. F. Meadows, D. Miles, C. Walinsky, M. Young, and R. Touzeau, "The Intel Paragon HPF compiler," Portland Group Inc., Wilsonville, OR, Tech. Rep., June 1995.
- [17] S. Chatterjee, J. R. Gilbert, F. J. E. Long, R. Schreiber, and S.-H. Teng, "Generating local addresses and communication sets for data-parallel programs," *J. Parallel Distrib. Comput.* vol. 26, pp. 72–84, April 1995. Also appeared in PPOPP'93.
- [18] H. Zima and B. Chapman, *Supercompilers for Parallel and Vector Computers*. ACM Press, New York, 1990.
- [19] C. Ancourt and F. Irigoien, "Automatic code distribution," presented at the Third Workshop on Compilers for Parallel Computers, Vienna, 1992.
- [20] R. Triolet, P. Feautrier, and F. Irigoien, "Direct parallelization of call statements," in *Proc. ACM Symp. on Compiler Construction*, 1986.
- [21] F. Irigoien, "Interprocedural analyses for programming environment," in *Environments and Tools for Parallel Scientific Computing*, J. J. Dongara and B. Tourancheau, Eds. North-Holland, Amsterdam: Saint-Hilaire-du-Touvet, NSF-CNRS, 1992, pp. 333–350.
- [22] B. Apvrille, "Calcul de régions de tableaux exactes," in P. Manneback, Ed., *Rencontres Francophones du Parallélisme*. Mons, Belgium: Faculté Polytechnique de Mons, 1994, pp. 65–68.
- [23] B. Apvrille-Creusillet, "Calcul de régions de tableaux exactes," *TSI, Numéro spécial RenPar'6*, May 1995, vol. 14, pp. 585–600.
- [24] B. Apvrille-Creusillet, "Régions exactes et privatisation de tableaux (exact array region analysis and array privatization)," Master's thesis, Université Paris VI, France, Sept. 1994. Available via: <http://www.cri.ensmp.fr/~creusil>.
- [25] B. Creusillet and F. Irigoien, "Interprocedural array region analyses," in *Language and Compilers for Parallel Computing*, 1995, pp. 4.1–4.15.
- [26] B. Creusillet, "IN and OUT array region analyses," presented at the Workshop on Compilers for Parallel Computers, Malaga, Spain, 1995.
- [27] B. Creusillet, "Analyse de flot de données: Régions de tableaux IN et OUT," in P. Manneback, Ed., *Rencontres Francophones du Parallélisme*. Mons, Belgium: Faculté Polytechnique de Mons.
- [28] C. Ancourt and F. Irigoien, "Scanning polyhedra with

- DO loops,” presented at *Symp. on Principles and Practice of Parallel Programming*, Williamsburg, VA, 1991.
- [29] F. Coelho and C. Ancourt. “Optimal compilation of HPF remappings,” CRI, École des mines de Paris, Tech. Rep. A 277, Oct. 1995.
- [30] P. Feautrier, “Parametric integer programming,” *RAIRO Recherche Opérationnelle*, vol. 22, pp. 243–268, Sept. 1988.
- [31] C. Koelbel, P. Mehrotra, and J. Van Rosendale, “Supporting shared data structures on distributed memory architectures,” Purdue University, Tech. Rep. ASD 915, Jan. 1990.
- [32] C. Koelbel and P. Mehrotra, “Compiling global namespace parallel loops for distributed execution,” *IEEE Trans. Parallel Distrib. Systems*, vol. 2, pp. 440–451, Oct. 1991.
- [33] C. van Reeuwijk, H. J. Sips, W. Denissen, and E. M. Paalvast, “Implementing HPF distributed arrays on a message-passing parallel computer system,” Delft University of Technology, Computational Physics Rep. Series, CP-95 006, Nov. 1994.
- [34] A. Schrijver, *Theory of Linear and Integer Programming*. New York: Wiley, 1986.
- [35] J.-F. Collard, P. Feautrier, and T. Risset, “Construction of DO loops from systems of affine constraints,” ENS-Lyon, LIP RR93 15, May 1993.
- [36] W. Kelly and W. Pugh, “A framework for unifying reordering transformations,” Institute for Advanced Computer Studies, University of Maryland, UMIACS-TR-93 134, April 1993.
- [37] S. P. Amarasinghe and M. S. Lam, “Communication optimization and code generation for distributed memory machines, in *ACM SIGPLAN Int. Conf. on Programming Language Design and Implementation*, pp. 126–138, 1993.
- [38] Z. Chamski, “Nested loop sequences: Towards efficient loop structures in automatic parallelization. Res. rep. 2094, INRIA, October 1993,” in *Proc. of the 27th Annual Hawaii Int. Conf. on System Sciences*, 1994, p. 14.
- [39] Z. Chamski, “Fast and efficient generation of loop bounds,” Res. Rep. 2095, INRIA, Oct. 1993.
- [40] H. Le Verge, V. Van Dongen, and D. K. Wilde, “Loop nest synthesis using the polyhedral library,” Publication Interne 830, IRISA, May 1994.
