

Compilation of a Specialized Functional Language for Massively Parallel Computers

Pascal Fradet, Julien Mallet

► **To cite this version:**

Pascal Fradet, Julien Mallet. Compilation of a Specialized Functional Language for Massively Parallel Computers. [Research Report] RR-3894, INRIA. 2000. <inria-00072760>

HAL Id: inria-00072760

<https://hal.inria.fr/inria-00072760>

Submitted on 24 May 2006

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

*Compilation of a Specialized Functional
Language for Massively Parallel Computers*

Pascal FRADET and Julien MALLET

N°3894

mars 2000

THÈME 2



*R*apport
de recherche



Compilation of a Specialized Functional Language for Massively Parallel Computers

Pascal FRADET* and Julien MALLET*

Thème 2 — Génie logiciel
et calcul symbolique
Projet Lande

Rapport de recherche n° 3894 — mars 2000 — 52 pages

Abstract: We propose a parallel specialized language that ensures portable and cost-predictable implementations on parallel computers. The language is basically a first-order, recursion-less, strict functional language equipped with a collection of higher-order functions or skeletons. These skeletons apply on (nested) vectors and can be grouped in four classes: computation, reorganization, communication, and mask skeletons. The compilation process is described as a series of transformations and analyses leading to SPMD-like functional programs which can be directly translated into real parallel code. The language restrictions enforce a programming discipline whose benefit is to allow a static, symbolic, and accurate cost analysis. The parallel cost takes into account both load balancing and communications, and can be statically evaluated even when the actual size of vectors or the number of processors are unknown. It is used to automatically select the best data distribution among a set of standard distributions. Interestingly, this work can be seen as a cross fertilization between techniques developed within the FORTRAN parallelization, skeleton, and functional programming communities.

Key-words: specialized language, skeleton-based language, program analyses and transformations, parallelism, cost analysis, data distribution

(Résumé : tsvp)

* {fradet,mallet}@irisa.fr

Unité de recherche INRIA Rennes
IRISA, Campus universitaire de Beaulieu, 35042 RENNES Cedex (France)
Téléphone : 02 99 84 71 00 - International : +33 2 99 84 71 00
Télécopie : 02 99 84 71 71 - International : +33 2 99 84 71 71

Compilation d'un langage fonctionnel spécialisé pour machine massivement parallèle

Résumé : Nous proposons un langage fonctionnel spécialisé dont l'implantation sur machine massivement parallèle est portable et disposant d'un coût prévisible. Le langage est un langage fonctionnel strict, du premier ordre à la récursivité bridée disposant d'une collection de fonctions d'ordre supérieur ou patrons. Ces patrons s'appliquent à des vecteurs (emboîtés) et peuvent être décomposés en quatre classes: patrons de calcul, de réorganisation, de communication et de masque. Le processus de compilation est décrit par une suite de transformations et d'analyses de programmes. Le langage cible est un langage à patrons décrivant des programmes SPMD directement traduisible en un langage séquentiel avec des appels à des fonctions de bibliothèque de communication. Les restrictions du langage imposent une discipline de programmation permettant une analyse statique, symbolique et précise du coût. Le coût parallèle tient compte de la répartition des calculs et des communications et peut être statiquement évalué même lorsque la taille des vecteurs d'entrée ou le nombre de processeurs sont inconnus. Il est utilisé pour choisir automatiquement la meilleure distribution de données parmi un ensemble de distributions standards. Ce travail est un exemple de fertilisation croisée entre les techniques développées dans les communautés de la parallélisation FORTRAN, des patrons et de la programmation fonctionnelle.

Mots-clé : langage spécialisé, langage à patrons, parallélisme, analyses et transformations de programmes, analyse de coût, distribution de données

1 Introduction

A good parallel programming model must be portable and cost predictable. General purpose languages such as FORTRAN achieve portability but cost estimations are often very approximate. A precise cost analysis is especially important in this context since the goal of parallelization is efficiency and its impact on the overall cost is, at best, a division by a constant. So, orders of magnitude or maximum complexities are insufficient to guide parallel implementation choices.

The approach described in this paper is based on a restricted functional language that is portable and allows us to design an automatic and accurate cost analysis. The language restrictions can be seen as enforcing a programming discipline that ensures a predictable performance on the target parallel computer (there will be no “performance bugs”). General recursion and conditionals are replaced by skeletons that encapsulate control and data flow in the sense of [Col88] or [DFH⁺93]. The skeletons, which act on (potentially nested) vectors, can be grouped into four classes: the *computation skeletons* (classical data parallel functions), the *reorganization skeletons* (creating and structuring vectors), the *communication skeletons* (data motion over vectors), and the *mask skeletons* (conditional data parallel functions). The skeletons and data structures have been chosen with scientific computing in mind. Matrix computations and nested for loops are easy to describe and many standard numerical algorithms have been expressed easily in our kernel language. Concerning the target parallel architecture, we aimed at MIMD (*Multiple Instructions Multiple Data*) computers with shared or distributed memory but SIMD (*Single Instruction Multiple Data*) computers could be accommodated as well.

The compilation process is described as a series of program transformations leading to SPMD-like (*Single Program Multiple Data*) functional programs which can be directly translated into true parallel code. Each compilation step transforms a skeleton-based language into another closer to a code with explicit parallelism. The main compilation steps consist in a size analysis, an update-in-place analysis, a transformation making all communications explicit, transformations implementing data distribution, and a symbolic code analysis. A key task in the compilation of data parallelism is the choice of the distribution since it determines both the load balancing and the communications between processors. In our approach, the restrictions imposed by the source language make it possible to choose automatically the best data distribution among a set of standard distributions. This choice relies on the cost analysis which evaluates accurate parallel execution times. One of the main challenges of our work was to tackle this problem without knowing the actual size of vectors or the number of processors. In other words, the analysis ought to be symbolic.

The contributions of this work are both technical and methodological.

- The compilation is efficient and original: it integrates very different analysis techniques in a sequence of program transformations. Maybe the most important contribution lies in the definition of the specialized language. We establish the necessary language restrictions to ensure the accuracy of a symbolic cost analysis. The compilation takes advantage of the analysis to choose the best parallel implementation

among a set of standard implementations. The analyses for FORTRAN nested loops [GB92, Fea94] do not evaluate precisely the cost of the communications. In particular, collective communications (diffusion, translation...) cannot be automatically detected in FORTRAN programs. In our approach, the analysis takes into account both load balancing and communication costs since the collective communications appear explicitly through communication skeletons. Most existing skeletons implementations [BHS⁺94, DGTJ95] are based on fixed implementations for each skeleton. This local view may lead to redistributing data before each skeleton application and therefore be very inefficient. Finding the best, global distribution is more complex but more efficient since it takes into account compromises (i.e. trading a local cost increase for a global improvement).

- At the methodology level, this work is a rare example of cross-fertilization between three different research fields: the automatic parallelization of FORTRAN, skeleton based languages, and functional programming. The FORTRAN community has worked on symbolic complexity analyses for subsets of FORTRAN. We adapt and extend this work to define the cost analysis in our context. Skeleton-based languages with their collection of data parallel functions provide a framework to define specialized parallel languages. We build upon this approach and provide two new collections of skeletons to the programmer: the communication skeletons and the mask skeletons. The community of functional programming has produced a large body of work on typing, program transformation, and analysis. Our compilation process, made of a collection of program analyses and transformation, relies on this work.

The article is structured as follows. Section 2 is an overview of the compilation process. Section 3 presents the source language and the target parallel language. Sections 4 to 8 describe each compilation step in turn. We report some experiments done on MIMD distributed memory computers in Section 9. Section 10 justifies *a posteriori* the source language restrictions, reviews related work, and suggests directions for further research.

2 Overview

The compilation process consists of a series of program transformations:

$$\mathcal{L}_1 \xrightarrow{GL} \mathcal{L}_2 \xrightarrow{EC} \mathcal{L}_3 \xrightarrow{ABS} \mathcal{L}_4 \xrightarrow{DIST} \mathcal{L}_5 \xrightarrow{OPT} \mathcal{L}_6 \xrightarrow{TRA} \text{Parallel Code}$$

Each arrow represents a transformation compiling a particular task by mapping skeleton programs from one intermediate language (\mathcal{L}_i) into another (\mathcal{L}_{i+1}). The source language (\mathcal{L}_1) is composed of a collection of higher-order functions (skeletons) acting on vectors (see Section 3.1). 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 embedded in a general sequential language (e.g. C). Only, parts of programs written in \mathcal{L}_1

will be executed in parallel whereas others parts will be executed in sequential, for example, on the host computer of the parallel machine.

The first compilation step is the *type/size analysis* of \mathcal{L}_1 programs. The analysis computes the shape (size) of all the vectors occurring in the program (Section 4). It must infer symbolic sizes because the sizes of input vectors may be unknown at compile time. As a byproduct, the analysis infers conditions ensuring that no vector access error occurs at runtime.

The first transformation ($\mathcal{L}_1 \rightarrow \mathcal{L}_2$) deals with in-place updating, a standard problem in functional programming with aggregates (Section 5). The program is analyzed to check that all vectors can be safely modified in place. If the program does not pass the analysis, it must be transformed by inserting explicit vector copies. This can be done automatically (the analysis indicates the places where to insert copies) or manually (the programmer may want to restructure the program to insert less copies). This step is also used to guarantee that vectors are either returned as result or are explicitly deallocated (i.e. a garbage collector is not needed).

The transformation \mathcal{EC} ($\mathcal{L}_2 \rightarrow \mathcal{L}_3$) makes all communications explicit (Section 6). 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 the language \mathcal{L}_3 , all communications are expressed through skeletons.

The transformations from \mathcal{L}_3 to \mathcal{L}_6 concerns automatic data distribution (Section 7). First, λ -abstractions and variables are removed by threading an explicit environment throughout the program (transformation \mathcal{ABS} , Section 7.1). This transformation, reminiscent of abstraction algorithms, prepares the distribution transformation \mathcal{DIST} ($\mathcal{L}_4 \rightarrow \mathcal{L}_5$). We consider a set of standard distributions of the input vectors. 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.). Distribution transforms programs so that they act on a single vector whose elements represent the processors. This implies, in particular, to change all vector accesses according to the distribution (Section 7.2). Finally, distributed programs are optimized (\mathcal{OPT} : $\mathcal{L}_5 \rightarrow \mathcal{L}_6$) using a set of local transformations (Section 7.3). After distribution, some vector copies become useless. They are removed in order to improve the sequential execution time. \mathcal{L}_6 programs apply on a vector of processors and simulate an SPMD code.

In order to choose the best distribution, a \mathcal{L}_4 program is transformed according to all the possible distributions of its input parameters leading to a set of \mathcal{L}_6 programs. The symbolic cost of each version is evaluated and the smallest one chosen (Section 8). 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 transformation \mathcal{TRA} ($\mathcal{L}_6 \rightarrow \text{Parallel Code}$) 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 (*Message Passing Interface*) library along with the

C compiler of the host machine (Section 9 presents experiments on an Intel Paragon XP/S and a CRAY T3E).

All the transformations are automatic. Nevertheless, the user can interact with the compiler, for example to insert explicit copies in the $\mathcal{L}_1 \rightarrow \mathcal{L}_2$ step, or to guide the choice of the best distribution in \mathcal{L}_5 .

For each transformation $\mathcal{T}_i: \mathcal{L}_i \rightarrow \mathcal{L}_{i+1}$, three correctness properties must be proved. First, it must be shown that the transformation \mathcal{T}_i transforms programs of \mathcal{L}_i into programs of \mathcal{L}_{i+1} , formally:

Property 1 $\forall Prog \in \mathcal{L}_i \Rightarrow \mathcal{T}_i[Prog] \in \mathcal{L}_{i+1}$.

Second, the transformation \mathcal{T}_i must preserve the semantics of programs, formally:

Property 2 $\forall Prog \in \mathcal{L}_i, \mathcal{T}_i[Prog] = Prog$

Third, it must be checked that the update-in-place property still holds on transformed programs.

Property 3 $\forall Prog \in \mathcal{L}_i (i > 1), \text{UnP}(Prog) \Rightarrow \text{UnP}(\mathcal{T}_i[Prog])$

Since transformations are defined on the structure of expressions, the proofs of these properties usually boil down to a routine inspection of the different cases. Due to their number and length, we do not describe them in this paper. A few examples of proofs are sketched in appendix B.

The source language comprises 16 skeletons, plus a number of other constructions (pairs, operators, affine expressions, ...). Further, new skeletons are added as the language gets closer to an SPMD language. Presenting the 6 compilation steps (analyses and transformations) for the whole language would be lengthy and tiresome. After presenting the complete source and target languages in the next section, we chose to focus on a tiny sublanguage for the rest of the presentation. A simple example (written in the sublanguage) is taken throughout the paper and illustrates the different steps.

The treatment of a more complete language (having at least one skeleton of each type) can be found in appendix A. The interested reader will find a description of transformations, analyses, and proofs for the whole language in [Mal98a]¹. A previous conference paper [Mal98b] focuses on the cost analysis and can be seen as a short introduction to this work.

3 Source and Target Languages

3.1 The source language \mathcal{L}_1

The source language \mathcal{L}_1 is basically a first-order, recursion-less, strict functional language, extended with a collection of higher-order functions (the skeletons). We have restricted

¹Actually, the language presented here is slightly different (it considers a larger class of mask skeletons) than the one in [Mal98a]

ourselves to a reasonable number of standard skeletons; new ones, especially among the computation and reorganization classes, could be integrated as well. The main data structure is the vector which can be nested to model multidimensional arrays.

| | | |
|------------------------|-----|--|
| Prog ₁ | ::= | Exp ₁ where Decl ₁ |
| Decl ₁ | ::= | Decl ₁ Decl ₁ <i>f</i> = Fun ₁ <i>x</i> :: Type ₁ |
| Type ₁ | ::= | (Type ₁ , ..., Type ₁) Vect LinF ₁ Type ₁ Int Float Bool |
| Exp ₁ | ::= | Fun ₁ Exp ₁ (Exp ₁ , ..., Exp ₁) <i>x</i> <i>k</i> |
| Fun ₁ | ::= | iterfor LinF ₁ Fun ₁ Op ₁ λ(<i>x</i> ₁ , ..., <i>x</i> _{<i>n</i>}).Exp ₁ <i>f</i> CompSkel ₁ ReorgSkel ₁ CommSkel ₁ MaskSkel ₁ |
| Op ₁ | ::= | + - * div exp log cos ... |
| LinF ₁ | ::= | LinF ₁ + LinF ₁ LinF ₁ - LinF ₁ <i>k</i> *LinF ₁ <i>x</i> <i>k</i> |
| CompSkel ₁ | ::= | map Fun ₁ fold Exp ₁ Op ₁ scan Exp ₁ Op ₁ |
| ReorgSkel ₁ | ::= | zip unzip concat makearray LinF ₁ |
| CommSkel ₁ | ::= | brdcast LinF ₁ transfer LinF ₁ LinF ₁ rotate LinF ₁ scatter LinF ₁ gather LinF ₁ allgather allbrdcast |
| MaskSkel ₁ | ::= | poly _{<i>n</i>} λ(<i>x</i> ₁ , ..., <i>x</i> _{<i>n</i>}).Ineq ₁ Fun ₁ Fun ₁ |
| Ineq ₁ | ::= | Ineq ₁ ∧ Ineq ₁ LinF ₁ < LinF ₁ LinF ₁ = LinF ₁ |

$x, x_1, \dots, x_n \in \text{VarIdent}$. $f \in \text{FunIdent}$. $\text{VarIdent} \cap \text{FunIdent} = \emptyset$. $k \in \text{Constant}$.

Figure 1: Skeleton language \mathcal{L}_1 .

The syntax of \mathcal{L}_1 is defined in Figure 1 (its type system will be described in Section 4). A program is a main expression followed by definitions. A definition is either a function definition or the declaration of input variables with their types. The types of input vectors bear their numerical or symbolic size. An expression (nonterminal Exp_1) is either an application, a tuple, a variable, or a constant. Functions are unary λ -abstractions, unary operators, or the predefined iterator **iterfor**. It can be defined in HASKELL[HJW92] as follows:

```
-- iterfor e f x = f(e ...,f(2,f(1,x))...)
iterfor e f x = let until p f x = if p x then x else until p f (f x)
                in snd.until(\(x,_) ->x>e)(\ (i,x) ->(i+1,f(i,x)))(1,x)
```

iterfor $n f a$ behaves like a loop; it applies n times its function argument f on a . Further, it makes the current loop index (henceforth called iterator index) accessible to its function argument.

Four classes of skeletons manipulate vectors: computation, reorganization, communication, and mask skeletons.

