Symbolic Cost Analysis and Automatic Data Distribution for a Skeleton-Based Language

Julien Mallet

IRISA, Campus de Beaulieu, 35042 Rennes, France
mallet@irisa.fr

Abstract. We present a skeleton-based language which leads to portable and cost-predictable implementations on MIMD computers. The compilation process is described as a series of program transformations. We focus in this paper on the step concerning the distribution choice. The problem of automatic mapping of input vectors onto processors is addressed using symbolic cost evaluation. Source language restrictions are crucial since they permit to use powerful techniques on polytope volume computations to evaluate costs precisely. The approach can be seen as a cross-fertilization between techniques developed within the FORTRAN parallelization and skeleton communities.

1 Introduction

A good parallel programming model must be portable and cost predictable. General sequential languages such as FORTRAN achieve portability but cost estimations are often very approximate. The approach described in this paper is based on a restricted language which is portable and allows an accurate cost analysis. The language enforces a programming discipline which ensures a predictable performance on the target parallel computer (there will be no “performance bugs”).

Language restrictions are introduced through skeletons which encapsulate control and data flow in the sense of [6], [7]. The skeletons act on vectors which can be nested. There are three classes defined as sets of restricted skeletons: *computation skeletons* (classical data parallel skeletons), *communication skeletons* (data motion over vectors), and *mask skeletons* (conditional data parallel computation). The target parallel computers are MIMD computers with shared or distributed memory. Our compilation process can be described as a series of program transformations starting from a source skeleton-based language to an SPMD-like skeleton-based language. There are three main transformations: in-place updating, making all communications explicit and distribution. Due to space concerns, we focus in this paper only on the step concerning the choice of the distribution.

We tackle the problem of automatic mapping of input vectors onto processors using a fixed set of classical distributions (e.g. row block, block cyclic,...). The goal is to determine mechanically the best distribution for each input vector

through cost evaluation. The restrictions of the source language and the fixed set of standard data distributions guarantee that the parallel cost (computation + communication) can be computed accurately for all source programs. Moreover, these restrictions allow us to evaluate and compare accurate execution times in a symbolic form. So, the results do not depend on specific vector sizes nor on a specific number of processors.

Even if our approach is rooted in the field of data-parallel and skeleton-based languages, one of its specificities is to reuse techniques developed for FORTRAN parallelization (polytope volume computation) and functional languages (single-threading analysis) in a unified framework.

The article is structured as follows. Section 2 is an overview of the whole compilation process. Section 3 presents the source language and the target parallel language. In Section 4, we describe the symbolic cost analysis through an example: LU decomposition. We report some experiments done on a MIMD distributed memory computer in Section 5. We conclude by a review of related work.

The interested reader will find additional details in an extended version of this paper ([12]).

2 · Compilation Overview

The compilation process consists of a series of program transformations:

$$\mathcal{L}_1 \longrightarrow \mathcal{L}_2 \longrightarrow \mathcal{L}_3 \equiv\equiv\equiv \mathcal{L}_4 \longrightarrow \mathcal{L}_5$$

Each transformation compiles a particular task by mapping skeleton programs from one intermediate language into another. The source language ($\mathcal{L}_1$) is composed of a collection of higher-order functions (skeletons) acting on vectors (see Section 3). It is primarily designed for a particular domain where high performance is a crucial issue: numerical algorithms. $\mathcal{L}_1$ is best viewed as a parallel kernel language which is supposed to be embedded in a general sequential language (e.g. C).

The first transformation ($\mathcal{L}_1 \rightarrow \mathcal{L}_2$) deals with in-place updating, a standard problem in functional programs with arrays. We rely on a type system in the spirit of [17] to ensure that vectors are manipulated in a single-threaded fashion. The user may have to insert explicit copies in order for his program to be well-typed. As a result, any vector in a $\mathcal{L}_2$ program can be implemented by a global variable.

The second transformation ($\mathcal{L}_2 \rightarrow \mathcal{L}_3$) makes all communications explicit. Intuitively, in order to execute an expression such as `map ( $\lambda x.x+y$ )` in parallel, `y` must be broadcast to every processor before applying the function. The transformation makes this kind of communication explicit. In $\mathcal{L}_3$, all communications are expressed through skeletons.