- [41] M. Le Fur, “Scanning parametrized polyhedron using Fourier-Motzkin elimination,” Publication Interne 858, IRISA, Sept. 1994.
- [42] W. Kelly, W. Pugh, and E. Rosser, “Code generation for multiple mappings,” in *5th Symp. on the Frontiers of Massively Parallel Computation*, 1995, p. 332.
- [43] J. Xue, “Constructing do loops for non-convex iteration spaces in compiling for parallel machines,” presented at the *Int. Parallel Processing Symp.*, 1995.
- [44] W. Pugh, “A practical algorithm for exact array dependence analysis,” *CACM*, vol. 35, pp. 102–114, Aug. 1992.
- [45] S. K. S. Gupta, S. D. Kaushik, S. Mufti, S. Sharma, C.-H. Huang, and P. Sadayappan, “On compiling array expressions for efficient execution on distributed-memory machines,” in *Int. Conf. on Parallel Processing*, 1993, p. II-301.
- [46] J. M. Stichnoth, “Efficient compilation of array statements for private memory multicomputers,” School of Computer Science, Carnegie Mellon University, CMU-CS-93 109, Feb. 1993.
- [47] D. K. Wilde and S. Rajopadhye, “Allocating memory arrays for polyhedra,” Res. Rep. 2059, INRIA, July 1993.
- [48] T. Granlund and P. L. Montgomery, “Division by invariant integers using multiplication,” in *ACM SIGPLAN Int. Conf. on Programming Language Design and Implementation*, 1994, p. 61.
- [49] P. Granger, “Analyses sémantiques de congruence,” PhD thesis, École Polytechnique, July 1991.
- [50] F. Masdupuy, “Array indices relational semantic analysis using rational cosets and trapezoids,” PhD thesis, École Polytechnique, Tech. Rep. EMP-CRI A/247, Nov. 1993.
- [51] D. Callahan and K. Kennedy, “Compiling programs for distributed-memory multiprocessors,” *J. Supercomput.* vol. 2, pp. 151–169, 1988.
- [52] R. Mirchandaney, J. S. Saltz, R. M. Smith, D. M. Nicol, and K. Crowley, “Principles of runtime support for parallel processors,” in *ACM Int. Conf. on Supercomputing*, 1988, p. 140.
- [53] H. Peter Zima, H. J. Bast, and H. M. Gerndt, “SUPERB: A tool for semi-automatic MIMD/SIMD parallelization,” *Parallel Comput.*, vol. 6, pp. 1–18, 1988.
- [54] H. M. Gerndt and H. P. Zima, “Optimizing communication in superb, in *CONPAR90*, 1990, p. 300.
- [55] J. Merlin, “Techniques for the automatic parallelisation of “Distributed Fortran 90,” in *SNARC 92 02*, University of Southampton, 1992.
- [56] T. Brandes, “Evaluation of high performance Fortran on some real applications, in *High Performance Computing and Networking*. New York: Springer-Verlag LNCS 797, 1994, pp. 417–422.
- [57] F. André, J.-L. Pazat, and H. Thomas, “Pandore: A system to manage data distribution,” Publication Interne 519, IRISA, Feb. 1990.
- [58] O. Lempel, S. S. Pinter, and E. Turiel, “Parallelizing a C dialect for distributed memory MIMD machines, in *Language and Compilers for Parallel Computing*. New Haven, CT: Springer-Verlag, pp. 369–390, 1992.
- [59] F. André, O. Chéron, and J.-L. Pazat, “Compiling sequential programs for distributed memory parallel computers with Pandore II,” Tech. Rep., IRISA, May 1992.
- [60] M. Le Fur, J.-L. Pazat, and F. André, “Commutative loop nests distribution,” presented at the Workshop on Compilers for Parallel Computers, Delft, 1993.
- [61] F. André, M. Le Fur, Y. Mahéo, and J.-L. Pazat, “Parallelization of a wave propagation application using

- a data parallel computer," Publication Interne 868, IRISA, Oct. 1994.
- [62] M. Le Fur, J.-L. Pazat, and F. André, "An array partitioning analysis for parallel loop distribution," in *Euro-Par'95, Stockholm, Sweden*. New York: Springer Verlag, LNCS 966, 1995, pp. 351–364.
- [63] A. Rogers and K. Pingali, "Process decomposition through locality of reference," in *ACM SIGPLAN Int. Conf. on Programming Language Design and Implementation*, 1989.
- [64] A. Rogers and K. Pingali, "Compiling for distributed memory architectures," June 1992. Sent by the first author when asked for her PhD.
- [65] J. Li and M. Chen, "Compiling communication-efficient programs for massively parallel machines," *IEEE Trans. Parallel Distrib. Systems*, vol. 2, pp. 361–376, July 1991.
- [66] E. M. Paalvast, H. J. Sips, and A. J. van Gemund, "Automatic parallel program generation and optimization from data decompositions, in "1991 Int. Conf. on Parallel Processing—Vol. II: Software," 1991, p. 124.