The *computation skeletons* are the classical higher order functions **map**, **fold** and **scan**. For example, the **fold** skeleton is defined, using the standard array library of HASKELL [HPWo99] as:

```
-- fold e f [a0;...;an] = f(...f(f(e,a0)...,an)
fold e f v = iterfor (n+1) (\(i,acc)->f(acc,v!(i-1))) e
      where n = sizeRange (bounds v) - 1
```

The four *reorganization skeletons* are **zip**, **unzip**, **makearray**, and **concat**. They allow to create and restructure vectors. The skeleton **zip** transforms a pair of vectors into a vector of pairs. The skeleton **unzip** transforms a vector of pairs into a pair of vectors. The skeleton **concat** appends its two vector arguments. The skeleton **makearray** a b creates a vector of size a of elements b , for example:

$$\text{makearray } 5 \ 1 = [1; 1; 1; 1; 1]$$

There are seven *communication skeletons* which describe families of data motion within vectors. When applied to a vector representing the parallel machine (as in \mathcal{L}_6 programs), these motions will denote (and will be implemented as) communications. They have been chosen because of their availability on parallel computers as hard-wired or optimized communication routines. The first three have type $Vect \ n \ \alpha \rightarrow Vect \ n \ \alpha$ (i.e. they map a vector of size n to a vector of the same size):

- **brdcast** $e \ v$ replaces all the elements of the vector v by the $e + 1$ th element of v . It can be expressed in HASKELL as follows:

```
-- brdcast i [a0;...;ai;...;an] = [ai;...;ai]
brdcast e v = array (0,n) [(i,v!e) |i<-[0..n]]
      where n = sizeRange (bounds v) - 1
```

- **transfer** $s \ d \ v$ updates the $d + 1$ th element of the vector v by the value of the $s + 1$ th element of v .
- **rotate** $d \ v$ performs a circular right shift of d positions on the vector v .

The next three communication skeletons have type $Vect \ m \ (Vect \ n \ \alpha) \rightarrow Vect \ m \ (Vect \ n \ \alpha)$.

- **gather** $i \ m$ copies the $i + 1$ th column of the matrix m to the $i + 1$ th row of m ,
- **scatter** $i \ m$ copies the $i + 1$ th row of the matrix m to the $i + 1$ th column of m ,
- **allgather** u transposes the matrix u .

Finally, **allbroadcast** v has type $Vect\ n\ \alpha \rightarrow Vect\ n\ (Vect\ n\ \alpha)$ and replaces all the elements of its argument v by a copy of v .

The *mask skeletons* are the only form of conditional provided by \mathcal{L}_1 . They are a family of data-parallel skeletons described as **poly** $_n$, where n is less or equal to the number of dimensions of its vector argument. That is to say, **poly** $_1$ applies on a vector of type $Vect\ n\ \alpha$ and **poly** $_k$ on a vector of type $Vect\ n_1\ (\dots\ (Vect\ n_k\ \alpha)\ \dots)$. The skeleton **poly** $_n\ p\ f_1\ f_2\ v$ applies the function f_1 to the elements of v contained in the n -dimensional polytope [Len93] described by a predicate p (a set of inequalities between affine expressions) and the function f_2 on the elements outside the polytope. For example, **poly** $_2$, which takes a matrix (vectors of vectors) as argument, can be defined in HASKELL as follows:

```
poly2 p f g v
= array (0,n-1)
      [(i,array (0,m-1) [(j,if p(i,j) then f(v!i!j) else g(v!i!j))
                        | j<-[0..m-1]]) | i<-[0..n-1]]
  where n = sizeRange (bounds v)
        m = sizeRange (bounds (v!0))
```

Note also that **poly** $_1\ (\lambda i.i = i)\ f\ Id$ is equivalent to **map** f .

In order to enable a precise symbolic cost analysis, additional syntactic restrictions are necessary. The scalar arguments of communication skeletons, of the iterfor operator, of mask skeletons, and of **makearray** must be affine expressions of variables denoting iterator indexes or sizes. This restriction is formalized by the nonterminal LinF_1 which defines affine expressions of such variables. We rely on the type system with simple subtyping (described in the next section) to ensure that variables in LinF_1 expressions are only index or size variables.

Also, if the grammar makes the first order restriction clear it does not restrict the use of recursion. The constraint that user-defined functions are not recursive must be checked separately (i.e. check that the call graph is acyclic).

Example 1 *The following simple program will be used throughout the paper to illustrate the compilation steps.*

```
f m where
m :: Vect n (Vect n (Float, Float))
f =  $\lambda m.(\text{poly}_2\ (\lambda(i,j).(i < j))\ +\ -\ (\text{brdcast}\ 0\ m), m)$ 
```

where m is a matrix of pairs of integers.

First, **brdcast** $0\ m$ builds a matrix made of copies of the first line of m (the data movements are depicted by arrows in the left matrix of Figure 2). Then, the pairs belonging to the upper triangle of the matrix are summed (represented by white dots in the right matrix of Figure 2) and the ones belonging to the lower triangle (represented by black dots) are subtracted (**poly** $_2\ (\lambda(i,j).(i < j))\ +\ -$). The result is a pair composed of the computed matrix and the initial one.

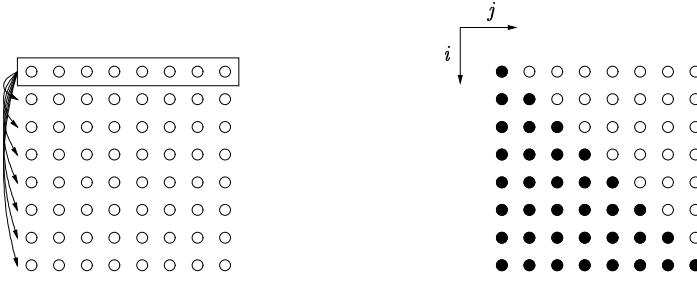


Figure 2: Data motion (**brdcast** 0) and mask skeleton (**poly**₂ ($\lambda(i, j).(i < j)$)).

The interested reader will find other examples of \mathcal{L}_1 programs in appendix C (namely, LU decomposition and the n-body problem).

3.2 The language \mathcal{L}_6

The target language \mathcal{L}_6 expresses SPMD computations. A program acts on a single vector (which is given the special type *Vectproc*) whose elements represent the processors or, more precisely, their local data spaces. An \mathcal{L}_6 program is a composition of parallel computations separated by communications. Typically a program has the following form:

$$\dots \mathbf{pimap} f \circ \mathit{Comm} \circ \mathbf{pimap} g \dots$$

where **pimap** is a version of **map** that makes the processor index accessible to its function. The language \mathcal{L}_6 introduces new versions of skeletons acting on the vector of processors. In particular, the new versions of communication skeletons express the same data motions as before but on the vector of processors. They model effective communications and will be implemented as such.

The syntax of \mathcal{L}_6 is defined Figure 3. A program is a parallel function $\mathit{FunComm}_6$ applied to a vector of processors. A parallel function is either the composition (\circ) of a parallel function with a communication, an iteration **piterfor** LinF_6 $\mathit{FunComm}_6$ (which applies LinF_6 times the parallel computation $\mathit{FunComm}_6$ to the vector of processors), or a parallel computation **pimap** Fun_6 (which applies Fun_6 on each processor). The skeleton **pimap** can be defined as:

```
-- pimap f [a0;...;ap] = [f(0,a0) ;...;f(p,ap)]
pimap f proc = array (0,p-1) [(i,f(i,(proc!i))) |i<-[0..p-1]]
      where p = sizeRange(bounds proc)
```

The nonterminal Comm_6 gathers the new versions of the communication skeletons (**pbrdcast**, **ptransfer**, **protate**, **pgather**, **pscatter**, **pallgather**, and

| | | |
|----------------------|-------|--|
| Prog_6 | $::=$ | FunComm_6 |
| FunComm_6 | $::=$ | $\text{FunComm}_6 \circ \text{Comm}_6 \mid \text{Comm}_6 \circ \text{FunComm}_6$ |
| | | \mid piterfor $\text{FunComm}_6 \mid$ pimap $\text{Fun}_6 \mid \text{Comm}_6$ |
| Comm_6 | $::=$ | pbrdcast \mid ptransfer \mid protate |
| | | \mid pscatter \mid pgather \mid pallgather \mid pallbrdcast |
| Exp_6 | $::=$ | $\text{Fun}_6 \text{Exp}_6 \mid (\text{Exp}_6, \dots, \text{Exp}_6) \mid x \mid k$ |
| Fun_6 | $::=$ | iterfor_{uc} $\text{Fun}_6 \mid \text{Op}_6 \mid \lambda(x_1, \dots, x_n).\text{Exp}_6$ |
| | | \mid $\text{CompSkel}_6 \mid \text{ReorgSkel}_6 \mid \text{CommSkel}_6 \mid \text{MaskSkel}_6$ |
| Op_6 | $::=$ | $+$ \mid $-$ \mid \dots \mid dealloc \mid copy \mid update \mid lookup |
| CompSkel_6 | $::=$ | map $\text{Fun}_6 \mid$ fold_{uc} $\text{Op}_6 \mid$ scan_{uc} Op_6 |
| ReorgSkel_6 | $::=$ | zip \mid unzip \mid concat \mid makearray_{uc} |
| CommSkel_6 | $::=$ | brdcast_{uc} \mid transfer_{uc} \mid rotate_{uc} \mid scatter_{uc} |
| | | \mid gather_{uc} \mid allgather_{uc} \mid allbrdcast_{uc} |
| MaskSkel_6 | $::=$ | poly_n $\lambda(x_1, \dots, x_n).\text{Ineq}_6 \text{Fun}_6 \text{Fun}_6$ |
| Ineq_6 | $::=$ | $\text{Ineq}_6 \wedge \text{Ineq}_6 \mid \text{LinF}_6 < \text{LinF}_6 \mid \text{LinF}_6 = \text{LinF}_6$ |
| LinF_6 | $::=$ | $\text{LinF}_6 + \text{LinF}_6 \mid \text{LinF}_6 - \text{LinF}_6$ |
| | | \mid $k * \text{LinF}_6 \mid p * \text{LinF}_6 \mid n * \text{Exp}_6 \mid x \mid k$ |
| Exp_6 | $::=$ | div $(\text{Exp}_6, p) \mid$ mod $(\text{Exp}_6, p) \mid ip$ |

$x, x_1, \dots, x_n \in \text{VarIdent}$. $ip \in \text{ProcIdent}$. $p \in \text{ProcNb}$. $k \in \text{Constant}$. $n \in \text{SizeIdent}$.

Figure 3: Skeleton language \mathcal{L}_6 .

pallbrdcast). They describe the same data motion as before but on the vector of processors. The previous communication skeletons remain in \mathcal{L}_6 in an uncurried version (**comm_{uc}**). They are used locally on the processors (they occur only inside **pimaps**) and represent local data motion.

Let us describe the communication skeleton **pbrdcast**. Before the call of this communication, each local processor memory is a 3-tuple of the form (index of the broadcasting processor, value, local memory). The broadcasting processor takes the value (second element of its triple) and send it to the other processors. The resulting vector of processors has elements of the form (broadcast value, local memory). Its functional semantics in HASKELL is:

```
-- pbrdcast [(e,a0,b0);...;(e,ap,bp)] = [(ae,b0);...;(ae,bp)]
pbrdcast proc = array (0,p-1) [let (_,_,md) = proc!i
                               in (i,(buf,md)) | i<-[0..p-1]]
  where e   = first(proc!0)
        buf = snd(proc!e)
        p   = sizeRange (bounds proc)
```

The sequential skeletons are similar to \mathcal{L}_1 . The only differences lie in the arguments of mask skeletons and new versions of communication and reorganization skeletons. The polytope argument of the mask skeleton may now include modulo and integer division (as expressed by the nonterminals LinF_6 and Expi_6 in \mathcal{L}_6). Uncurried versions of the communication and reorganizing skeletons are introduced (e.g. **brdcast_{uc}**(e, v) = **brdcast** $e v$). Furthermore, Op_6 includes four new functions: explicit vector copy and deallocation (**copy**, **dealloc**), **lookup** which accesses one element of its vector argument (**lookup** (e, v) = $v!e$) and **update** which modifies one vector element:

```
update(e,x,v)
  = array (0,n-1) [(i,if i = e then x else v!i) | i<-[0..n-1]]
  where n = sizeRange (bounds v)
```

The implementation of \mathcal{L}_6 programs is relatively straightforward. Functions inside **pimaps** represent the sequential programs to be executed by all processors in parallel. They will be compiled into sequential code using standard techniques. Communication skeletons (such as **pbrdcast**) will be implemented as calls to a standard communication library.

Example 2 *After a row-block distribution, our simple example is transformed into the following \mathcal{L}_6 program:*

$$\begin{aligned} & \mathbf{pimap} \lambda(ip, (buf, x)). (\mathbf{poly}_2 (\lambda(i, j). ip * b + i < j) + - \\ & \quad \circ \mathbf{makearray}_{uc}(n, buf), x) \\ & \circ \mathbf{pbrdcast} \circ \mathbf{pimap} \lambda(ip, x).(0, \mathbf{lookup}(0, x), x) \end{aligned}$$

where the input matrix is supposed to have n lines and b is the block size (i.e. n divided by the number of processors).

The first parallel computation prepares the communication (`pbroadcast`) which broadcasts the first row of the matrix to all the processors. Then, the vector creation (`makearrayuc`) and the mask skeleton (`poly2`) are applied locally in parallel.

4 Type and Size Inference

Type/Size inference has several objectives:

- it checks that programs are well typed.
- it ensures that all vector computations are well defined. For example, `brdcast` is defined only if its index argument is within the bounds of its vector argument.
- it makes sure that some arguments are affine expressions depending only on vector sizes and iterator indexes.
- it computes the symbolic size of each vector expression.

Size inference can be seen both as a static analysis to compute the size of each vector-typed expression and as a type system to enforce constraints. The static and symbolic evaluation of sizes is made possible by several restrictions:

1. the argument and result of `iterfor` have the same size,
2. indexes and sizes, used by accesses and creations, are affine expressions,
3. vectors are homogeneous; that is, all the elements of a vector have the same size.

Types are represented as *size types* whose syntax is described in Figure 4. Vectors are associated with their size and basic types are either scalar types (integers *Int*, floating point numbers *Float*, or booleans *Bool*), *Size*, or *Index* types. Intuitively, an expression has type *Size* if it is a constant, the size of an input vector, or an affine expression of *Size* variables. Similarly, an expression has type *Index* if it is an iterator index, has type *Size* or is an affine expression *Index* variables. *Size* and *Index* types bear an affine expression denoting their symbolic value.

For example, the size type $Size^n \rightarrow Vect\ n\ Int$ indicates that the function takes an integer n and returns a vector of size n .

The type system integrates subtyping based on the following hierarchy:

$$Size^a \subseteq Index^a \subseteq Int \subseteq Float$$

The subtype relation $Size^a \subseteq Index^a$ indicates that loop indexes may depend on sizes but vector sizes cannot be defined in terms of loop indexes. The subtyping rules for functions, pairs, etc. are the usual ones. The inference rules, described in Figure 5, are of the form

$$C, \Gamma \vdash e : T, C_1$$

$$\begin{aligned}
T & ::= T_{\text{Exp}} \rightarrow T_{\text{Exp}} \\
T_{\text{Exp}} & ::= (T_{\text{Exp}}, T_{\text{Exp}}) \mid \text{Vect } A \ T_{\text{Exp}} \mid \alpha \mid B \\
B & ::= \text{Int} \mid \text{Float} \mid \text{Bool} \mid \text{Index}^A \mid \text{Size}^A \\
A & ::= A + A \mid A - A \mid k * A \mid x \mid k \\
& x \in \text{SizeVar}, \alpha \in \text{SizeType} \text{ and } k \in \text{Constant}.
\end{aligned}$$

Figure 4: Size types

which means that e has size type T with size constraints C_1 in the environment Γ with the subtyping constraints C . If the size constraints C_1 are satisfied by the input parameters then the evaluation of e will not produce any vector access error.

Some rules just express unification in terms of an equality constraint between size types. For example, in the rule [APPL], the constraint $\alpha = \gamma$ unifies the actual and formal parameters. Some others adds constraints on sizes.

- The rule [CONST] gives the smallest (most precise) type according to the subtyping relation (i.e. *Size*) to constants.
- The type of an addition or a subtraction of affine expressions (rule [LINF₁]), is the most general type of its two subexpressions. Further, it introduces a constraint denoting the value of the result as an affine expression. Note that the rule for other operators (e.g. multiplication) must be different since they do not preserve affinity in general.
- The rule [BRD] forces the first argument of `brdcast` to have type *Index*. It also introduces the constraints $(0 \leq s_1 \leq s)$ which enforce that this index is within the bounds of the second (vector) argument.
- The rule [POLY2] states that the polytope must be defined using arguments of type *Index*. It also constraints the two functions to have the same type (including sizes if they have a vector argument) $(\alpha_1 = \alpha_2 \text{ and } \beta_1 = \beta_2)$.

As described the type system is monomorphic. This can be too harsh a limitation since user defined functions could be used only for specific vector sizes. However, since user defined functions are not recursive, all the calls can be replaced by the corresponding function definitions. Such a preliminary unfolding amounts to making the type system polymorphic.

The rules in Figure 5 constitute a small representative subset of the inference system. The rules for other skeletons are similar and the rules for subtyping are the usual ones. The inference rules can be turned into an algorithm using standard techniques of subtype inference (e.g. [Mit91]) and of polyhedral computations (e.g. [Wil93]).