The transformation $\mathcal{L}_3 \rightarrow \mathcal{L}_4$ concerns automatic data distribution. We restrict ourselves to a small set of standard distributions. A vector can be distributed cyclicly, by contiguous blocks or allocated to a single processor. For a

matrix (vector of vectors), this gives 9 possible distributions (cyclic cyclic, block cyclic, line cyclic, etc.). The transformation constructs a single vector from the input vectors according to a given data distribution. This means, in particular, that all vector accesses have to be changed according to the distributions. Once the distribution is fixed, some optimizations (such as copy elimination) become possible and are performed on the program. The transformation yields an SPMD-like skeleton program acting on a vector of processors.

In order to choose the best distribution, we transform the $\mathcal{L}_3$ program according to all the possible distributions of its input parameters. The symbolic cost of each transformed program can then be evaluated and the smallest one chosen (see Section 4). For most numerical algorithms, the number of input vectors is small and this approach is practical. In other cases, we would have to rely on the programmer to prune the search space.

The $\mathcal{L}_4 \rightarrow \mathcal{L}_5$ step is a straightforward translation of the SPMD skeleton program to an imperative program with calls to a standard communication library. We currently use C with the MPI library.

Note that all the transformations are automatic. The user may have to interact only to insert copies ($\mathcal{L}_1 \rightarrow \mathcal{L}_2$ step) or, in $\mathcal{L}_4$, when the symbolic formulation does not designate a single best distribution (see 4.4).

3 The Source and Target Languages

The source language $\mathcal{L}_1$ is basically a first-order functional language without general recursion, extended with a collection of higher-order functions (the skeletons). The main data structure is the vector which can be nested to model multi-dimensional arrays. An $\mathcal{L}_1$ program is a main expression followed by definitions. A definition is either a function definition or the declaration of the (symbolic) size of an input vector. Combination of computations is done through function composition (.) and the predefined iterator `iterfor`. `iterfor n f a` acts like a loop applying `n` times its function argument `f` on `a`. Further, it makes the current loop index accessible to its function argument.

The $\mathcal{L}_1$ program implementing LU decomposition is:

```

LU(M) where
  M :: Vect n (Vect n Float)
  LU(a)                = iterfor (n-1) (calc.fac.apivot.colrow) a
  calc(i,a,row,piv)    = (i,(map (map first)
                                .rect i (n-1) (rect i (n-1) fcalc)
                                .map zip3.zip3)(a,row,piv))
  fcalc(a,row,piv)     = (a-row*piv,row,piv)
  fac(i,a,row,col,piv) = (i,a,row,(map(map /).map zip.zip)(col,piv))
  apivot(i,a,row,col)  = (i,a,row,col,map (brdcast (i-1)) row)
  colrow(i,a)          = (i,a,brdcast (i-1) a,map (brdcast (i-1)) a)
  zip3(x,y,z)          = map p2t (zip(x,zip(y,z)))
  p2t(x,(y,z))         = (x,y,z)

```

For handling vectors, there are three classes of skeletons: computation, communication, and mask skeletons. The *computation skeletons* include the classical higher order functions **map**, **fold** and **scan**. The *communication skeletons* describe restricted data motion in vectors. In the definition of **colrow** in LU, **brdcast** (i-1) copies the (i-1)th vector element to all the other elements. Another communication skeleton is **gather** i M which copies the ith column of the matrix M to the ith row of M. There are seven communication skeletons which have been chosen because of their availability on parallel computers as hard-wired or optimized communication routines. The *mask skeletons* are data-parallel skeletons which apply their function argument only to a selected set of vector elements. In the definition of **calc** in LU, **rect** i (n-1) **fcalc** applies **fcalc** on vector elements whose index ranges from i to (n-1). Note that our approach is not restricted to these sole skeletons: even if it requires some work and caution, more iterators, computation or mask skeletons could be added.

In order to enable a precise symbolic cost analysis, additional syntactic restrictions are necessary. The scalar arguments of communication skeletons, the **iterfor** operator and mask skeletons must be linear expressions of **iterfor** indexes and variables of vector size. We rely on a type system with simple subtyping (not described here) to ensure that each variable appearing in such expression, are only **iterfor** and vector size variables.