- [67] S. Benkner, P. Brezany, and H. Zima, "Processing array statements and procedure interfaces in the prepare high performance Fortran compiler," in P. Fritzson, Ed., *5th International Conference on Compiler Construction*. New York and Edinburgh: Springer-Verlag LNCS 786, 1994, pp. 324–338.
- [68] S. Benkner, "Handling block-cyclic distributed arrays in Vienna Fortran. Institute for Software Technology and Parallel Systems, University of Vienna, Tech. Rep. 9, Nov. 1994.
- [69] K. Kennedy, N. Nedeljković, and A. Sethi, "A linear time algorithm for computing the memory access sequence in data-parallel programs," *Sigplan Notices*, vol. 30, pp. 102–111, 1995.
- [70] K. Kennedy, N. Nedeljković, and A. Sethi, "Efficient address generation for block-cyclic distributions," in *ACM Int. Conf. on Supercomputing*, 1995, p. 180.
- [71] S. Hiranandani, K. Kennedy, J. Mellor-Crummey, and A. Sethi, "Advanced compilation techniques for Fortran D," Center for Research on Parallel Computation, Rice University, CRPC-TR 93338, Oct. 1993.
- [72] S. Hiranandani, K. Kennedy, J. Mellor-Crummey, and A. Sethi, "Compilation techniques for block-cyclic distributions," in *ACM Int. Conference on Supercomputing*, 1994.
- [73] A. Thirumalai and J. Ramanujam, "Fast address sequence generation for data-parallel programs using integer lattices," in *Language and Compilers for Parallel Computing*. Columbus: The Ohio State University, pp. 13.1–13.19.
- [74] J. Stichnoth, D. O'Hallaron, and T. Gross, "Generating communication for array statements: Design, implementation, and evaluation," *J. Parallel Distrib. Comput.* vol. 21, pp. 150–159, April 1994.
- [75] S. K. S. Gupta, S. D. Kaushik, C.-H. Huang, and P. Sadayappan, "On compiling array expressions for efficient execution on distributed-memory machines," Department of Computer and Information Science, The Ohio State University, Tech. Rep. 19, 1994.
- [76] S. D. Kaushik C.-H. Huang, and P. Sadayappan, "Compiling array statements for efficient execution on distributed-memory machines: Two-level mappings," in *Language and Compilers for Parallel Computing*. Columbus: The Ohio State University, 1995, pp. 14.1–14.5.
- [77] C. Ancourt, "Génération automatique de codes de transfert pour multiprocesseurs à mémoires locales," PhD thesis, Université Paris VI, March 1991.
- [78] V. Van Dongen, "Compiling distributed loops onto SPMD code," *Parallel Proc. Lett.*, vol. 4, pp. 301–312, March 1994.
- [79] F. Irigoien, "Code generation for the hyperplane method and for loop interchange," CRI, École des mines de Paris, ENSMP-CAI-88 E102/CAI/I, Oct. 1988.
- [80] F. Balasa, F. H. M. Franssen, F. V. M. Catthoor, and H. J. De Man, "Transformation on nested loops with modulo indexing to affine recurrences," *Parallel Proc. Lett.* vol. 4, pp. 271–280, March 1994.
- [81] A. V. Aho, R. Sethi, and J. D. Ullman, *Compilers Principles, Techniques, and Tools*. Reading: Addison-Wesley, 1986.
- [82] F. Gasperoni and U. Scheiegelshohn, "Scheduling loop on parallel processors: A simple algorithm with close to optimum performance," Lecture Note, INRIA, 1992.
- [83] J. Wang and C. Eisenbeis, "Decomposed software pipelining: A new approach to exploit instruction level parallelism for loop programs," presented at the Working Conf. on Architectures and Compilation Techniques for Fine and Medium Grain Parallelism, 1993.
- [84] F. Irigoien, P. Jouvelot, and R. Triolet, "Semantical interprocedural parallelization: An overview of the PIPS project," in *ACM Int. Conf. on Supercomputing*, pp. 144–151, 1991.
- [85] D. Callahan and K. Kennedy, "Analysis of interprocedural side effects in a parallel programming environment," *J. Parallel Distrib. Comput.*, vol. 5, pp. 517–550, 1988.
- [86] V. Balasundaram and K. Kennedy, "A technique for summarizing data access and its use in parallelism enhancing transformations," in *ACM SIGPLAN Int. Conf. on Programming Language Design and Implementation*, 1989.
- [87] P. Havlak and K. Kennedy, "An implementation of interprocedural bounded regular section analysis," *IEEE Trans. Parallel Distrib. Systems*, vol. 2, pp. 350–360, July 1991.
- [88] F. Coelho, "Experiments with HPF compilation for a network of workstations," in *High Performance Computing and Networking*. New York: Springer-Verlag LNCS 797, 1994, pp. 423–428.
- [89] F. Coelho, "Compilation of I/O communications for HPF," in *5th Symp. on the Frontiers of Massively Parallel Computation*, 1995, p. 102.



Hindawi

Submit your manuscripts at
<http://www.hindawi.com>