Satisfiability is made easy by the fact that constraints are inequalities between affine expressions. The solutions of a set of such constraints can be seen as the points of a convex

$$\frac{C, \Gamma \vdash e : \alpha, C_1 \quad C, \Gamma \vdash f : \gamma \rightarrow \beta, C_2}{C, \Gamma \vdash fe : \beta, C_1 \cup C_2 \cup \{\alpha = \gamma\}} \text{[APP]} \quad \frac{}{C, \Gamma \cup \{y : \alpha\} \vdash y : \alpha, \{\}} \text{[VAR]}$$

$$\frac{C, \Gamma \vdash e_1 : \alpha_1, C_1 \quad C, \Gamma \vdash e_2 : \alpha_2, C_2}{C, \Gamma \vdash (e_1, e_2) : (\alpha_1, \alpha_2), C_1 \cup C_2} \text{[PAIR]} \quad \frac{}{C, \Gamma \vdash k : \text{Size}^k, \{\}} \text{[CONST]}$$

$$\frac{\forall k \in [1, 2] \ C, \Gamma \vdash e_k : \alpha^{s_k}, C_k \quad C \vdash \alpha \subseteq \text{Index}}{C, \Gamma \vdash e_1 \text{ Op } e_2 : \alpha^{s_3}, C_1 \cup C_2 \cup \{s_1 \text{ Op } s_2 = s_3\}} \text{[LINF}_1\text{]}$$

where $\text{Op} \in \{+, -\}$

$$\frac{C, \Gamma \vdash e : \beta, C_1 \quad C \vdash \beta \subseteq \alpha}{C, \Gamma \vdash e : \alpha, C_1} \text{[COER]}$$

$$\frac{C, \Gamma \vdash e : \text{Index}^{s_1}, C_1}{C, \Gamma \vdash \text{brdcast } e : \text{Vect } s \ \alpha \rightarrow \text{Vect } s \ \alpha, C_1 \cup \{0 \leq s_1 < s\}} \text{[BRD]}$$

$$\frac{\forall k \in [1, 2] \ C, \Gamma \cup \{i : \text{Index}, j : \text{Index}\} \vdash e_k : \text{Index}, _ \quad C, \Gamma \vdash f_k : \alpha_k \rightarrow \beta_k, C_k}{C, \Gamma \vdash \text{poly}_2 \ \lambda(i, j). e_1 \leq e_2 \ f_1 \ f_2 : \text{Vect } s_1 (\text{Vect } s_2 \ \alpha_1) \rightarrow \text{Vect } s_1 (\text{Vect } s_2 \ \beta_1), C_1 \cup C_2 \cup \{\alpha_1 = \alpha_2, \beta_1 = \beta_2\}} \text{[POLY}_2\text{]}$$

Figure 5: Size inference (extract).

polyhedron. The dimension of this polyhedron are the variables denoting sizes and occurring in the inequalities. These variables can be the size variables introduced by the inference (denoted by s_i in the rules) or the symbolic sizes used to express input vector sizes. If s_1, \dots, s_p are the p size variables and n_1, \dots, n_q the q symbolic sizes occurring in the constraints, the tuple $(s_1, \dots, s_p, n_1, \dots, n_q)$ belongs to the convex polyhedron defined by the affine inequalities.

Satisfiability amounts to normalizing the constraints to get only inequalities of affine expressions and checking that the obtained polyhedron is not empty. The expression of a size variable s is found by projecting the polyhedron on the dimensions n_i of the symbolic sizes and the dimension corresponding to s . This projection permits to express s as an affine expression of symbolic sizes. Note that testing for emptiness and projection are standard operations of polyhedral libraries (e.g. [Wil93]).

If the set of constraints is satisfiable, then no vector error will occur at runtime. In this case, each size variable is determined and we use size types to annotate each program expression (i.e. sizes are symbolically expressed in terms of the size of input vectors).

Example 3 *Let us illustrate the size analysis on our small example:*

$$(\mathbf{poly}_2 (\lambda(i, j).i < j) + - (\mathbf{brdcast} 0 m), m)$$

We have

$$\begin{aligned} m &:: \mathit{Vect} \ n \ (\mathit{Vect} \ n \ (\mathit{Float}, \mathit{Float})) \\ \mathbf{poly}_2 (\lambda(i, j).i < j) + - &: \mathit{Vect} \ s_1 \ (\mathit{Vect} \ s_2 \ (\mathit{Float}, \mathit{Float})) \\ &\quad \rightarrow \mathit{Vect} \ s_1 \ (\mathit{Vect} \ s_2 \ \mathit{Float}) \\ (\mathbf{brdcast} 0 m) &: \mathit{Vect} \ s_3 \ (\mathit{Vect} \ s_4 \ (\mathit{Float}, \mathit{Float})) \end{aligned}$$

with the following set of constraints:

$$\{s_1 = s_3, s_3 = n, s_2 = s_4, s_4 = n, 0 \leq 0 \leq s_2\}$$

By projecting on the n dimension we get: $s_1 = n, s_2 = n, s_3 = n, s_4 = n$ and, for example, the mask skeleton has the size type:

$$\begin{aligned} \mathbf{poly}_2 (\lambda(i, j).i < j) + - &: \mathit{Vect} \ n \ (\mathit{Vect} \ n \ (\mathit{Float}, \mathit{Float})) \\ &\quad \rightarrow \mathit{Vect} \ n \ (\mathit{Vect} \ n \ \mathit{Float}) \end{aligned}$$

5 Update-in-place Analysis

As we use a pure functional language and want to manipulate vectors as efficiently as possible, we have to tackle the update-in-place problem. The solution must be precise and fully automatic. But, since this step may have a tremendous impact on runtime costs, it should also yield useful feedback and allow programmers to keep complete control of their programs. As a consequence, the update-in-place analysis is the first step of our compilation chain. The programmer is only able to read and change programs at this stage.

Our approach relies on extracting an abstract representation of possible execution traces which are then used to check whether all vectors can be modified in place or not. The abstract execution trace represents the vector access sequences carried out during the program evaluation. If these access sequences do not satisfy some criteria, two possibilities arise:

- either **copy** operations are automatically inserted in the program so that the criteria hold. Functionally, this new operator is the identity function but will be implemented by copying its vector argument.
- either the faulty subexpressions and violated criteria are pointed out to the programmer who may insert copies manually or restructure the program before a new analysis.

The same abstract sequences are used to ensure that that all vectors are either returned as result or explicitly deallocated (using a new operation **dealloc**). In the following, we present the update-in-place analysis (at a fairly high and intuitive level) and return to the deallocation problem at the end of this section.

In order to check the updated in place property, we statically compute a runtime trace which corresponds to the sequence of vector accesses during execution of program. This information is similar to the sequences calculated in [KS86] which, in general, are described by grammars.

Each vector is annotated by labels (e.g. l_1) which serve to represent sharing. The basic vector accesses may be a read (l_1^r denotes a read of the vector with label l_1), a write (l_1^w) which produces a vector annotated with a fresh label (here l_2), a copy (l_1^c where l_2 is a fresh label), a vector deallocation (l_1^f), or vector display (l_1^p) which indicates that the vector is returned in the final result. A concrete access sequence is either a sequential composition of sequences ($EV_2 \circ EV_1$, where the accesses of EV_1 are made before the accesses of EV_2), a parallel composition ($EV_1 \parallel EV_2$ which represents any interleaving of EV_1 and EV_2 accesses), or a vector access.

In order to represent finitely all the possible sequences of accesses, abstract access sequences are described as regular expressions of concrete sequences. An abstract sequence can be a concrete sequence, the union of abstract sequences $EV_1 + EV_2$ (i.e. EV_1 or EV_2), or the repeated composition of abstract sequences $(EV_1 \circ)^* EV_2$ (i.e. EV_2 followed by zero or more occurrences of EV_1). In [Mal98a], we present computation of the access sequences as an abstraction of an instrumented semantics for the language \mathcal{L}_1 . This allows us to prove the correctness of the static analysis with respect to the instrumented semantics.

Intuitively, vectors can be safely updated in place if no access of the form l_1^w , l_1^f , or l_1^p is followed by an access on the form l_1^r , l_1^c , l_1^w , l_1^f , or l_1^p in the program access sequence. In other words, after a write to a vector (say, labeled l_1) there should not be accesses to the previous version (i.e. l_1). If this holds, it is clear that the update/write can be done in place.

We say that an access follows another access if they appear, either in a sequential composition (\circ), or, since it does not enforce any specific order, in a parallel composition (\parallel). The update-in-place condition amounts to checking for each created, copied or modified vector that no access follows a write on this same vector. Let EV be the access sequence of the

program, the condition is expressed formally by

$$\mathcal{UP}(EV) = \forall l. \mathcal{UP}_l(EV)$$

The condition for a specific vector (\mathcal{UP}_l) is defined in Figure 6. The condition \mathcal{UP}_l for an access sequence $EV_1 \circ EV_2$, and a l -labeled vector, is that either the vector is not written in EV_2 and the condition holds on EV_1 , or the vector is written in EV_2 and the vector is not accessed in EV_1 and the condition holds on EV_2 . In the parallel composition case, the condition is that either one sequence does not contain an access to l and the condition holds for the other, or no write to l occurs in both sequences. Finally, the condition \mathcal{UP}_l holds on the sequence $EV_1 + EV_2$ if it holds on both sequences EV_1 and EV_2 . For a sequence of the form $(EV_1 \circ)^* EV_2$, it holds if, either the vector is not accessed in EV_2 and the condition holds in EV_1 , either the vector is not written in EV_1 and EV_2 .

$$\begin{aligned} \mathcal{UP}_l(EV_1 \circ EV_2) &= (\mathcal{UP}_l(EV_2) \wedge \beta l^A \in EV_1) \vee (\mathcal{UP}_l(EV_1) \wedge \beta l^W \in EV_2) \\ \mathcal{UP}_l(EV_1 \parallel EV_2) &= (\mathcal{UP}_l(EV_2) \wedge \beta l^A \in EV_1) \vee (\mathcal{UP}_l(EV_1) \wedge \beta l^A \in EV_2) \\ &\quad \vee (\beta l^W \in (EV_1 \parallel EV_2)) \\ \mathcal{UP}_l(EV_1 + EV_2) &= \mathcal{UP}_l(EV_1) \wedge \mathcal{UP}_l(EV_2) \\ \mathcal{UP}_l((EV_1 \circ)^* EV_2) &= (\beta l^W \in EV_1 \wedge \beta l^W \in EV_2) \vee (\mathcal{UP}_l(EV_2) \wedge \beta l^A \in EV_1) \\ \mathcal{UP}_l(_) &= \text{true} \quad \text{otherwise} \end{aligned}$$

with $\mathcal{A} \in \{p, c, r, w, f\}$, $\mathcal{W} \in \{w, f, p\}$.

Figure 6: Conditions on Sequences of Accesses

Programs can be automatically transformed so that they respect the property \mathcal{UP} . Any sequence not satisfying the property \mathcal{UP} contains accesses to a vector (annotated by) l following a write to l . Since each access is associated to a unique operation, the faulty writes are easy to find. To ensure the update in place property, it is sufficient to insert an explicit `copy` before each such write operation in the program.

Another possibility is to let the user restructure the program in order to make the program respect the property. This way, the user keeps complete control of the costs induced by the program and may potentially minimize the number of necessary copies.

Example 4 *The analysis of our example*

$$(\mathbf{poly}_2 (\lambda(i, j). i < j) + - (\mathbf{brdcast} \ 0 \ m), m)$$

produces the following access information:

$$(l_3^p \parallel l_3^p \parallel l_1^p \parallel l_1^p) \circ (l_3^w \parallel l_2^w \parallel l_3^w \parallel l_2^w) \circ (l_2^w \parallel l_1^w \parallel l_2^w \parallel l_1^w)$$

where the outer vector of m is annotated by the label l_1 , and its inner vectors by l_1' .

The condition UP is violated for the vector annotated l_1 which is written and returned as result as well. So, an explicit copy is inserted just before the write on l_1 (done by the skeleton `brdcast`) leading to:

$$(\mathbf{poly}_2 (\lambda(i, j).i < j) + - (\mathbf{brdcast} 0 (\mathbf{copy} m)), m)$$

The functional language community has proposed many different update-in-place analyzes. They are either syntactic criteria [Sch85], semantics-based analyses [KS86, Ses89], or type systems [Wad90, GH90]. Most of them analyze whether a function parameters or instances of specific types can be implemented as a global variable. In our context, we want to check that each vector (taken as input or dynamically created) could be updated in place. The type-based approach of [Wad90] meets this requirement, however it was not precise enough to prove that the property still held on programs obtained after the distribution transformation. Our main motivation to design a more precise analysis was to be able to prove that the update-in-place property was preserved by the successive transformations.

The same abstract sequences are used to enforce that vectors are explicitly deallocated. This allows us to avoid the need for a garbage collector. This is particularly important to guarantee that our cost analysis evaluates real runtime costs. On abstract sequences, the criteria are that every vector is either written, deallocated, or displayed as result. Any vector which is not accessed anymore must be explicitly deallocated. If the criteria are not satisfied, `dealloc` operations are automatically inserted just after the last use of vectors. One remaining problem is that pairs are allocated in the heap. The language \mathcal{L}_1 being based on vectors, the memory allocated for pairs is quite small in practice and our implementation does not include a garbage collector. However, it would be more satisfactory either to explicitly deallocate pairs (as we do for vectors), or to implement them in the runtime stack.

6 Making Communications Explicit

Recall that the target language corresponds to SPMD code, that is to say, a program is a sequence made of local computations (a function applied by each processor) followed by communications (a communication skeleton) between processors. All data appearing in computations must be local to the processors. In programming terms, it means that a function to be executed in parallel must be closed. However, there might be expressions in \mathcal{L}_2 including free variables that induce communications not expressed by communication skeletons. For example, in the expression

$$\mathbf{poly}_2 (\lambda(i, j).i < j) (\mathbf{fold} e +) (\mathbf{fold} 0 +) m$$

the variable e occurs free in an expression `(fold e +)` supposed to be executed locally on each processor.

The transformation \mathcal{EC} produces expressions in the intermediate language \mathcal{L}_3 , very similar to \mathcal{L}_2 , but in which no free variable occurs in the functional arguments of `map` or `polyn`. It aims at making parallel functions closed and is related to λ -lifting [Joh85]. The transformation \mathcal{EC} is defined by local transformations applied iteratively to the program until a

fixpoint is reached. In Figure 7, which describes the rule for **poly**₂, C denotes a context and $\mathcal{FV}[\![Fun]\!]$ denotes the free variable of Fun .

The skeleton **poly**₂ is transformed so that its free variables are abstracted. Each element of the matrix argument is now associated with the free variables values. This is done using the size (type) of the matrix argument $m \times n$ (previously inferred) and a composition of the skeletons **makearray**, **map**, and **zip**.

When the free variables are vectors, they must be explicitly deallocated after function application to preserve the explicit deallocation property (Section 5).

$$\begin{aligned} & \mathcal{EC}[\![C[\mathbf{poly}_2 \text{ P Fun}_1 \text{ Fun}_2]\!]\!] \\ &= C[\lambda v. \mathbf{poly}_2 \text{ P } (\lambda((x_1, \dots, x_p), x). \mathbf{dealloc} (x_1, \dots \mathbf{dealloc} (x_q, \text{Fun}_1 x) \dots)) \\ & \quad (\lambda((x_1, \dots, x_p), x). \mathbf{dealloc} (x_1, \dots \mathbf{dealloc} (x_q, \text{Fun}_2 x) \dots)) \\ & \quad (\mathbf{map} \mathbf{zip} (\mathbf{zip} (\mathbf{makearray} m (\mathbf{makearray} n (x_1, \dots, x_p)), v)))] \\ & \text{with } \mathbf{poly}_2 \text{ P Fun}_1 \text{ Fun}_2 : \text{Vect } m (\text{Vect } n \alpha) \rightarrow \text{Vect } m (\text{Vect } n \beta) \\ & \text{and } \mathcal{FV}[\![\text{Fun}_1]\!] \cup \mathcal{FV}[\![\text{Fun}_2]\!] = \{x_1, \dots, x_p\}, 1 \leq p, \\ & \text{and } \forall 1 \leq i \leq q, x_i : \text{Vect } n_i \alpha, \forall q+1 \leq i \leq p, x_i : \alpha \neq \text{Vect } n_i \alpha \end{aligned}$$

Figure 7: Transformation \mathcal{EC} (extract).

Example 5 Assuming that the matrix m has size $n \times n$, the above example

$$\mathbf{poly}_2 (\lambda(i, j). i < j) (\mathbf{fold} e +) (\mathbf{fold} 0 +) m$$

is transformed into

$$\begin{aligned} & (\lambda v. \mathbf{poly}_2 (\lambda(i, j). i < j) (\lambda(e, v). \mathbf{fold} e + v) (\lambda(e, v). \mathbf{fold} 0 + v) \\ & \quad (\mathbf{map} \mathbf{zip} (\mathbf{zip} (\mathbf{makearray} n (\mathbf{makearray} n e), v)))) m \end{aligned}$$

A $n \times n$ matrix of integers e is built and zipped with m . The distribution of the value e to local processors will be made explicit by the next compilation step which distributes matrices and introduces communication skeletons.

7 Distribution

This step is decomposed in three transformations. The transformation *ABS* replaces program variables by combinators. *DIST* transforms the resulting program according to a distribution choice for each input vector. Afterwards, some optimizations become possible and are described as program transformations. This transformation chain produces SPMD programs and can be described as

$$\mathcal{L}_3 \xrightarrow{ABS} \mathcal{L}_4 \xrightarrow{DIST} \mathcal{L}_5 \xrightarrow{OPT} \mathcal{L}_6$$

7.1 ABS Transformation

The transformation ABS is comparable to the abstraction algorithms used to compile the β -reduction with combinators [Tur79]. Each expression E with is transformed into an expression such that

$$(ABS[[E]]\vec{X})\vec{X} = (E, \vec{X})$$

where \vec{X} is made of nested pairs representing the free variables of E . Initially \vec{X} represents the program input variables. More generally, \vec{X} represents the environment which is explicitly threaded throughout the program (taken as argument and returned as result by each sub-expression). This is the data structure that will be distributed over the vector of processors. The transformation ABS also unfolds the program. Each function call is replaced by its definition. A \mathcal{L}_4 program is a variable-less, call-less expression.