It is easy to check that these restrictions are verified in LU. The integer argument of **iterfor** is a linear expression of an input vector size. The arguments of **rect** and **brdcast** skeletons are made of linear expressions of **iterfor** variable i or vector size variable n.

The target language $\mathcal{L}_4$ expresses SPMD (Single Program Multiple Data) computations. A program acts on a single vector whose elements represent the data spaces of the processors. Communications between processors are made explicit by data motion within the vector of processors in the spirit of the source parallel language of [8]. $\mathcal{L}_4$ introduces new versions of skeletons acting on the vector of processors.

An $\mathcal{L}_4$ program is a composition of computations, with possibly a communication inserted between each computation. More precisely, a program *FunP* has the following form:

$$FunP ::= (\text{pimap } Fun \mid \text{piterfor } E \text{ } FunP) (. Com)^{0/1} (. FunP)^{0/1}$$

A computation is either a parallel operation **pimap** *Fun* which applies *Fun* on each processor, or an iteration **piterfor** *E* *FunP* which applies *E* times the parallel computation *FunP* to the vector of processors. The functions *Fun* are similar to functions of $\mathcal{L}_1$. The only differences are the skeleton arguments which may include modulo, integer division, minimum and maximum functions. *Com* denotes the communication between processors. It describes data motion in the same way as the communication skeletons of $\mathcal{L}_1$ but on the vector of processors. For example, the corresponding parallel communication skeleton of **brdcast** is **pbrdcast** which broadcasts values from one processor of the vector of processors to each other.

4 Accurate Symbolic Cost Analysis

After transformation of the program according to different distributions, this step aims at automatically evaluating the complexity of each $\mathcal{L}_4$ program obtained in order to choose the most efficient. Our approach to get an accurate symbolic cost is to reuse results on polytope volume computations ([16], [5]). It is possible because the restrictions of the source language and the fixed set of data distributions guarantee that the abstracted cost of all transformed source programs can be translated into a polytope volume description.

First, an abstraction function $\mathcal{CA}$ takes the program and yields its symbolic parallel cost. The cost expression is transformed so that it can be seen as the definition of a polytope volume. Then, classic methods to compute the volume of a polytope are applied to get a symbolic cost in polynomial form. Finally, a symbolic math package (such as Maple) can be used to compare symbolic costs and find the smallest cost among the $\mathcal{L}_4$ programs corresponding to the different distribution choices.

4.1 Cost Abstraction and Cost Language $\mathcal{LC}$

The symbolic cost abstraction $\mathcal{CA}$ is a function which extracts cost information from parallel programs. The main rules are shown below. We use a non-standard notation for indexed sums: instead of $\sum_{i=1}^n$, we note $\sum_i \{ \begin{smallmatrix} 1 \leq i \\ i \leq n \end{smallmatrix} \}$. Communication costs are expressed as polynomials whose constants depend on the target computer. For example, the cost of `pbrdcast` involves the parameters α_{br} and β_{br} which denote respectively the time of one-word transfer between two processors and the message startup time on the parallel computer considered. One just has to set those constants to adapt the analysis for a specific parallel machine.

$\mathcal{CA} \llbracket f1 . f2 \rrbracket$	$= \mathcal{CA} \llbracket f1 \rrbracket + \mathcal{CA} \llbracket f2 \rrbracket$
$\mathcal{CA} \llbracket \text{piterfor } e \ (\lambda i.f) \rrbracket$	$= \sum_i \{ \begin{smallmatrix} 1 \leq i \\ i \leq e \end{smallmatrix} \} \mathcal{CA} \llbracket f \rrbracket$
$\mathcal{CA} \llbracket \text{pimap } (\lambda i.f) \rrbracket$	$= \max_{i=0}^{p-1} \mathcal{CA} \llbracket f \rrbracket$ where p is the number of processors
$\mathcal{CA} \llbracket \text{pbrdcast} \rrbracket$	$= (\alpha_{br} * b + \beta_{br}) * p$ where $\begin{cases} p \text{ is the number of processors} \\ b \text{ is the size of broadcast data} \end{cases}$
$\mathcal{CA} \llbracket \text{rect el } e2 \ f \rrbracket$	$= \sum_i \{ \begin{smallmatrix} e1 \leq i \\ i \leq e2 \end{smallmatrix} \} \mathcal{CA} \llbracket f \rrbracket$