The rules of ABS for our sublanguage are given in Figure 8. In order to express the transformation, new functions are introduced. First, as can be expected, environment management is expressed using combinators. The family of *restructuring combinators* $\mathbf{extract}^{\vec{X}, \vec{Y}}$ (defined as $\lambda\vec{X}.\vec{Y}$) permits to restructure the environment. The function \mathbf{ftuple}_2 takes two functions and a pair and applies the first (resp. second) function to the first (resp. second) pair component. Composition \circ ($Fun_1 \circ Fun_2 = \lambda x.Fun_1(Fun_2 x)$) and a curried pair operator ($\mathbf{pair} = \lambda x.\lambda y.(x, y)$) are introduced. Finally, new uncurried versions of functions are needed (e.g. $\mathbf{brdcast}_{uc}(e, v) = \mathbf{brdcast} e v$).

The transformation ABS propagates variable values to the places they are used. An access to a variable x now returns a pair made of its value extracted from the environment \vec{X} and the environment itself. For functions with several arguments (e.g. $\mathbf{brdcast}$), the initial function is substituted by its version with one argument. For example, the expression $\mathbf{brdcast} v x$ is transformed into $\mathbf{brdcast}_{uc}(v, x)$.

Note that, due to the previous transformation \mathcal{EC} , the rule for \mathbf{poly}_2 does not apply recursively ABS on Fun_1 and Fun_2 because they are closed.

Example 6 *Our simple example*

$$(\mathbf{poly}_2 (\lambda(i, j).i < j) + - (\mathbf{brdcast} 0 (\mathbf{copy} m)), m)$$

is transformed into the variable-less expression

$$\begin{aligned} & \mathbf{ftuple}_2 (\mathbf{poly}_2 (\lambda(i, j).i < j) + - \circ \mathbf{brdcast}_{uc} \circ \mathbf{pair} 0 \circ \mathbf{copy}) id \\ & \circ \mathbf{extract}^{x, (x, x)}(m) \end{aligned}$$

It duplicates (combinator $\mathbf{extract}^{x, (x, x)}$) its argument, then applies the composition of skeletons ($\mathbf{poly}_2 \dots \circ \dots \circ \mathbf{copy}$) to the first component and the identity to the second one (returned as result).

7.2 DIST Transformation

A distribution d can be seen as a function restructuring vectors. Intuitively, transforming the program P according to the distribution d amounts to starting from the equivalent program

$$\begin{aligned}
\mathcal{ABS}[\text{Fun Exp}] \vec{X} &= \mathcal{ABS}[\text{Fun}] \vec{X} \circ \mathcal{ABS}[\text{Exp}] \vec{X} \quad \text{if Fun} \neq \text{brdcast} \\
\mathcal{ABS}[\text{brdcast Exp}] \vec{X} &= \text{ftuple}_2 \text{ brdcast}_{\text{uc}} \text{ Id} \circ \text{extract}^{(a,(b,c)),((a,b),c)} \\
&\quad \circ \text{ftuple}_2 \text{ Id} (\mathcal{ABS}[\text{Exp}] \vec{X}) \\
\mathcal{ABS}[(E_1, E_2)] \vec{X} &= \text{extract}^{((a_1, b_1), (a_2, b_2)), ((a_1, a_2), b_1)} \\
&\quad \circ \text{ftuple}_2 (\mathcal{ABS}[E_1] \vec{X}) (\mathcal{ABS}[E_2] \vec{X}) \circ \text{extract}^{x,(x,x)} \\
\mathcal{ABS}[x] \vec{X} &= \text{extract}^{\vec{x},(x,\vec{x})} \quad (x \in \vec{X}) \\
\mathcal{ABS}[k] \vec{X} &= \text{pair } k \\
\mathcal{ABS}[\text{copy}] \vec{X} &= \text{ftuple}_2 \text{ copy Id} \\
\mathcal{ABS}[\text{poly}_2 P \text{ Fun}_1 \text{ Fun}_2] \vec{X} &= \text{ftuple}_2 (\text{poly}_2 P \text{ Fun}_1 \text{ Fun}_2) \text{ Id}
\end{aligned}$$

Figure 8: Transformation \mathcal{ABS} (extract).

$P \circ d^{-1} \circ d$ and pushing d^{-1} to the left until the program is of the form $d'^{-1} \circ P' \circ d$. After transformation, the program (P') takes and returns a single vector whose elements can be seen as the local memory of each processor. We give to this data structure the special type *Vectproc*.

Distributions

The data distributions define the allocation of data on the processors. They are functions of the type $\alpha \rightarrow \text{Vectproc } p \alpha$ where p is the number of processors.

We consider a fixed set of standard distributions. For a vector, there are three distributions: **block**, **cyc** and **seq**.

- The distribution **block** p breaks up the vector in p blocks of contiguous elements and allocates each block to a processor (e.g. **block** 2 [1; 2; 3; 4] = [[1; 2]; [3; 4]]).
- The distribution **cyc** p distributes cyclicly each vector element on the p processors (e.g. **cyc** 2 [1; 2; 3; 4] = [[1; 3]; [2; 4]]).
- **seq** yields a vector with a single processor containing the data vector (e.g. **seq** [1; 2; 3; 4] = [[1; 2; 3; 4]]).

They can be described using the general distribution **dist** defined as

```

dist b p v
= array (0,p-1)
  [(i,array (0,m-1) [(j,v!((div j b)*p*b+i*b+mod j b))
                    |j<-[0..m-1]]) |i<-[0..p-1]]
  where m = div (rangeSize (bounds v)) p

```

This function distributes cyclicly blocks of size b on p processors. So, let s be the size of v , we have:

- **block** $p v = \mathbf{dist} (\mathbf{div} s p) p v$,
- **cyc** $p v = \mathbf{dist} 1 p v$,
- **seq** $v = \mathbf{dist} s 1 v$.

A degenerate case is the distribution of a scalar data. The distribution **const** p allocates its scalar argument to each processor in a vector of processors of size p . For example, **const** 5 0 = [0; 0; 0; 0; 0].

These distributions can be combined using higher-order functions (**dp** and **de**) to deal with nested vectors and pairs.

The distribution **dp** $d_1 d_2$ allocates a pair on the processors. The first pair component is allocated according to the distribution d_1 and the second one according to d_2 . The result is a vector of processors containing pairs. For example,

$$\mathbf{dp} \mathbf{seq} \mathbf{seq} ([1; 2], [10; 20]) = [[([1; 2], [10; 20])]]$$

The distribution **de** $d_1 d_2$ allocates the top-level vector according to d_1 and the inner vectors according to d_2 . This distribution takes a vector of vectors and returns a processor vector of vectors of vectors. For matrices (vector of vectors), this entails 9 different distributions: sequential (**de seq seq**), row block (**de (block p) seq**), ...

Example 7 A row cyclic distribution on two processors:

$$\mathbf{de} (\mathbf{cyc} 2) \mathbf{seq} [[1; 2]; [3; 4]; [5; 6]; [7; 8]] = [[([1; 2]; [5; 6]); ([3; 4]; [7; 8])]]$$

A row cyclic, column block distribution on 4 processors:

$$\mathbf{de} (\mathbf{cyc} 2) (\mathbf{block} 2) [[1; 2]; [3; 4]; [5; 6]; [7; 8]] = [[([1]; [5]); ([2]; [6]); ([3]; [7]); ([4]; [8])]]$$

Each distribution defines a bijection. We write **blocki**, **cyci**, **seqi**, **consti**, **dei**, and **dpi** for the corresponding inverse distributions.

Transformation Rules

The transformation *DIST* assumes that *ABS* has been applied and that the program is of the form $Fun \vec{X}$ where \vec{X} represent the program input parameters. Let d be the distribution of input parameters considered, then Fun is rewritten into the equivalent program $Fun \circ d^{-1} \circ d$. We do not consider different choices of distribution for dynamically allocated vectors. Vectors created by **makearray_{uc}** are allocated on a single processor and copied vectors (produced by **copy**) are distributed as their copy. Considering other distributions for these vectors would potentially entail communications. In this respect, it is very similar to the problem of redistribution that we mention in the future work section (Section 10.3).

Figure 9 presents the transformation rules required by our simple example. It uses functions (**pimap**, **pbrdcast**, **lookup**, ...) presented along with the target language in Section 3.2. The rules propagate the inverse distribution d^{-1} to the left. Note that since that an inverse distribution d^{-1} has type $Vectproc\ p\ \alpha \rightarrow \alpha$, a function Fun having the type $\tau_1 \rightarrow \tau_2$ is transformed into a function $FunComm$ with type $Vectproc\ p\ \tau_1 \rightarrow Vectproc\ p\ \tau_2$.

$$\begin{aligned}
& \mathcal{DIST}\llbracket F_1 \circ F_2 \circ di \rrbracket \\
&= \text{let } di' \circ F_2' = \mathcal{DIST}\llbracket F_2 \circ di \rrbracket \\
&\quad \text{in } \mathcal{DIST}\llbracket F_1 \circ di' \rrbracket \circ F_2' \\
& \mathcal{DIST}\llbracket \mathbf{extract}^{x,(x,x)} \circ di \rrbracket \\
&= \mathbf{dpi}\ di\ di \circ \mathbf{pimap}\ (\mathbf{extract}^{(ip,x),(x,x)}) \\
& \mathcal{DIST}\llbracket Fun \circ di \rrbracket \\
&= di \circ \mathbf{pimap}\ (Fun \circ \mathbf{extract}^{(ip,x),x}) \quad \text{where } Fun \in \{\mathbf{copy}, \mathbf{Id}\} \\
& \mathcal{DIST}\llbracket \mathbf{pair}\ k \circ di \rrbracket \\
&= \mathbf{dpi}\ (\mathbf{consti}\ p)\ di \circ \mathbf{pimap}\ (\mathbf{pair}\ k \circ \mathbf{extract}^{(ip,x),x}) \\
&\hspace{15em} \text{with } di : Vectproc\ p\ \alpha \rightarrow \alpha \\
& \mathcal{DIST}\llbracket \mathbf{poly}_2\ P\ F_1\ F_2 \circ \mathbf{dei}\ (\mathbf{blocki}\ p)\ \mathbf{seqi} \rrbracket \\
&= \mathbf{dei}\ (\mathbf{blocki}\ p)\ \mathbf{seqi} \\
&\quad \circ \mathbf{pimap}\ (\lambda(ip, v). \mathbf{poly}_2\ (P \circ \lambda(i_1, i_2). (ip * b + i_1, i_2))\ F_1\ F_2\ v) \\
&\text{with } \mathbf{blocki}\ p : Vectproc\ p\ (Vect\ b\ \alpha) \rightarrow Vect\ n\ \alpha \\
& \mathcal{DIST}\llbracket \mathbf{brdcast}_{uc} \circ \mathbf{dpi}\ (\mathbf{consti}\ p)\ (\mathbf{dei}\ (\mathbf{blocki}\ p)\ \mathbf{seqi}) \rrbracket \\
&= \mathbf{dei}\ (\mathbf{blocki}\ p)\ \mathbf{seqi} \\
&\quad \circ \mathbf{pimap}\ (\lambda(ip, (buf, (e, v))). \mathbf{brdcast}_{uc}(0, (\mathbf{update}\ (0, buf, v)))) \\
&\quad \circ \mathbf{pbrdcast} \circ \mathbf{pimap}\ (\lambda(ip, (e, v)). (\mathbf{div}\ (e, b), \mathbf{lookup}\ (\mathbf{mod}\ (e, b), v), (e, v))) \\
&\text{with } \mathbf{blocki}\ p : Vectproc\ p\ (Vect\ b\ \alpha) \rightarrow Vect\ n\ \alpha \\
& \mathcal{DIST}\llbracket \mathbf{ftuple}_2\ F_1\ F_2 \circ \mathbf{dpi}\ di_1\ di_2 \rrbracket \\
&= \text{let } di'_1 \circ F'_1 = \mathcal{DIST}\llbracket F_1 \circ di_1 \rrbracket \\
&\quad di'_2 \circ F'_2 = \mathcal{DIST}\llbracket F_2 \circ di_2 \rrbracket \\
&\quad \text{in } \mathbf{dpi}\ di'_1\ di'_2 \circ \mathcal{FUS}\llbracket (F'_1, F'_2) \rrbracket
\end{aligned}$$

Figure 9: Transformation \mathcal{DIST} (extract).

The rule for the composition $F_1 \circ F_2 \circ di$ consists in transforming first the function $F_2 \circ di$. This yields an equivalent expression of the form $di' \circ F'_2$. Then, the transformation is recursively called on $F_1 \circ di'$. The final expression is of the form $di'' \circ F'_1 \circ F'_2$ where F'_1 and F'_2 are SPMD functions acting on a unique vector of processors.

The transformation of the restructuring combinator $\mathbf{extract}^{x,(x,x)}$ (which duplicates its argument) consists in duplicating the local memory of each processor. The resulting inverse distribution is an inverse distribution that distributes its pair argument identically (**dpi** di).

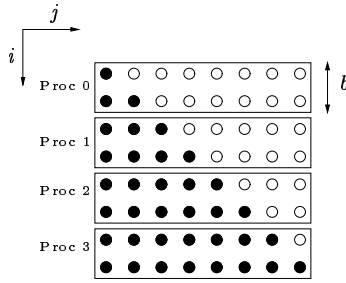


Figure 10: Impact of a row block distribution on $\mathbf{poly}_2 (\lambda(i, j).i < j)$

The rules for `copy` and the identity function `Id` consist in applying it to the local memory of each processor. The inverse distribution is unchanged. In the same way, distributing the `pair k` function (where k is a constant) consists in applying the function to each local memory. The resulting inverse distribution distributes the first component of its pair argument according to `consti p` (where p is the number of processors) and the second according to the original distribution (`di`).

The rule for the mask skeleton \mathbf{poly}_2 modifies the index in the inequalities describing the polytope (P) according to the distribution. After transformation, each processor applies the mask skeleton to its local vector. The inequalities still define a polytope thanks to the restricted set of distributions. This key property would not hold for more general forms of distribution. Figure 9 presents the transformation rule for a row block distribution only (a more complicated generic rule exists also). Note that the symbolic size computed by the size inference (Section 4) is used by this rule. The rule is based on the fact that an element whose indexes are (i_1, i_2) in the distributed matrix has indexes $(ip * b + i_1, i_2)$ in the initial matrix (where ip denotes the processor index and b the size of blocks, that is the number of rows n divided by the number of processors p). For example, the expression

$$\mathbf{poly}_2 (\lambda(i, j).i < j) f g$$

applied to a 8×8 matrix returns the matrix given in Figure 10 where the function f has been applied on white dots and g on black ones. After a row block distribution on 4 processors, each processor must apply \mathbf{poly}_2 on the elements contained in their local memory. The local function f must be applied to the elements (i, j) such that $2 * ip + i < j$.

In order to give the idea of the transformation for the communication skeleton `brdcastuc`, let us consider a specific distribution: `dei (blocki 2) seqi` i.e. the inverse of a row block distribution on 2 processors. The transformed expression consists in a local access to the matrix (preparing the communication), a communication `pbrdcast` and a `brdcastuc` executed by each processor. For example, let b be the size of blocks, the expression `brdcastuc o dei (blocki 2) seqi` is transformed into:

```

dei (blocki 2) seqi
  ◦ pimap  $\lambda(ip, (buf, (e, v))).(\mathbf{brdcast}_{uc}(0, \mathbf{update}(0, buf, v)))$ 
  ◦ pbrdcast ◦ pimap  $\lambda(ip, (e, v)).(\mathbf{div}(e, b), \mathbf{lookup}(\mathbf{mod}(e, b), v, (e, v)))$ 

```

The first function builds a triple containing the number of the broadcasting processor, a value $\mathbf{lookup}(\mathbf{mod}(e, b), x)$ (on the broadcasting processor it will be the value to broadcast), and the local memory. Then, the value is broadcast ($\mathbf{pbrdcast}$), bound to buf , integrated in the local matrix (\mathbf{update}) to be copied locally to all the vector elements ($\mathbf{brdcast}_{uc}$).

```

 $\mathcal{FUS}[(\mathbf{pimap} F_1, \mathbf{pimap} F_2)]$ 
  =  $\mathbf{pimap}(\mathbf{ftuple}_2 F_1 F_2 \circ (\lambda(ip, (x, y)).(F_1(ip, x), F_2(ip, y))))$ 
 $\mathcal{FUS}[(F_1 \circ F_2, \mathbf{pimap}(\lambda(ip, x).x))]$ 
  =  $\mathcal{FUS}[(F_1, \mathbf{pimap}(\lambda(ip, x).x))] \circ \mathcal{FUS}[(F_2, \mathbf{pimap}(\lambda(ip, x).x))]$ 
 $\mathcal{FUS}[(\mathbf{pbrdcast}, \mathbf{pimap}(\lambda(ip, x).x))]$ 
  =  $\mathbf{pimap}(\lambda(ip, (a, (c, d))).((a, c), d)) \circ \mathbf{pbrdcast}$ 
  ◦  $\mathbf{pimap}(\lambda(ip, ((a, b, c), d)).(a, b, (c, d)))$ 
  ...
 $\mathcal{FUS}[(F_1, F_2)]$ 
  =  $\mathcal{FUS}[(F_1, \mathbf{pimap}(\lambda(ip, x).x))] \circ \mathcal{FUS}[(\mathbf{pimap}(\lambda(ip, x).x), F_2)]$ 

```

Figure 11: Transformation \mathcal{FUS} (extract).

The distribution of the function pair \mathbf{ftuple}_2 consists in propagating the distributions associated with each function and merging the resulting SPMD functions in a single parallel function.

Merging is performed by the auxiliary transformation \mathcal{FUS} (Figure 11). When each function of the pair is a \mathbf{pimap} then the pair is rewritten into a single \mathbf{pimap} function. Otherwise, we have to sequentialize the two functions (last rule of Figure 11). Each function is associated with the identity function for the vector of processors (i.e. $\mathbf{pimap}(\lambda(ip, x).x)$) and composed together. The fusion of a communication function and the identity requires to restructure the local data of the processors before and after the communication operation (to thread the unused element of the pair argument). For example, in the case of a $\mathbf{pbrdcast}$, the number of the broadcasting processor and the value to broadcast must be in the first and second components of the local processor memory. The reorganization is carried out by $\mathbf{pimap}(\lambda(ip, (a, (c, d))).((a, c), d))$ and $\mathbf{pimap}(\lambda(ip, ((a, b, c), d)).(a, b, (c, d)))$.

Example 8 According a row block distribution $\mathbf{de}(\mathbf{block} p) \mathbf{seq}$, our example is transformed into the SPMD function:

```

pimap  $\lambda(ip, (x, y)).(\text{poly}_2 (\lambda(i, j).ip * b + i < j) + - x, y)$ 
  ◦ pimap  $\lambda(ip, (buf, (x, y))).\text{brdcast}_{uc}(0, \text{update} (0, buf, x), y)$ 
  ◦ pbrdcast ◦ pimap  $\lambda(ip, (x, y)).(0, \text{lookup} (0, x), (x, y))$ 
  ◦ pimap  $\lambda(ip, x).(\text{copy } x, x)$ 
    
```

According a column block distribution (**de seq** (**block** p)), we would get the following function:

```

pimap  $\lambda(ip, (x, y)).(\text{poly}_2 (\lambda(i, j).i < ip * b + j) + - x, y)$ 
  ◦ pimap  $\lambda(ip, x).(\text{brdcast}_{uc}(0, (\text{copy } x)), x)$ 
    
```

7.3 Optimizing Transformations

The transformations described in this section aim at simplifying and optimizing local (sequential) and parallel computations. This step is described as a set of local program transformations. For example, the following rules are applied:

- merging of **pimaps**

$$\mathbf{pimap} F \circ \mathbf{pimap} G = \mathbf{pimap} (\lambda(ip, v).F(ip, G(ip, v)))$$

- merging of **rotates**

$$\mathbf{rotate}_{uc} \circ \mathbf{ftuple}_2 f \mathbf{rotate}_{uc} = \mathbf{rotate}_{uc} \circ \lambda(e_1, (e_2, v)).(f e_1 + e_2, v)$$

In the rest of this section, we concentrate on an optimization that removes vector allocations (**copy** or **makearray_{uc}**) that have been made useless by the distribution. The impact of this optimization on performances can be considerable.

The local optimization rules are described as rewrite rules in Figure 12. In order to preserve the update-in-place property (Section 5), these local transformations are applied only if they do not violate a criterion on access sequences. This condition remains implicit in the rules.

There are two classes of rules: transformations propagating the vector allocation to the left of the expression (**P.n** rules) and transformations performing the elimination (**E.n** rules). Elimination becomes possible when an allocation has been shifted next to its deallocation.

Propagation transformations amount to delaying vector allocations. The rule [P.1] transforms function applications whose argument contains an allocation.

For example, the expression

$$\mathbf{dealloc} (\mathbf{brdcast}_{uc}(1, x), \mathbf{copy } y)$$

is transformed into

$$(\mathbf{dealloc} \circ \mathbf{ftuple}_2 Id \mathbf{copy})(\mathbf{brdcast}_{uc}(1, x), y)$$

$$\begin{aligned}
& \text{Fun}_1(\text{Exp}_1, \dots, \text{Exp}_{i-1}, \text{Fun}_2 \text{Exp}_i, \dots, \text{Exp}_n) \\
& \quad \Leftrightarrow \text{Fun}_1 \circ \mathbf{ftuple}_n \text{Id} \dots \text{Id} \text{Fun}_2 \dots \text{Id} \\
& \quad \quad (\text{Exp}_1, \dots, \text{Exp}_{i-1}, \text{Exp}_i, \dots, \text{Exp}_n) \\
& \quad \quad \mathbf{if copy, makearray}_{\text{uc}} \in \text{Fun}_2 \quad \quad \quad \text{[P.1]} \\
\mathbf{pimap} (\lambda(ip, (x_1, \dots, x_n)).\text{Exp}) \circ \mathbf{pimap} (\mathbf{ftuple}_n \text{Fun}_1 \dots \text{Fun}_n \circ \text{Fun}) \\
& \quad \Leftrightarrow \mathbf{pimap} (\lambda(ip, (x_1, \dots, x_n)).\text{Exp}[x_i \mapsto \text{Fun}_i x_i]) \\
& \quad \quad \circ \mathbf{pimap} \mathbf{ftuple}_n \text{Fun}_1 \dots \text{Id} \dots \text{Fun}_n \circ \text{Fun} \\
& \quad \quad \mathbf{if Fun}_i \text{ is closed and } \mathbf{copy, makearray}_{\text{uc}} \in \text{Fun}_i \quad \quad \quad \text{[P.2]} \\
\mathbf{brdcast}_{\text{uc}} \circ \mathbf{ftuple}_2 \text{Fun } \mathbf{copy} \\
& \quad \Leftrightarrow \lambda(e, v).\mathbf{makearray}_{\text{uc}}(n, \mathbf{lookup} (\text{Fun } e, v)) \\
& \quad \quad \text{where } \mathbf{copy} : \text{Vect } n \ \alpha \rightarrow \text{Vect } n \ \alpha \quad \quad \quad \text{[P.3]} \\
\mathbf{dealloc} \circ \mathbf{ftuple}_2 \mathbf{copy} \text{Fun} \\
& \quad \Leftrightarrow \text{Fun} \circ \mathbf{extract}^{(x,y),y} \quad \quad \quad \text{[E.1]}
\end{aligned}$$

Figure 12: Copy Elimination (extract)

The rule [P.2] applies to a composition of **pimap**s and propagates the vector allocation occurring in the first **pimap** into the second one.

The rule [P.3] indicates that copying a vector and broadcasting its i th value is similar to fetching and duplicating the i th value. The **copy** is transformed into a **makearray_{uc}** and the allocation is shifted to the left.

The rule [E.1] applies to the case where an allocated vector is immediately freed. It eliminates both the vector allocation and its deallocation.

Example 9 *The distributed version of our example according a column block distribution*

$$\begin{aligned}
& \mathbf{pimap} \lambda(ip, (x, y)).(\mathbf{poly}_2 (\lambda(i, j).i < ip * b + j) + - x, y) \\
& \quad \circ \mathbf{pimap} \lambda(ip, x).(\mathbf{brdcast}_{\text{uc}}(0, \mathbf{copy} x), x)
\end{aligned}$$

is transformed into

$$\begin{aligned}
& \mathbf{pimap} \lambda(ip, (x, y)).(\mathbf{poly}_2 (\lambda(i, j).i < ip * b + j) + - x, y) \\
& \quad \circ \mathbf{pimap} \lambda(ip, x).(\mathbf{makearray}_{\text{uc}}(n, \mathbf{lookup} (0, x), x)
\end{aligned}$$

using rules [P.1] and [P.3]. The copy of the complete matrix is avoided and replaced by the allocation of a single line.

8 Cost Analysis

This step aims at automatically evaluating the cost of \mathcal{L}_6 programs in order to select the most efficient distribution. The complexity analysis is based on polytope volume computations

([Taw94], [Pug94], [Cla96]), and yields accurate symbolic costs. This approach is made possible by the restrictions of \mathcal{L}_1 and the fixed set of data distributions considered. Together, they guarantee that the cost of all transformed source programs can be expressed as a sum of polytope volumes.

The goal is to find the most efficient \mathcal{L}_6 program among the different distribution choices considered. First, an abstraction function \mathcal{CA} transforms programs into their symbolic parallel cost. Costs are expressed using inequalities, sums, and maxima. Then, standard methods to compute the volume of polytopes are applied to get a symbolic cost in polynomial form. Finally, symbolic costs are simplified and compared using a symbolic math package.

8.1 Cost Abstraction

The abstraction function \mathcal{CA} extracts cost information from parallel programs. The main rules are shown in Figure 13. We use a non-standard notation for indexed sums: we write $\sum_i \{i \leq n\}$ instead of $\sum_{i=1}^n$. This notation is needed because polytopes (introduced by the \mathbf{poly}_n skeletons) may be defined by more than two inequalities.

The abstraction relies heavily on the size information present in types (see Section 4). The cost of (**pimap** Fun) is the maximum of the costs of Fun on each processor. Communication costs are expressed as polynomials whose constants depend on the target computer. For example, the cost of **pbrdcast** involves the parameters α_{transf} and α_{init} which denote respectively the time of one-word transfer between two processors and the message startup time on the parallel computer considered. Further, we use the function `size` which returns a polynomial representing the symbolic size of the argument type. Basic arithmetic operators are also given a machine dependent cost (α_{Op}). One has to set those constants to adapt the analysis for a specific parallel machine. The cost of the mask skeleton \mathbf{poly}_2 is the sum of the cost of the first function for elements belonging to the polytope $\lambda(i, j).Ineq$ and of the cost of the second one for elements belonging to the complementary (written \overline{Ineq}). The complementary polytope is expressed as the difference between the whole (i, j) -domain and the polytope $\lambda(i, j).Ineq$. This is a standard operation of polyhedral libraries that yields a union of polytopes.

The obtained cost expression is then reduced in order to remove all the λ -abstractions. Since we deal with terminating programs, any abstracted expression will have a normal form. Moreover, their normal form boils down to generalized sums (G-sum) and generalized maxima (G-max). A G-sum is of the form $\sum_{i_1, \dots, i_n} \{Ineq\} Poly$ where $Ineq$ are inequalities made of affine expressions of loop indexes and vector sizes and $Poly$ is a polynomial whose variables are vector sizes or the processor number. A G-max is of the form $\max_{i=0}^k E_i$ and denotes the maximum among the expressions E_0, \dots, E_k .

Let us 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 the mask skeletons (which limit conditional application to polytope domains), the communication and the computation skeletons all have a complexity which depends polynomially on vector

$$\begin{aligned}
\mathcal{CA}[\text{Fun}_1 \circ \text{Fun}_2] &= (\lambda x. \mathcal{CA}[\text{Fun}_1](\text{Fun}_2 x) + \mathcal{CA}[\text{Fun}_2] x) \\
\mathcal{CA}[\mathbf{pimap} \text{ Fun}] &= (\lambda x. \max_{i_p=0}^{p-1} \mathcal{CA}[\text{Fun}](i_p, x!i_p)) \\
\text{where } \mathbf{pimap} \text{ Fun} : \text{Vectproc } p \alpha &\rightarrow \text{Vectproc } p \beta \\
\mathcal{CA}[\mathbf{pbrdcast}] &= (\lambda x. (\alpha_{\text{transf}} * \text{size } \beta + \alpha_{\text{init}}) * p) \\
\text{where } \mathbf{pbrdcast} : \text{Vectproc } p (\alpha, \beta, \gamma) &\rightarrow \text{Vectproc } p (\beta, \gamma) \\
\mathcal{CA}[\mathbf{Op}] &= (\lambda x. \alpha_{\text{Op}}) \\
\mathcal{CA}[\mathbf{poly}_2 \lambda(i, j). \text{Ineq Fun}_1 \text{ Fun}_2] &= (\lambda x. \sum_{i,j} \left\{ \begin{array}{l} 0 \leq i < n_1 \\ 0 \leq j < n_2 \\ \text{Ineq} \end{array} \right\} \mathcal{CA}[\text{Fun}_1](x!i) \\
&\quad + \sum_{\text{Ineq}' \in \text{Ineq}} \sum_i \left\{ \begin{array}{l} 0 \leq i < n_1 \\ 0 \leq j < n_2 \\ \text{Ineq}' \end{array} \right\} \mathcal{CA}[\text{Fun}_2](x!i)) \\
\text{where } \mathbf{poly}_2 \lambda(i, j). \text{Ineq Fun}_1 \text{ Fun}_2 : \text{Vect } n_1 (\text{Vect } n_2 \alpha) &\rightarrow \text{Vect } n_1 (\text{Vect } n_2 \beta)
\end{aligned}$$

Figure 13: Cost Abstraction Function \mathcal{CA} (extract).

sizes. Their costs can be described as nested sums. Another important restriction is that expressions involving iteration indexes (mask skeletons, communication skeletons and iterfor bounds) are affine. This restriction, expressed by the nonterminal $\text{Lin}F_1$ in the definition of \mathcal{L}_1 , along with the standard distributions considered which keep vector accesses affine, ensures that the inequalities of G-sums are affine.

Example 10 *Our example after a column block distribution is*

$$\mathbf{pimap} \lambda(ip, (x, y)). (\mathbf{poly}_2 (\lambda(i, j). i < ip * b + j) + - x, y)$$

Its cost is expressed as

$$\mathcal{C}_{col} \equiv \max_{i_p=0}^{p-1} \left(\sum_{i,j} \left\{ \begin{array}{l} 0 \leq j < b \wedge 0 \leq i < n \\ i < i_p * b + j \end{array} \right\} \alpha_+ + \sum_{i,j} \left\{ \begin{array}{l} 0 \leq j < b \wedge 0 \leq i < n \\ i_p * b + j \leq i \end{array} \right\} \alpha_- \right)$$

8.2 Symbolic Cost Analysis

If the vector sizes and the number of processors are known at compile time, the cost comparison is done directly by evaluating the value of the abstracted cost expression. An enumeration of the polytope points makes it possible to obtain the accurate numerical value of the execution time on the target machine. The best distribution is that whose cost value is the smallest.

For unknown vector sizes, the cost computation consists in a symbolic evaluation which yields a polynomial expression. This computation is decomposed into two transformations:

the first one evaluates polytope volumes symbolically by reusing existing techniques and the second one removes G-maxs by calculating the maximum value of the polynomial.

Parametrized Polytope Volume Computation

The method of Clauss [Cla96] makes it possible to obtain an accurate computation of the polytopes volumes. The computation principle is based on the fact that such a volume is equal to an extension of polynomials (Ehrhart's pseudo-polynomials), such that the coefficients of its monomials vary according to the remainder of integer division of the symbolic sizes by a constant.

However, this method cannot be simply extended when the processor number is unknown at compile time. In this case, the coefficients of affine expressions in the inequalities can contain an unknown. The volume of this kind of polytope cannot be represented by a pseudo-polynomial.

When we want to keep the number of processors as a parameter, the technique described in [Taw94] can be used. It consists in cutting out the polytope by breaking up the inequalities into several subsets such as each of them contains only two inequalities for each variable and that the lower limit is lower than the upper limit (to rule out null polytopes). After this step, traditional formulae of symbolic summation can be applied and polytope volumes are expressed by polynomials whose variables are symbolic sizes. Integer divisions are approximated by their real division minus a constant representing the average difference between the two operations (e.g. $\lfloor \frac{i}{2} \rfloor$ is approximated by $\frac{i}{2} - \frac{1}{4}$). The evaluated cost is no longer absolutely accurate but the introduced approximations appear to be negligible in practice.

Example 11 Using Tawbi's technique, C_{col} is simplified into

$$\max_{i_p=0}^{p-1} \left(b^2 (\alpha_+ - \alpha_-) i_p + b^2 \left(\frac{\alpha_+}{2} + \frac{\alpha_+}{2} + (p-1) \alpha_- \right) \right)$$

G-max removal

Before comparing symbolic costs, G-maxs occurring in cost expressions must be removed. This can be done by computing the maximum value taken by the polynomial on the symbolic interval of the G-max. This calculation is based on the fact that a polynomial defined over a finite interval $[0..n]$ reaches its maximum value in 0, n , or one of the zeros of its derivative. There can be several maximums depending on conditions on symbolic sizes. A G-max can be rewritten as a collection of polynomials, each polynomial being defined on an interval (i.e. the conditions on sizes) where a zero of the derivative maximizes the original expression. This technique removes G-maxs without any approximations. It may however creates complex symbolic expressions for high-degree polynomials. In such cases, simpler, approximated solutions exist (e.g. considering only the bounds of the interval).

Example 12 With the hypothesis $\alpha_+ \equiv \alpha_-$, the former expression of C_{col} is simplified into

$$\max_{i_p=0}^{p-1} \left(b^2 (\alpha_+ - \alpha_-) i_p + b^2 \left(\frac{\alpha_+}{2} + \frac{\alpha_-}{2} + (p-1) \alpha_- \right) \right) = b^2 p \alpha_+$$

8.3 Symbolic Cost Comparison

The last task is to compare the symbolic costs obtained for different distribution choices. It amounts to computing the symbolic intervals where the difference of two costs (i.e. polynomials) is positive or negative. Symbolic math packages such as Maple [CGG⁺92] can be used for solving this problem. It may be the case that a distribution is definitely better than another and Maple will determine it. But in general, it will depend on sizes and the number of processors. In such cases, Maple can be seen as a simplifier that will produce symbolic conditions (e.g. $Cost_1 > Cost_2$ iff $n > p$). The programmer may have to indicate if these conditions are satisfied or not. Another (automatic) solution is to use these conditions as run-time tests which choose between several versions of the program.