The cost abstraction yields accurate costs expressed in a simple language $\mathcal{LC}$. $\mathcal{LC}$ is made of sums of generalized sums (G-sums): $\sum_i \{ \dots \}$ (...) and generalized maxs (G-maxs): $\max_{i=e}^{e'} (\dots)$ possibly nested. A G-max denotes the maximum expression one can get when the max variable ranges over the integer interval. Inequalities in G-sums involve only linear expressions and maximum and minimum function of them. There may be polynomial included in the body of G-sums and G-maxs. The variables of polynomials are only denoting vector sizes or the number of processors. In order to illustrate the analysis, we consider here only the two best distribution choices for LU: row block and row cyclic. The ab-

stracted costs are given below; note that only the part of the costs which differs is detailed:

$$\begin{aligned} \mathcal{C}_{bloc}^1 &\equiv \sum_i \{ \{ \frac{1 \leq i}{i \leq n-1} \} \left(\max_{i_p=0}^{p-1} \sum_j \left\{ \begin{array}{c} \max(0, i - i_p * b) \leq j \\ j \leq \min(b-1, n-1 - i_p * b + b) \end{array} \right\} (\sum_k \{ \frac{i \leq k}{k \leq n-1} \} 1) \right) \} + C \\ \mathcal{C}_{cyc}^1 &\equiv \sum_i \{ \{ \frac{1 \leq i}{i \leq n-1} \} \left(\max_{i_p=0}^{p-1} \sum_j \left\{ \begin{array}{c} \max(0, \lfloor \frac{i - i_p}{p} \rfloor) \leq j \\ j \leq \min(b, \lceil \frac{n-1 - i_p}{p} \rceil) \end{array} \right\} (\sum_k \{ \frac{i \leq k}{k \leq n-1} \} 1) \right) \} + C \end{aligned}$$

We emphasize that writing the source program in $\mathcal{L}_1$ is crucial to get an accurate symbolic cost. First of all, without a severe limitation of the use of recursion no precise cost could be evaluated in general. Further, the restrictions imposed by $\mathcal{L}_1$ ensure that every abstract parallel cost belongs to $\mathcal{LC}$. The mask skeletons (which limit conditional application to index intervals), the communication and the computation skeletons all have a complexity which depends polynomially on the vector size. Their costs can be described by nested sums. Another important restriction is that expressions involving iteration indexes (mask skeletons and iterfor bounds) are linear. This restriction, in addition to the form of the standard distributions which keep the linearity of the vector accesses, ensure that the inequalities occurring in $\mathcal{LC}$ are linear.

4.2 Transformation to Descriptions of Polytope Volume

Polytope volumes are only defined in terms of nested G-sums with linear inequalities containing neither max nor min ([5]). In order to apply methods to compute polytope volumes, we must remove the G-max, min, max occurring in the cost expressions.

The transformation takes the form of a structural recursive scan of the cost expression. The min and max appearing in inequalities are transformed into conjunctions of inequalities. For example, the expression $\max(0, i - i_p * b) \leq j$, in $\mathcal{C}_{bloc}^1$, becomes $0 \leq j \wedge i - i_p * b \leq j$. G-max expressions are propagated through their subexpressions. For example, in $\mathcal{C}_{bloc}^1$, $\max_{i_p=0}^{p-1} \sum_j \left\{ \begin{array}{c} 0 \leq j \wedge i - i_p * b \leq j \\ j \leq b-1 \wedge j \leq n-1 - i_p * b + b \end{array} \right\} (\dots)$