For our example, the column block distribution is always better than the row block one because local computations are identical in both cases but the column distribution does not entail any communication between processors.

9 Translation and Experiments

9.1 Translation

The SPMD programs produced by the compilation chain are translated in a sequential language with calls to a communication library. In the current implementation, we use the C language and the MPI library [CGH94]. These are *de facto* standards which guarantee portability across many different parallel machines.

The single program executed by all the processors is made of the local functions (argument of **pimaps**), the loops (**piterfor**), and calls to MPI. The composition (**o**) is simply translated, locally, by the sequence (**;**); it does not represent a synchronization barrier. The only synchronization between processors is introduced through parallel communications (e.g. **pbroadcast**).

The transformation \mathcal{TRA} of local functions is a straightforward translation into C. For example, the expression $e1 + e2$ will be translated in

```
{ int tmp1;
  <evaluation of e1>; tmp1 = res;
  <evaluation of e2>;
  res = tmp1+res}
```

Since functions are strict and, in essence, first order, they are implemented by C functions. Sequential skeletons such as **brdcast_{uc}**, **copy**, etc. are translated into nested loops. Memory management is implemented using the C library functions **malloc** and **free**. The

translation of communications amounts to calling the corresponding communication function of MPI (recall that we have restricted ourselves to communication skeletons having a direct counterpart in the functions implemented in standard libraries such as MPI).

In all cases, symbolic sizes present in types are needed by the code generation. For example, the vector size is required by the translation of the sequential communication skeleton `brdcastuc` (to set the corresponding loop bound), and by the translation of the parallel communication `pbrdcast` (to set the size of the message).

9.2 Prototype

The compiler has 9 stages. A *parser* produces an abstract syntactic tree starting from the initial program. The *size analysis* is similar to a type inference. It uses the library of polyhedric computations PolyLib [Wil93] to check the coherence of constraints and to determine the overall size constraints. The *update in place transformation* first approximates sharing and the access sequence of the execution. It then checks the property \mathcal{UP} for the abstract access sequences and may insert automatically vector copies and deallocations. The transformations *making communications explicit* and *distribution* are simple syntactic transformations. The *symbolic cost analysis* is relatively complex to implement. Indeed, it is necessary to make inter-operate the cost computation with PolyLib and Maple. The *optimization* and the *translation* towards the target language (C+MPI) are also syntactic transformations which can be implemented directly. Lastly, the *production of object code* is carried out by the C compiler available on the target machine and by linking with the MPI library.

The compilation stages made up of syntactic transformations are easily written in HASKELL. The static analyses are more complex to implement. For our experiments, some compilation steps, such as the destructive update step and most of the symbolic cost computation, were done manually.

9.3 Experiments

We have performed experiments on an Intel Paragon XP/S and a CRAY T3E with a handful of standard linear algebra programs (LU, Cholesky factorization, Jacobi iteration, ...). The experiments have two objectives. First, we study the adequacy between the theoretical execution times produced by the cost analysis and those measured in practice. Then, we compare the performances of \mathcal{L}_1 programs with other implementations (standard sequential C code, HPPF, the skeleton language NESL, and the linear algebra library ScaLAPACK).

Figure 14 gives the execution times, both measured and theoretical (i.e. statically evaluated), for LU decomposition on a CRAY T3E with a cyclic line distribution and a block line distribution. The number of processors varies between 1 and 16 and we give times for two matrix sizes (1024x1024 and 2048x2048). The differences between the theoretical cost and the measured cost are less than 6 %. The execution times on the Paragon with the same example show the same differences with the theoretical cost. Our experiments suggest that

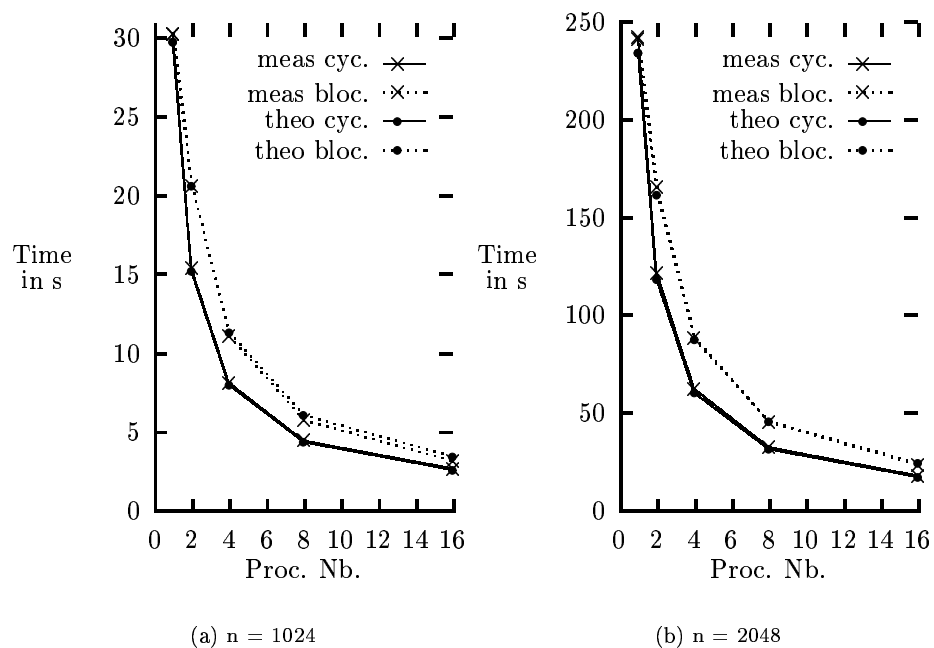


Figure 14: Theoretical (theo) and Measured (meas) Times for LU Decomposition on a CRAY T3E.

the cost analysis is portable. It is sufficient to change the constants representing the basic operations costs to obtain a fairly faithful estimation of the actual execution times.

Figure 15 gathers the execution times obtained for LU decomposition, Cholesky factorization, the Householder method, the Jacobi iteration, and the n-body problem on the Intel Paragon. For all programs, the distribution chosen by the cost analysis proved to be the best one in practice.

| LU (n=256) | | | LU (n=512) | | | Cholesky (n=1024) | | | |
|----------------------|--------|--------|----------------|-------|-------|-------------------|--------|--------|--------|
| Proc. | Skel. | Seq. | HPF | Skel. | Seq. | HPF | Skel. | Seq. | HPF |
| 1 | 2.14 | 1.73 | 2.16 | 14.77 | 13.61 | 15.36 | 67.45 | 53.17 | 65.40 |
| 2 | 1.34 | × | 1.38 | 8.43 | × | 8.67 | 35.54 | × | 35.97 |
| 4 | 0.93 | × | 0.95 | 5.25 | × | 5.41 | 21.35 | × | 20.65 |
| 8 | 0.76 | × | 0.77 | 3.23 | × | 3.38 | 14.81 | × | 13.10 |
| 16 | 0.66 | × | 0.67 | 2.97 | × | 3.06 | 11.55 | × | 9.53 |
| 32 | 0.62 | × | 0.61 | 2.57 | × | 2.67 | 9.91 | × | 7.83 |
| Householder (n=1024) | | | Jacobi (n=512) | | | N body (n=2048) | | | |
| Proc. | Skel. | Seq. | HPF | Skel. | Seq. | HPF | Skel. | Seq. | HPF |
| 1 | 318.16 | 308.17 | 325.44 | 63.42 | 55.10 | 56.20 | 125.15 | 122.76 | 127.86 |
| 2 | 159.04 | × | 164.13 | 32.98 | × | 29.81 | 62.99 | × | 91.77 |
| 4 | 82.12 | × | 84.62 | 17.19 | × | 16.83 | 31.53 | × | 49.04 |
| 8 | 44.59 | × | 46.32 | 7.66 | × | 10.19 | 15.51 | × | 27.39 |
| 16 | 26.68 | × | 27.03 | 3.90 | × | 6.93 | 7.59 | × | 17.35 |
| 32 | 17.98 | × | 18.23 | 1.98 | × | 5.19 | 3.64 | × | 12.68 |

Figure 15: Times (in s) for Skeletons and HPF on Intel Paragon XP/S.

We compared the sequential execution of \mathcal{L}_1 programs with standard (and portable) C versions taken from [PTVF86]. We also compared 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 believe that the important differences for Jacobi and Nbody are due to the inability of the HPF compiler to recognize collective communications in general.

We compared our code with the parallel implementation of NESL, a skeleton-based language [BHS⁺94]. 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).

Finally, we have considered ScaLAPACK, an optimized library of linear algebra programs designed for distributed memory MIMD parallel computers [CD95]. In ScaLAPACK, the user may explicitly indicate the data distribution. So, we indicated the best distribution found

| Proc. | LU (n=512) | | Householder (n=1024) | | Cholesky (n=1024) | |
|-------|------------|-----------|----------------------|-----------|-------------------|-----------|
| | Skel. | ScaLAPACK | Skel. | ScaLAPACK | Skel. | ScaLAPACK |
| 1 | 14.77 | 3.78 | 318.16 | 56.35 | 67.45 | 55.80 |
| 2 | 8.43 | 2.4 | 159.04 | 35.23 | 35.54 | 34.95 |
| 4 | 5.25 | 1.84 | 82.12 | 22.57 | 21.35 | 21.80 |
| 8 | 3.23 | 1.66 | 44.59 | 16.27 | 14.81 | 15.56 |
| 16 | 2.97 | 1.50 | 26.68 | 12.82 | 11.55 | 12.56 |
| 32 | 2.57 | 1.41 | 17.98 | 10.83 | 9.91 | 11.32 |

Figure 16: Times (in s) for Skeletons and ScaLAPACK on Intel Paragon XP/S.

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 6 times slower), the difference decreases as the number of processors increases (typically, 1.8 times slower on 32 processors). We believe that 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 extend \mathcal{L}_1 with new skeletons corresponding to the BLAS operations in order to benefit from these machine specific routines.

Note that ScaLAPACK allows block cyclic distributions with a variable size of blocks which are a more general form of distribution than ours. This enables the programmer to sometimes find a better compromise between communication costs and load balancing by guessing the right block size. This is the case for the Cholesky factorization where the optimal distribution is block cyclic with a size of blocks between one and the vector size divided by the number of processors. Integrating block cyclic distributions (with a variable size of blocks) within our framework is not obvious. First, the combination of such distributions and the poly_n skeletons gives rise to expressions whose cost cannot be expressed as polytopes. Second, it is clear that an exhaustive analysis of all possible distributions would become unrealistic in this case. Extending the cost analysis to this kind of distributions would presumably require approximations and interactions with the user.

10 Conclusions

We have presented the compilation of a skeleton-based language for parallel computers. Our compilation process makes use of a variety of techniques: typing, static analyses, program transformations, polytope volume computation. Working by program transformations in a unified framework simplifies the correctness proofs 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. But, the most important characteristic of our approach is the source language restrictions. We now review and justify them.

10.1 Source language restrictions

Most \mathcal{L}_1 restrictions were guided by our need for an accurate, symbolic cost analysis. Relaxing one of the following restrictions would not be possible without changing drastically the approach.

- *Restricted recursion.* Disallowing general recursion is requisite for an accurate cost analysis. The existence of a fixpoint operator in the language makes complexity analysis undecidable. Skeletons are a way to tame recursion since the data and control flow are known *a priori*.
- *Restricted vector manipulations.* The symbolic cost analysis requires the symbolic size of all vectors. This information is inferred by the size analysis. In order to do so, the size of a vector should not depend on scalar values but only on constants or parameter sizes. A program observing this condition is called “shapely” [JS97]. For example, a skeleton *filter p* producing a vector made only of the elements of its vector argument satisfying the predicate *p* cannot be allowed. It would produce vectors whose size is unpredictable.
- *Restricted conditionals.* One cannot associate an accurate cost to conditionals whose test depends on scalar values. Only the maximum, the minimum, or, with probabilistic information, the average complexity could be evaluated in this case. Moreover, to reuse the powerful tools based on polytopes, costs must be defined by affine (in)equations depending on symbolic sizes and iterator indexes. In order to ensure the affinity of cost expressions, conditionals are constrained to be mask skeletons whose condition characterizes a convex polytope. Typing enforces that the condition is an affine expression of symbolic sizes and iterator indexes only. If this condition were expressed as a polynomial, its cost could not be expressed as the volume of a polytope.
- *Restricted communication skeletons.* The communications involved by \mathcal{L}_1 programs must be statically predictable. This is ensured by enforcing the arguments of the communication skeletons to be of type *Index* or *Size* in order to infer their symbolic values at compile time.

The following restrictions were chosen because they entailed simplifications or a more efficient implementation. They could be relaxed to a certain extent.

- The communication skeletons considered are standard collective communication primitives (broadcast, translation...) which are either hard-wired or optimized on many parallel machines. However, new communication skeletons could be taken into account. Similarly, the collections of computation and reorganization skeletons could be extended.

- We have considered only nested vectors, not multi-dimensional arrays. Nested vectors are general but make a distinction between dimensions. For example, some operations on columns are less easily expressed than the same operations on lines. It would be possible to include new families of skeletons acting on 2D or 3D matrices.
- User defined functions are first order. We feel that higher order functions (and the use of closures) would make the analyses and the implementation much more complex. However, a simple solution to relax this restriction would be to use a preliminary transformation removing higher order functions (e.g. [Chi90]).

10.2 Related work

The existing specialized languages for data parallelism are generally based on restricted forms of recursion. These restrictions are either syntactic constraints on the form of recursive calls or the fixpoint operator is replaced by a collection of skeletons.

The ALPHA language [Wil94a] is a first order, strongly typed, functional language. An ALPHA program is a system of recursive functions where the arguments within recursive calls are restricted to be affine expressions of the function parameters. These restrictions make it possible to define precise static analyses and lead to efficient implementations on MIMD machines [QRW95]. ALPHA was initially introduced to express systolic algorithms and may seem too restricted to express less regular programs. A possibility is to relax the affinity restriction, as in the systolic language CRYSTAL [CCL91], but the static analyses lose their precision.

The FORTRAN community has studied automatic data distribution through parallel cost estimation [GB92, CGST93]. 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. [Taw94], [Pug94] and [Cla96] focus on a subset of FORTRAN: loop bound and array indexes are affine expressions of the loop variables. This restriction allows them to compute a precise symbolic computation cost based on polytopes. 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 [Fea94]. These approaches estimate real costs too roughly to ensure that a good distribution is chosen.

The skeletons of [Col88, DFH⁺93] can be seen as hard-wired parallel schemes (e.g. divide-and-conquer, pipe, etc.). Each skeleton comes with a fixed, optimal, distribution and implementation; this may entail a (costly) redistribution before each skeleton. This approach cannot exploit the nested parallelism expressed by the combination of skeletons (e.g. only the parallelism of the top-level **map** would be taken into account in the expression **map (map (+1))**). The skeletons of [DGTJ95] include classic higher-order functions (**map**, **fold**, **scan**) and coordination and distribution skeletons. In the same vein, Südholt [Süd97] introduces high-level distribution skeletons much more general than HPF distributions. In these approaches, data distribution is explicit and chosen by the programmer. Finally, the

skeletons of [BHS⁺94, SC93] can be nested and facilitate the expression of algorithms having several levels of parallelism.

The skeleton community has produced several work on size analysis. Jay [JCSS97] defines a shape analysis for a specialized functional language with vectors (the Vec language). Functions are restricted to be “shapely”, i.e. the shape (size) of the results should depend only on that of the inputs and not of the value of the data. This restriction makes it possible to evaluate the size of vectors. However, the numerical size of input vectors must be known at compile time. The analysis is not symbolic and may be expensive since it amounts to evaluating an abstract program of the same complexity that the source program. Nitsche [Nit98] aims at detecting “shapely” expressions in a standard functional language. The “shapeliness” property is undecidable in general and that the analysis only finds a subset of “shapely” expressions. Herrmann [HL98] defines a size analysis for functional language with nested lists. The analysis is symbolic but incomplete (recursive functions makes the problem undecidable).

The implementation of skeleton languages are based on cost analyses. [GL98] defines precise communication costs for combination of scan and fold skeletons on several parallel topologies (hypercube, mesh, ...). [SC93], [JCSS97] and [Ran96] define cost analyses for skeleton-based languages. Their skeletons are less restricted than ours leading to an 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. [Bra93] uses cost estimations based on profiling to choose the distribution for each skeleton. Such experimental approaches do not ensure good and portable parallel performances for different machines, number of processors, or sizes of inputs.

Most implementations use cost information to apply local, cost-reducing, transformations [DFH⁺93] or to choose the best distribution for each skeleton [Pel93]. In both cases, implementation decisions are local and no arbitration of tradeoffs is possible.

In our approach, we start from a high level language (\mathcal{L}_1) where skeletons can be freely nested and obtain a skeleton language (\mathcal{L}_6) with explicit distribution, communication, and allocation similar to the source language of [DGTJ95]. Contrary to local optimization approaches, we consider a global distribution and cost analysis, and we are able to select the best implementation (among a restricted set of choices).

10.3 Future Work

The preliminary results obtained by our prototype are promising but more experiments are necessary to assess both the expressiveness of the language and the efficiency of the compilation. The expressivity may be evaluated by encoding a significant set of examples requiring high performances. The algorithms introduced into the Cowinchan set [Wil94b] seem a good starting point since they were conceived to test the expressivity and the elegance of parallel programming languages. Experiments are also necessary to evaluate more thoroughly the precision of the cost analysis, in particular, the costs associated with the communication primitives. Another field of experimentation is the evaluation of our compilation process for SIMD machines. For such parallel machines, the cost of synchronization is negligible, but the

load balancing of computations is of primary importance to obtain good performances. Our cost analysis seems reusable in this context. Experiments are needed to support this claim and may also suggest new optimizations specific to SIMD machines.

Another research direction is to study dynamic redistributions chosen at compile-time. Some parallel algorithms (e.g. Alternative Direction Implicit Integration [GlO93]) 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 to this problem is to consider (high-level) interactions with the user.

The language \mathcal{L}_1 introduces only data parallelism. However, certain algorithms are more easily expressed in the form of control parallelism such as “divide and conquer” algorithms. The extension of the language \mathcal{L}_1 with such skeletons would be useful to increase the expressivity of the language. [HL97] describes the compilation of “divide and conquer” skeletons into nested sequential and parallel loops. This can be seen as transforming control parallelism into data parallelism. So, a solution to accommodate control parallelism in our approach would be to characterize a class of control parallelism skeletons that can be transformed into \mathcal{L}_1 .

More generally, skeletons appear to be an interesting technique to the design of domain specific languages (DSLs). They make it possible to describe high-level languages enjoying important properties without preventing further extensions (by adding new skeletons). We believe that this approach to DSLs deserves more consideration.

Acknowledgements: Thanks to Rémi Douence and Mario Südholt for their useful comments.

References

- [BHS⁺94] G. E. Blelloch, J. C. Hardwick, J. Sipelstein, M. Zaghera, and S. Chatterjee. Implementation of a portable nested data-parallel language. *Journal of Parallel and Distributed Computing*, 21(1):4–14, April 1994.
- [Bra93] T. Bratvold. A Skeleton-Based Parallelising Compiler for ML. In *Proc. of the International Workshop on Parallel Implementation of Functional Languages*, pages 23–33, 1993.
- [CCL91] M. Chen, Y. Choo, and J. Li. Crystal: Theory and Pragmatics of Generating Efficient Parallel Code. In Boleslaw K. Szymanski, editor, *Parallel Functional Languages and Compilers*, chapter 7, pages 255–308. ACM Press, 1991.
- [CD95] J. Choi and J. J. Dongarra. Scalable linear algebra software libraries for distributed memory concurrent computers. In *Proc. of the IEEE Workshop on Future Trends of Distributed Computing Systems*, pages 170–177, 1995.

- [CGG⁺92] B. W. Char, K. O. Geddes, G. H. Gonnet, B. L. Leong, M. B. Monagan, and S. M. Watt. *Maple V Language Reference Manual*. Springer-Verlag, 1992.
- [CGH94] L. Clarke, I. Glendinning, and R. Hempel. The MPI Message Passing Interface Standard. In *Programming environments for massively parallel distributed systems: working conference of the IFIP*, pages 213–218. Birkhäuser, 1994.
- [CGST93] S. Chatterjee, J. R. Gilbert, R. Schreiber, and S. Teng. Automatic array alignment in data-parallel program. In *Proc. of the ACM Symposium on Principles of Programming Languages*, pages 16–28, 1993.
- [Chi90] W. N. Chin. *Automatic Methods for Program Transformation*. PhD thesis, Imperial College, 1990.
- [Cla96] P. Clauss. Counting solutions to linear and nonlinear constraints through Ehrhart polynomials: Applications to analyze and transform scientific programs. In *Proc. of the ACM International Conference on Supercomputing*, pages 278–285, 1996.
- [Col88] M. Cole. A skeletal approach to the exploitation of parallelism. In *Proc. of the CONPAR*, pages 667–675. Cambridge University Press, 1988.
- [DFH⁺93] 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 *Proc. of the PARLE*, pages 146–160. LNCS 694, 1993.
- [DGTJ95] J. Darlington, Y. K Guo, H. W. To, and Y. Jing. Skeletons for structured parallel composition. In *Proc. of ACM Symposium on Principle and Practice of Parallel Programming*, pages 19–28, 1995.
- [Fea94] P. Feautrier. Toward automatic distribution. *Parallel Processing Letters*, 4(3):233–244, 1994.
- [FJL⁺88] G. Fox, M. Johnson, G. Lyzenga, S. Otto, J. Salmon, and D. Walker. *Solving Problems on Concurrent Processors*. Prentice-Hall, 1988.
- [GB92] 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.
- [GH90] J. C. Guzmán and P. Hudak. Single-threaded polymorphic lambda calculus. In *Proc. of the IEEE Symposium on Logic in Computer Science*, pages 333–343. IEEE Computer Society Press, 1990.
- [GL98] S. Gorlatch and C. Lengauer. (De)Composition rules for parallel scan and reduction. In *Proc. of the IEEE International Working Conference on Massively Parallel Programming Models*, 1998.

- [GlO93] G. H. Golub and J. M. Ortega. *Scientific Computing: An Introduction with Parallel Computing*. Academic Press, Boston, MA, USA, 1993.
- [HJW92] P. Hudak, S. P. Jones, and P. Wadler. Report on the programming language haskell. *SIGPLAN Notices*, 27(5), 1992.
- [HL97] C. Herrmann and C. Lengauer. Transformation of divide conquer to nested parallel loops. In *Programming Languages: Implementation, Logics, and Programs*, LNCS 1292, pages 95–109, 1997.
- [HL98] C. A. Herrmann and C. Lengauer. Size inference of nested lists in functional programs. In *Proc. of the International Workshop on Implementation of Functional Languages*, pages 347–364. University College, London, 1998.
- [HPWo99] P. Hudak, S. Peyton Jones, P. Wadler, and other. Standard Libraries for the Haskell 98 Programming Language. Research report, Yale university, 1999.
- [JCSS97] 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.
- [Joh85] T. Johnsson. Lambda lifting: Transforming programs to recursive equations. In *Functional Programming Languages and Computer Architecture*, pages 190–203. LNCS 201, 1985.
- [JS97] C. B. Jay and M. Sekanina. Shape checking of array programs. In *Computing: the Australasian Theory Seminar*, Australian Computer Science Communications, pages 113–121, 1997.
- [KS86] U. Kastens and M. Schmidt. Lifetime analysis for procedure parameters. In *European Symposium on Programming*, pages 53–69. LNCS 213, 1986.
- [Len93] C. Lengauer. Loop parallelization in the polytope model. In *International Conference on Concurrency Theory*, LNCS 715, pages 398–416, Hildesheim, Germany, 1993.
- [Mal98a] J. Mallet. *Compilation d'un langage spécialisé pour machine massivement parallèle*. Doctorat d'université, Dec. 1998, Rennes I, Ifsic, 1998.
- [Mal98b] J. Mallet. Symbolic cost analysis and automatic data distribution for a skeleton-based language. In *Euro-Par'98 Parallel Processing*, LNCS 1470, pages 688–697, 1998.
- [Mit91] J.C. Mitchell. Type inference with simple sub-types. *Journal of Functional Programming*, 1(3):245–286, 1991.

- [Nit98] T. Nitsche. Shape analysis of functional programs with algebraic data types. In *Proc. of the Workshop on Constructive Methods for Parallel Programming*, pages 16–30, 1998.
- [Pel93] S. Pelagatti. *A Methodology for the Development and the Support of Massively Parallel Programs*. PhD thesis, Pise University, 1993.
- [PTVF86] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and Brian P. Flannery. *Numerical Recipes in FORTRAN The Art of Scientific Computing*. Cambridge University Press, 1986.
- [Pug94] W. Pugh. Counting solutions to presburger formulas: How and why. In *Proc. of the Conference on Programming Language Design and Implementation*, pages 121–134, 1994.
- [QRW95] P. Quinton, S. Rajopadhye, and D. Wilde. On deriving data parallel code from a functional program. In *Proc. of the International Symposium on Parallel Processing*, pages 766–773. IEEE Computer Society Press, 1995.
- [Ran96] R. Rangaswami. *A Cost Analysis for a Higher-order Parallel Programming Model*. PhD thesis, Edinburgh University, 1996.
- [SC93] D. B. Skillicorn and W. Cai. A cost calculus for parallel functional programming. Technical report, Queen’s University, 1993.
- [Sch85] D. A. Schmidt. Detecting global variables in denotational definitions. *ACM TOPLAS*, 7(2):299–310, 1985.
- [Ses89] P. Sestoft. Replacing functional parameters by global variables. In *Proc. of the Conference on Functional Programming Languages and Computer Architecture*, pages 39–53. ACM, 1989.
- [Süd97] M. Südholt. *The Transformational Derivation of Parallel Programs using Data-Distribution Algebras and Skeletons*. PhD thesis, Technische Universität Berlin, 1997.
- [Taw94] N. Tawbi. Estimation of nested loops execution time by integer arithmetic in convex polyhedra. In *Proc. of International Symposium on Parallel Processing*, pages 217–223, 1994.
- [Tur79] D. A. Turner. A new implementation technique for applicative languages. *Software – Practice and Experience*, 9:31–49, 1979.
- [Wad90] P. Wadler. Linear types can change the world! In *Programming Concepts and Methods*, pages 561–581. North Holland, 1990.
- [Wil93] D. Wilde. A library for polyhedral operations. Publication Interne 785, IRISA, 1993.

- [Wil94a] D. Wilde. The ALPHA language. Technical Report 2295, INRIA, 1994.
- [Wil94b] G. Wilson. Assessing the usability of parallel programming systems: The cowichan problems. In *Proc. of the IFIP Working Conference on Programming Environments for Massively Parallel Distributed Systems*, 1994.

A Transformations & Analyses (Addendum)

We describe here additional rules for the size inference and the transformations. Along with the rules shown in the main text, there is at least one rule for each skeleton class. The reader is referred to [Mal98a] for complete descriptions.

A.1 Size Inference

$$\frac{C \vdash \text{Int} \subseteq \alpha \quad \text{Op} \in \{+, -\}}{C, \Gamma \vdash \text{Op} : (\alpha, \alpha) \rightarrow \alpha, \{ \}}_{[\text{OP}]} \quad \frac{C, \Gamma \vdash e : \alpha^{s_1}, C_1 \quad C \vdash \alpha \subseteq \text{Float}}{C, \Gamma \vdash k * e : \alpha^{s_2}, C_1 \cup \{s_2 = k * s_1\}}_{[\text{LINF}_2]}$$

$$\frac{C, \Gamma \vdash e : \text{Index}^{s_1}, C_1 \quad C, \Gamma \vdash f : (\text{Index}^{s_2}, \beta) \rightarrow \gamma, C_2 \quad C_1 \cup C_2 \Rightarrow \{1 \leq s_2 \leq s_1\}}{C, \Gamma \vdash \text{iterfor } e f : \alpha \rightarrow \alpha, C_1 \cup C_2 \cup \{\alpha = \beta = \gamma\}}_{[\text{ITE}]}$$

$$\frac{C, \Gamma \cup \{x_1 : \alpha_1, \dots, x_n : \alpha_n\} \vdash e : \beta, C_1}{C, \Gamma \vdash \lambda(x_1, \dots, x_n).e : (\alpha_1, \dots, \alpha_n) \rightarrow \beta, C_1}_{[\text{ABS}]}$$

$$\overline{C, \Gamma \vdash \text{zip} : (\text{Vect } s_1 \alpha, \text{Vect } s_2 \beta) \rightarrow \text{Vect } s_1 (\alpha, \beta), \{s_1 = s_2\}}_{[\text{ZIP}]}$$

$$\frac{C, \Gamma \vdash e : \text{Size}^s, C_1}{C, \Gamma \vdash \text{makearray } e : \alpha \rightarrow \text{Vect } s \alpha, C_1}_{[\text{MAK}]}$$

$$\frac{C, \Gamma \vdash e : \beta_1, C_1 \quad C, \Gamma \vdash f : (\beta_2, \alpha) \rightarrow \beta_3, C_2}{C, \Gamma \vdash \text{fold } e f : \text{Vect } s \alpha \rightarrow \beta_1, C_1 \cup C_2 \cup \{\beta_1 = \beta_2 = \beta_3\}}_{[\text{FOLD}]}$$

A.2 Distribution

Abstraction Transformation

$$\begin{aligned} \mathcal{ABS}[\lambda \vec{X}. \text{Exp}] \vec{Y} &= \mathcal{ABS}[\text{extract}^{(a,b),a} \circ \mathcal{ABS}[\text{Exp}] \vec{X}] \vec{Y} \\ \mathcal{ABS}[\text{iterfor } \text{Exp } \text{Fun}] \vec{X} &= \text{iterfor}_{\text{uc}} (\mathcal{ABS}[\text{Fun}] \vec{X}) \\ &\quad \circ \text{extract}^{(a,(b,c)),(b,(c,a))} \circ \text{ftuple}_2 \text{ Id } (\mathcal{ABS}[\text{Exp}] \vec{X}) \\ \mathcal{ABS}[\text{Op}] \vec{X} &= \text{ftuple}_2 \text{ Op } \text{Id} \\ \mathcal{ABS}[\text{zip}] \vec{X} &= \text{ftuple}_2 \text{ zip } \text{Id} \\ \mathcal{ABS}[\text{makearray } \text{Exp}] \vec{X} &= \text{ftuple}_2 \text{ makearray}_{\text{uc}} \text{Id} \\ &\quad \circ \text{extract}^{(a,(b,c)),((a,b),c)} \circ \text{ftuple}_2 \text{ Id } (\mathcal{ABS}[\text{Exp}] \vec{X}) \\ \mathcal{ABS}[\text{fold } \text{Exp } \text{Op}] \vec{X} &= \text{ftuple}_2 (\text{fold}_{\text{uc}} \text{Op}) \text{Id} \\ &\quad \circ \text{extract}^{(a,(b,c)),((b,a),c)} \circ \text{ftuple}_2 \text{ Id } (\mathcal{ABS}[\text{Exp}] \vec{X}) \end{aligned}$$

Distribution Transformation

$$\begin{aligned}
& \mathit{DIST}[\mathit{iterfor}_{\text{uc}} F \circ \mathit{dpi} (\mathit{loci} p) \mathit{di}] \\
&= \text{let } \mathit{di} \circ F' = \mathit{DIST}[F \circ \mathit{dpi} (\mathit{loci} p) \mathit{di}] \\
&\quad \text{in } \mathit{di} \circ \mathit{piterfor} F' \quad \text{with } \mathit{di} : \mathit{Vect} p \alpha \rightarrow \alpha \\
& \mathit{DIST}[\mathit{dealloc} \circ \mathit{dpi} \mathit{di}_1 \mathit{di}_2] \\
&= \mathit{di}_2 \circ \mathit{pimap} (\lambda(ip, x). \mathit{dealloc} x) \\
& \mathit{DIST}[\mathit{Op} \circ \mathit{consti} k p] \\
&= \mathit{consti} k p \circ \mathit{pimap} (\lambda(ip, x). \text{if } ip = k \text{ then } \mathit{Op} x \text{ else } x) \\
& \mathit{DIST}[\mathit{zip} \circ \mathit{dpi} \mathit{fdi} \mathit{fdi}] \\
&= \mathit{fdi} \circ \mathit{pimap} (\lambda(ip, x). \mathit{zip} x) \\
& \mathit{DIST}[\mathit{zip} \circ \mathit{dpi} (\mathit{dei} \mathit{fdi} \mathit{di}_1) (\mathit{dei} \mathit{fdi} \mathit{di}_2)] \\
&= \mathit{dei} \mathit{fdi} (\mathit{dpi} \mathit{di}_1 \mathit{di}_2) \circ \mathit{pimap} (\lambda(ip, x). \mathit{zip} x) \\
& \mathit{DIST}[\mathit{fold}_{\text{uc}} \mathit{Op} \circ \mathit{dpi} \mathit{seqi} (\mathit{consti} k p)] \\
&= \mathit{consti} 0 p \circ \mathit{pimap} (\lambda(ip, (e, v)). \text{if } ip = 0 \text{ then } \mathit{fold}_{\text{uc}} \mathit{Op}(e, v) \text{ else } e) \\
&\quad \circ \mathit{ptransfer} \circ \mathit{pimap} (\lambda(ip, (v, e)). (k, 0, e, v)) \\
& \mathit{DIST}[\mathit{fold}_{\text{uc}} \mathit{Op} \circ \mathit{dpi} (\mathit{blocki} p) (\mathit{consti} k p)] \\
&= \mathit{consti} k p \\
&\quad \circ \mathit{pimap} (\lambda(ip, (v_1, v_2)). \text{if } ip = k \text{ then } \mathit{fold}_{\text{uc}} \mathit{Op} (v_1, v_2) \text{ else } 1_{\mathit{Op}}) \\
&\quad \circ \mathit{pgather} \circ \mathit{pimap} (\lambda(ip, (v_1, v_2)). (k, \mathit{fold}_{\text{uc}} \mathit{Op} (v_1, 1_{\mathit{Op}}), v_2)) \\
&\quad \quad \text{if } \mathit{Op} \text{ is commutative} \\
& \mathit{DIST}[\mathit{fold}_{\text{uc}} \mathit{Op} \circ \mathit{dpi} (\mathit{cyci} p) (\mathit{consti} k p)] \\
&= \mathit{consti} k p \\
&\quad \circ \mathit{pimap} (\lambda(ip, (v_1, v_2)). \text{if } ip = k \text{ then } \mathit{fold}_{\text{uc}} \mathit{Op} (v_1, v_2) \text{ else } 1_{\mathit{Op}}) \\
&\quad \circ \mathit{pgather} \circ \mathit{pimap} (\lambda(ip, (v_1, v_2)). (k, \mathit{fold}_{\text{uc}} \mathit{Op} (v_1, 1_{\mathit{Op}}), v_2)) \\
&\quad \quad \text{if } \mathit{Op} \text{ is commutative} \\
& \mathit{DIST}[\mathit{fold}_{\text{uc}} \mathit{Op} \circ \mathit{dpi} (\mathit{cyci} p) (\mathit{consti} k p)] \\
&= \mathit{consti} k p \\
&\quad \circ \mathit{pimap} (\lambda(ip, (v_1, v_2)). \text{if } ip = k \text{ then } \mathit{fold}_{\text{uc}} \mathit{Op} (\mathit{cyci} p v_1, v_2) \text{ else } 1_{\mathit{Op}}) \\
&\quad \circ \mathit{pgather} \circ \mathit{pimap} (\lambda(ip, x). (k, x)) \\
& \mathit{DIST}[\mathit{makearray}_{\text{uc}} \circ \mathit{dpi} (\mathit{consti} k p) (\mathit{consti} k p)] \\
&= \mathit{consti} k p \circ \mathit{pimap} (\lambda(ip, (e, x)). \mathit{makearray}_{\text{uc}}(e, x))
\end{aligned}$$