is temporarily transformed into $\sum_j \left\{ \begin{array}{c} 0 \leq j \wedge \min_{i_p=0}^{p-1} (i - i_p * b) \leq j \\ j \leq b-1 \wedge j \leq \max_{i_p=0}^{p-1} (n-1 - i_p * b + b) \end{array} \right\} \max_{i_p=0}^{p-1} (\dots)$. The max value of a G-sum is a sum where the upper bounds are maximized and the lower ones minimized. For a sum inequality $i \leq e$, we propagate the G-max inside e until it reaches a variable or a constant. We can finally use the facts that $\max_{i_p=0}^{p-1} i_p = p-1$ and $\max_{i_p=0}^{p-1} k = k$ if $k \neq i_p$. For an inequality $e \leq i$, the transformation propagates a G-min analogously. Since no polynomial in $\mathcal{LC}$ may contain a G-max index, the G-max of a polynomial is just this polynomial. In the case of distributed LU, the cost expressions are transformed into

$$\begin{aligned} \mathcal{C}_{bloc}^1 &= \sum_{(i,j,k)} \left\{ \begin{array}{c} 1 \leq i \leq n-1 \\ 0 \leq j \wedge i - (p-1) * b \leq j \\ j \leq b-1 \wedge j \leq n-1 + b \wedge i \leq k \leq n-1 \end{array} \right\} 1 + C \equiv \mathcal{C}_{bloc}^2 \\ \mathcal{C}_{cyc}^1 &= \sum_{(i,j,k)} \left\{ \begin{array}{c} 1 \leq i \leq n-1 \wedge 0 \leq j \wedge \lfloor \frac{i-p+1}{p} \rfloor \leq j \\ j \leq b-1 \wedge j \leq \lceil \frac{n}{p} \rceil - 1 \wedge i \leq k \leq n-1 \end{array} \right\} 1 + C \equiv \mathcal{C}_{cyc}^2 \end{aligned}$$

The technique described to remove G-max expressions may maximize the real cost. This is due to the duplication of G-max in G-sums. A more complicated

but accurate symbolic solution also exists. The idea is to evaluate (as described in the next section) the nested G-sums first. When the G-sum is reduced to a polynomial the problem amounts to computing symbolic maximum of polynomials over symbolic intervals. This can be done using symbolic differentiation, solution and comparison of polynomials. At this point, we have only used the simpler approximated solution because the approximation is minor for our examples (e.g. none for LU) and the accurate symbolic solution is not implemented as such in existing symbolic math packages.

4.3 Parametrized Polytope Volume Computation

A parametrized polytope is a set of points whose coordinates satisfy a system of linear inequalities with possibly unknown parameters. After the previous transformation, the cost expression denotes a polytope volume, that is the number of points in the associated polytope. The works of [16] and [5] describe algorithms for computing symbolically the volume of parametrized polytopes. The result is a polynomial whose variables are the system parameters. Further, in [5], Clauss presents an extension for systems of non-linear inequalities with floor ($\lfloor \cdot \rfloor$) and ceiling ($\lceil \cdot \rceil$) functions as appearing in $\mathcal{C}_{cyc}^2$.

Thus, the algorithm of [5] allows one to transform each parallel cost expression into a polynomial. This method applied to LU yields:

$$\begin{aligned}\mathcal{C}_{bloc}^2 &= n^3 \left(\frac{3p^2-1}{6p^3} \right) - \frac{n^2}{2p} + \frac{n}{6p} + C && \equiv \mathcal{C}_{bloc}^3 \\ \mathcal{C}_{cyc}^2 &= \frac{n^3}{3p} + n^2 \left(\frac{p-2}{2p} \right) + n \left(\frac{-3p^2+6p-2}{6p} \right) + \frac{p^2-3p+2}{6} + C && \equiv \mathcal{C}_{cyc}^3\end{aligned}$$

4.4 Symbolic Cost Comparison

The last step is to compare the symbolic costs of different distribution choices. It amounts to computing the symbolic intervals where the difference of cost (polynomial) is positive or negative. Symbolic math packages such as Maple can be used for solving this problem. In the case of LU, Maple produces the following condition:

$$\mathcal{C}_{bloc}^3 - \mathcal{C}_{cyc}^3 \geq 0 \Leftrightarrow n \geq p$$

The programmer may have to indicate if the relations given by Maple are satisfied or not. In our example, he must indicate if $n \geq p$. Another (automatic) solution is to use the relations as run-time tests which choose between several versions of the program.

In our example, the difference in the two costs can be explained by the fact that the cyclic distribution provides a much better load balancing than the block distribution whereas communications are identical.

5 Experiments

We have performed experiments on an Intel Paragon XP/S with a handful of standard linear algebra programs (LU, Cholesky factorization, Householder, Ja-

cobi elimination, ...). Our implementation is not completed and some compilation steps, such as the destructive update step and part of the symbolic cost computation, were done manually. Below, a table gathers the execution times obtained for LU decomposition. They are representative of the results we got for the other few programs. For all programs, the distribution chosen by the cost analysis proved to be the best one in practice.

Processors	Skel. cyclic	Skel. bloc	C Seq.	HPF cyclic	ScaLAPACK cyclic
1	14.77	15.07	13.61	15.36	3.78
4	5.25	6.75	×	5.41	1.84
16	2.97	5.33	×	3.06	1.50
32	2.57	5.58	×	2.67	1.41

We compared the sequential execution of skeleton programs with standard (and portable) C versions and our parallel implementation with High Performance Fortran (a manual distribution approach). No significant sequential or parallel runtime penalty seems to result from programming using skeletons, at least for such regular algorithms.

We compared our code with the parallel implementation of NESL, a skeleton-based language [1]. The work on the implementation of NESL has mostly been directed towards SIMD machines. On the Paragon, the NESL compiler distributes vectors uniformly on processors and communications are not optimized. Not surprisingly, the parallel code is very inefficient (at least fifty times slower than our code).

We also compared our implementation with ScaLAPACK, an optimized library of linear algebra programs designed for distributed memory MIMD parallel computers [4]. In ScaLAPACK, the user may explicitly indicate the data distribution. So, we indicated the best distribution found by the cost analysis in each ScaLAPACK program considered. If our code on 1 processor is much slower than its ScaLAPACK equivalent (between 3 to 5 times slower), the difference decreases as the number of processors increases (typically, 1.8 times slower on 32 processors). Much of this difference comes from the machine specific routines used by ScaLAPACK for performing matrix operations (the BLAS library). This suggests a possible interesting extension of our source language. The idea would be to introduce new skeletons corresponding to the BLAS operations in order to benefit from these machine specific routines.

These preliminary results are promising but more experiments are necessary to assess both the expressiveness of the language and the efficiency of the compilation. We believe that these experiments may also indicate useful linguistic extensions (e.g. new skeletons) and new optimizations.

6 Related Work and Conclusion

The community interested in the automatic parallelization of FORTRAN has studied automatic data distribution through parallel cost estimation ([10], [3]).

If the complete FORTRAN language (unrestricted conditional, indexing with runtime value, ...) is to be taken into account, communication and computation costs cannot be accurately estimated. In practice, the approximated cost may be far from the real execution time leading to a bad distribution choice. [16], [5] focus on a subset of FORTRAN: loop bound and array indexes are linear expressions of the loop variables. This restriction allows them to compute a precise symbolic computation cost through their computations of polytope volume. Unfortunately, using this approach to estimate communication costs is not realistic. Indeed, the cost would be expressed in terms of point-to-point communications without taking into account hard-wired communication primitives. All these works estimate real costs too roughly to ensure that a good distribution is chosen.

The skeleton community has studied the transformation of restricted computation patterns into lower-level parallel primitives. [7] defines a restricted set of skeletons which are transformed using cost estimation. Only cost-reducing transformations are considered. It is, however, well known that often intermediary cost-increasing transformations are necessary to derive a globally cost optimal algorithm. [15], [11] and [14] define cost analyses for skeleton-based languages. Their skeletons are more general than ours leading to approximate parallel cost (communication or/and computation). Furthermore, the costs are not symbolic (the size of input matrices and the number of processors are supposed to be known). [9] define a precise communication cost for scan and fold skeletons on several parallel topologies (hypercube, mesh, ...) which allows them to apply optimizations of communications through cost-reducing transformations. There are also a few real parallel implementations of skeleton-based languages. [2] uses cost estimations based on profiling which does not ensure a good parallel performance for different sizes of inputs. [13] uses a finer cost estimation but implementation decisions are taken locally and no arbitration of tradeoffs is possible.

We have presented in this paper the compilation of a skeleton-based language for MIMD computers. Working by program transformations in a unified framework simplifies the correctness proof of the implementation. One can show independently for each step that the transformation preserves the semantics and that the transformed program respects the restrictions enforced by the target language. The overall approach can be seen as promoting a programming discipline whose benefit is to allow precise analyses and a predictable parallel implementation. The source language restrictions are central to the approach as well as the techniques to evaluate the volume of polytopes. We regard this work as a rare instance of cross-fertilization between techniques developed within the FORTRAN parallelization and skeleton communities. A possible research direction is to study dynamic redistributions chosen at compile-time. Some parallel algorithms are much more efficient in the context of dynamic data redistribution. A completely automatic and precise approach to this problem would be possible in our framework. However, this would lead to a search space of exponential size. A possible solution is to have the user select the redistribution places.

Acknowledgements : Thanks to Pascal Fradet, Daniel Le Métayer, Mario Südholt for commenting on an earlier version of this paper.

References

1. G. E. Blelloch, S. Chatterjee, J. C. Hardwick, J. Sipelstein, and M. Zagha. Implementation of a portable nested data-parallel language. In *4th ACM Symp. on Princ. and Practice of Parallel Prog.*, pages 102–112. 1993.
2. T. Bratvold. A Skeleton-Based Parallelising Compiler for ML. In *5th Int. Workshop on the Imp. of Fun. Lang.*, pages 23–33, 1993.
3. S. Chatterjee, J. R. Gilbert, R. Schreiber, and S. Teng. Automatic array alignment in data-parallel program. In *20th ACM Symp. on Princ. of Prog. Lang.*, pages 16–28, 1993.
4. J. Choi and J. J. Dongarra. Scalable linear algebra software libraries for distributed memory concurrent computers. In *Proc. of the 5th IEEE Workshop on Future Trends of Distributed Computing Systems*, pages 170–177, 1995.
5. P. Clauss. Counting solutions to linear and nonlinear constraints through Ehrhart polynomials: Applications to analyze and transform scientific programs. In *ACM Int. Conf. on Supercomputing*, 1996.
6. M. Cole. A skeletal approach to the exploitation of parallelism. In *CONPAR'88*, pages 667–675. Cambridge University Press, 1988.
7. J. Darlington, A. J. Field, P. G. Harrison, P. H. J. Kelly, D. W. N. Sharp, Q. Wu, and R. L. While. Parallel programming using skeleton functions. In *PARLE '93*, pages 146–160. LNCS 694, 1993.
8. J. Darlington, Y. K. Guo, H. W. To, and Y. Jing. Skeletons for structured parallel composition. In *5th ACM Symp. on Princ. and Practice of Parallel Prog.*, pages 19–28, 1995.
9. S. Gorlatch and C. Lengauer. (De)Composition rules for parallel scan and reduction. In *3rd IEEE Int. Conf. on Massively Par. Prog. Models*, 1998.
10. M. Gupta and P. Banerjee. Demonstration of automatic data partitioning techniques for parallelizing compilers on multicomputers. *IEEE Transactions on Parallel and Distributed Systems*, 3(2):179–193, 1992.
11. C. B. Jay, M. I. Cole, M. Sekanina, and P. Steckler. A monadic calculus for parallel costing of a functional language of arrays. In *Euro-Par'97 Parallel Processing*, pages 650–661. LNCS 1300, 1997.
12. J. Mallet. Compilation for MIMD computers of a skeleton-based language through symbolic cost analysis and automatic data distribution. Technical Report 1190, IRISA, May 1998.
13. S. Pelagatti. *A Methodology for the Development and the Support of Massively Parallel Programs*. PhD thesis, Pise University, 1993.
14. R. Rangaswami. *A Cost Analysis for a Higher-order Parallel Programming Model*. PhD thesis, Edinburgh University, 1996.
15. D. B. Skillicorn and W. Cai. A cost calculus for parallel functional programming. Technical report, Queen's University, 1993.
16. N. Tawbi. Estimation of nested loops execution time by integer arithmetic in convex polyhedra. In *Int. Symp. on Par. Proc.*, pages 217–223, 1994.
17. P. Wadler. Linear types can change the world! In *Programming Concepts and Methods*, pages 561–581. North Holland, 1990.