A.3 Cost Analysis

$$\begin{aligned}
\mathcal{CA}[\lambda(x_1, \dots, x_n). e] &= \lambda(x_1, \dots, x_n). \mathcal{CA}[e] \\
\mathcal{CA}[\mathit{Fun} \mathit{Exp}] &= \mathcal{CA}[\mathit{Fun}] (\mathit{Exp}) + \mathcal{CA}[\mathit{Exp}] \\
\mathcal{CA}[(\mathit{Exp}_1, \mathit{Exp}_2)] &= \mathcal{CA}[\mathit{Exp}_1] + \mathcal{CA}[\mathit{Exp}_2] \\
\mathcal{CA}[x] &= 0 \\
\mathcal{CA}[k] &= 0 \\
\mathcal{CA}[\mathit{iterfor}_{\text{uc}} \mathit{Fun}] &= (\lambda x. \sum_i \{ \overset{1}{\leq} \overset{i}{s} \} \mathcal{CA}[\mathit{Fun}](i, x)) \\
&\quad \text{where } \mathit{iterfor}_{\text{uc}} \mathit{Fun} : (\mathit{Index}^s, \alpha) \rightarrow \alpha \\
\mathcal{CA}[\mathit{copy}] &= (\lambda x. \alpha_{\mathit{copy}} * \text{size } \alpha) \quad \text{where } \mathit{copy} : \alpha \rightarrow \beta \\
\mathcal{CA}[\mathit{dealloc}] &= (\lambda x. \alpha_{\mathit{dea}} * \text{size } \alpha) \quad \text{where } \mathit{dealloc} : (\alpha, \beta) \rightarrow \beta
\end{aligned}$$

$$\begin{aligned}
 \mathcal{CA}[\mathbf{zip}] &= (\lambda x. \alpha_{zip} * n) \quad \text{where } \mathbf{zip} : (Vect\ n\ \alpha, Vect\ n\ \beta) \rightarrow \gamma \\
 \mathcal{CA}[\mathbf{makearray}_{uc}] &= (\lambda x. \alpha_{make} * e * \text{size } \alpha) \quad \text{where } \mathbf{makearray}_{uc} : (\beta, \alpha) \rightarrow Vect\ e\ \alpha \\
 \mathcal{CA}[\mathbf{fold}_{uc}\ Op] &= (\lambda x. \sum_i \{ \begin{smallmatrix} 1 \leq i \\ i \leq n \end{smallmatrix} \} (\mathcal{CA}[Op](x!i) + \alpha_{fol})) \\
 &\quad \text{where } \mathbf{fold}_{uc}\ Op : Vect\ n\ \alpha, \beta \rightarrow \gamma \\
 \mathcal{CA}[\mathbf{brdcast}_{uc}] &= (\lambda x. \alpha_{br} * n * \text{size } \alpha) \quad \text{where } \mathbf{brdcast}_{uc} : (\alpha, Vect\ n\ \alpha) \rightarrow Vect\ n\ \alpha
 \end{aligned}$$

B Examples of Proofs

The compilation process is proven correct by showing for each transformation step three properties (cf. Section 2). Proofs are mainly simple structural inductions. We sketch here the proof for the property 1 for the abstraction \mathcal{ABS} and the property 2 for the distribution transformation \mathcal{DIST} .

B.1 Proof of the Property 1

The source language of the transformation \mathcal{ABS} is \mathcal{L}_3 which is the language \mathcal{L}_1 with two additional functions **copy** and **dealloc** and such that the argument functions of the skeleton \mathbf{poly}_n are closed.

The target language of \mathcal{ABS} is \mathcal{L}_4 whose syntax is:

| | | |
|----------------------|-------|--|
| Prog_4 | $::=$ | $\text{Fun}_4 \text{ VarTuple}_4$ |
| Fun_4 | $::=$ | $\text{Fun}_4 \circ \text{Fun}_4 \mid \mathbf{iterfor}_{\text{uc}} \text{Fun}_4 \mid \text{Op}_4 \mid \mathbf{extract}^{(\text{VarTuple}_4, \text{VarTuple}_4)}$ |
| | | $\mid \mathbf{ftuple}_n \text{Fun}_4 \dots \text{Fun}_4$ |
| | | $\mid \text{CompSkel}_4 \mid \text{ReorgSkel}_4 \mid \text{CommSkel}_4 \mid \text{MaskSkel}_4$ |
| Op_4 | $::=$ | $+ \mid - \mid * \mid \mathbf{div} \mid \mathbf{exp} \mid \mathbf{log} \mid \mathbf{cos} \mid \dots \mid \mathbf{copy} \mid \mathbf{dealloc}$ |
| CompSkel_4 | $::=$ | $\mathbf{map} \text{Fun}_4 \mid \mathbf{fold}_{\text{uc}} \text{Op}_4 \mid \mathbf{scan}_{\text{uc}} \text{Op}_4$ |
| ReorgSkel_4 | $::=$ | $\mathbf{zip} \mid \mathbf{unzip} \mid \mathbf{concat} \mid \mathbf{makearray}_{\text{uc}}$ |
| CommSkel_4 | $::=$ | $\mathbf{brdcast}_{\text{uc}} \mid \mathbf{transfer}_{\text{uc}} \mid \mathbf{rotate}_{\text{uc}} \mid \mathbf{scatter}_{\text{uc}}$ |
| | | $\mid \mathbf{gather}_{\text{uc}} \mid \mathbf{allgather}_{\text{uc}} \mid \mathbf{allbrdcast}_{\text{uc}}$ |
| MaskSkel_4 | $::=$ | $\mathbf{poly}_n \lambda(x_1, \dots, x_n). \text{Ineq}_4 \text{Fun}_4 \text{Fun}_4$ |
| Ineq_4 | $::=$ | $\text{Ineq}_4 \wedge \text{Ineq}_4 \mid \text{LinF}_4 < \text{LinF}_4 \mid \text{LinF}_4 = \text{LinF}_4$ |
| LinF_4 | $::=$ | $\text{LinF}_4 + \text{LinF}_4 \mid \text{LinF}_4 - \text{LinF}_4 \mid k * \text{LinF}_4 \mid x \mid k$ |
| VarTuple_4 | $::=$ | $(\text{VarTuple}_4, \dots, \text{VarTuple}_4) \mid x$ |

The property 1 is expressed as:

$\forall P \in \text{Prog}_3, (\mathcal{ABS}[[P]] \vec{X}) \vec{X} \in \text{Prog}_4$ where \vec{X} are the free variables of P .

The proof of this property boils down to the proof of the corresponding properties on (recursive) non-terminals of \mathcal{L}_3 . That is:

- $\forall E \in \text{Exp}_3, \mathcal{ABS}[[E]] \vec{X} \in \text{Fun}_4$ where \vec{X} contains the free variables of E
- $\forall F \in \text{Fun}_3, \mathcal{ABS}[[F]] \vec{X} \in \text{Fun}_4$ where \vec{X} contains the free variables of F

This proof is done by structural induction. We show only the case of pairs.

Let $E_1, E_2 \in \text{Exp}_3$, we have

$$\begin{aligned} \mathcal{ABS}[[E_1, E_2]] \vec{X} &= \mathbf{extract}^{((a_1, b_1), (a_2, b_2)), ((a_1, a_2), b_1)} \\ &\quad \circ \mathbf{ftuple}_2 (\mathcal{ABS}[[E_1]] \vec{X}) (\mathcal{ABS}[[E_2]] \vec{X}) \circ \mathbf{extract}^{x, (x, x)} \end{aligned}$$

By induction hypothesis, we have $\mathcal{ABS}[\llbracket E_1 \rrbracket \vec{X}] \in \text{Fun}_4$ and $\mathcal{ABS}[\llbracket E_2 \rrbracket \vec{X}] \in \text{Fun}_4$. So $\text{ftuple}_2 (\mathcal{ABS}[\llbracket E_1 \rrbracket \vec{X}]) (\mathcal{ABS}[\llbracket E_2 \rrbracket \vec{X}]) \in \text{Fun}_4$ and the right hand side expression belongs to Fun_4 .

B.2 Proof of the Property 2

To prove that the distribution transformation preserves the semantics, we show, for each transformation rule, that the left-hand side expression is semantically equal to the right-hand side one.

For example, in order to prove the poly_2 rule, we use the HASKELL definition of the inverse distribution $\text{dei} (\text{blocki } p) \text{ seqi}$:

```
dei (blocki p) seqi proc
= array (0,p*b-1)
  [(ip*b+i, array (0,n-1) [(j,proc!ip!i!j) | j<-[0..n-1]])
   | i<-[0..b-1], ip<-[0..p-1]]
  where p = sizeRange(bounds proc)
        b = sizeRange(bounds proc!0)
        n = sizeRange(bounds proc!0!0)
```

Let $\text{dei} (\text{blocki } p) \text{ seqi} : \text{Vectproc } p (\text{Vect } b (\text{Vect } n \alpha)) \rightarrow \text{Vect } m (\text{Vect } n \alpha)$,

Then, for all $\text{proc} : \text{Vectproc } p (\text{Vect } b (\text{Vect } n \alpha))$

$$\begin{aligned}
 & (\text{poly}_2 P F_1 F_2 \circ \text{dei} (\text{blocki } p) \text{ seqi}) \text{proc} \\
 &= \text{poly}_2 P F_1 F_2 (\text{dei} (\text{blocki } p) \text{ seqi } \text{proc}) && \text{definition} \\
 &= \text{poly}_2 P F_1 F_2 \\
 & \quad (\text{array } (0,p*b-1) [(ip*b+i, \text{array } (0,n-1) [(j, \text{proc}!ip!i!j) \\
 & \quad \quad \quad | j<-[0..n-1]]) | i<-[0..b-1], ip<-[0..p-1]]) && \text{dei definition} \\
 &= \text{array } (0,p*b-1) \\
 & \quad [(ip*b+i, \text{array } (0,n-1) \\
 & \quad \quad [(j, \text{if } P(ip*b+i,j) \text{ then } F_1(\text{proc}!ip!i!j) \text{ else } F_2(\text{proc}!ip!i!j)) \\
 & \quad \quad \quad | j<-[0..n-1]]) | i<-[0..b-1], ip<-[0..p-1]] && \text{poly}_2 \text{ definition} + ! \text{ definition} \\
 &= \text{dei} (\text{blocki } p) \text{ seqi} \\
 & \quad (\text{pimap } \lambda(ip,v).(\text{array } (0,b-1) [(i, \text{array } (0,n-1) \\
 & \quad \quad [(j, \text{if } P(ip*b+i,j) \text{ then } F_1(v!i!j) \text{ else } F_2(v!i!j)) \\
 & \quad \quad \quad | j<-[0..n-1]] | i<-[0..b-1]]) \text{proc}) && \text{dei definition} + \text{pimap definition} \\
 &= \text{dei} (\text{blocki } p) \text{ seqi} \\
 & \quad (\text{pimap } (\lambda(ip,v).\text{poly}_2 P \circ \lambda(i,j).(ip*b+i,j) F_1 F_2 v) \text{proc}) && \text{poly}_2 \text{ definition}
 \end{aligned}$$

$$= \text{dei (blocki } p) \text{ seqi} \\ \circ \text{pimap } (\lambda(ip, v). \text{poly}_2 P \circ \lambda(i, j).(ip * b + i, j) F_1 F_2 v)) \text{ proc} \quad \circ \text{definition}$$

Thus,

$$\text{poly}_2 P F_1 F_2 \circ \text{dei (blocki } p) \text{ seqi} \\ = \text{dei (blocki } p) \text{ seqi} \circ \text{pimap } (\lambda(ip, v). \text{poly}_2 P \circ \lambda(i, j).(ip * b + i, j) F_1 F_2 v) \\ = \text{DIST}[\text{poly}_2 P F_1 F_2 \circ \text{dei (blocki } p) \text{ seqi}]$$

C Examples of Programs

C.1 LU decomposition

The LU decomposition computes the decomposition of a square matrix M in two matrices L (triangular lower) and U (triangular higher) such as $M = L * U$.

The \mathcal{L}_1 program implementing this algorithm is given Figure 17. This program takes a matrix M of float of size $n \times n$ and returns a matrix containing the matrices L and U . The algorithm consists of an iteration (main function LU). At the k th step, the loop body determines the k th line and the k th column of the matrix result. This calculation is carried out by elimination by taking the element of the line $k - 1$ and the column $k - 1$ as pivot. The function $calc$ applies the function $fcalc$ to all the elements having indexes of line and column greater than k . This computation is iterated until the calculation of the n th line and the n th column. A FORTRAN version of this program can be found in [PTVF86].

```

LU(M) where
    M :: Vect n (Vect n Float)
    LU = iterfor (n - 1) loop
    loop = λ(k, a).calc(fac(apivot(colrow(k, a))))
    calc = λ(k, a, row, piv).poly2 (λ(i, j).k ≤ i ∧ k ≤ j) fcalc first
      (map zip3(zip3(a, row, piv)))
    fcalc = λ(a, row, piv).a - row * piv
    fac = λ(k, a, row, col, piv).(k, a, row, map (map /) (map zip (zip (col, piv))))
    apivot = λ(k, a, row, col).(k, a, row, col, map (brdcast (k - 1)) row)
    colrow = λ(k, a).(k, a, brdcast (k - 1) a, map (brdcast (k - 1)) a)
    zip3 = λ(x, y, z).map p2t (zip (x, zip (y, z)))
    p2t = λ(x, (y, z)).(x, y, z)
    first = λ(x, y, x).x
    
```

Figure 17: LU decomposition in \mathcal{L}_1 .

C.2 The n-body problem

The n-body problem concerns the simulation of N particles interacting via a long-range force such as gravity.

The \mathcal{L}_1 program given Figure 18 evaluates the new configuration of N particles after one unit of time. It takes a vector V of size N of a 7-tuple of floats. Each tuple element represents parameters of one particle (coordinates, speed, weight, ...) and returns a vector containing the new parameters of particles. The algorithm is an iteration (main function

NBODY) computing the interaction between each pair in turn. A FORTRAN version of this program can be found in [FJL⁺88].

```

NBODY(V) where
  V :: Vect N (Float, Float, Float, Float, Float, Float, Float)
  NBODY = λv.end(iterfor (div (N - 1) 2) loop (v, v))
  loop = λ(i, (p, q)).floop (p, rotate 1 q)
  floop = λ(p, q).unzip (map assign(zip3(p, q, (calc(zip (p, q))))))
  calc = λv.map (txyz.fac.sq.dxyz) v
  dxyz = λ((pm, px, py, pz, a, b, c), (qm, qx, qy, qz, a, b, c)).
    (pm * qm, px - qx, py - qy, pz - qz, 0)
  sq = λ(a, b, c, d, e).(a, b, c, d, ((b * b) + c * c) + (d * d))
  fac = λ(a, b, c, d, e).(a / (e * (sqrt e)), b, c, d, e)
  txyz = λ(a, b, c, d, e).(a * b, a * c, a * d)
  assign = λ((pm, px, py, pz, pfx, pfy, pfz), (qm, qx, qy, qz, qfx, qfy, qfz), (tx, ty, tz)).
    ((pm, px, py, pz, pfx - tx, pfy - ty, pfz - tz),
     (qm, qx, qy, qz, qfx + tx, qfy + ty, qfz + tz))
  end = λ(p, q).map plus(zip (p, rotate (- (÷ N 2)) q))
  plus = λ((a1, b1, c1, d1, pfx, pfy, pfz), (a2, b2, c2, d2, qfx, qfy, qfz)).
    (pfx + qfx, pfy + qfy, pfz + qfz)
  zip3 = λ(x, y, z).(map p2t(zip (zip (x, y), z)))
  p2t = λ((x, y), z).(x, y, z)

```

Figure 18: N-bodies in \mathcal{L}_1 .



Unité de recherche INRIA Lorraine, Technopôle de Nancy-Brabois, Campus scientifique,
615 rue du Jardin Botanique, BP 101, 54600 VILLERS LÈS NANCY
Unité de recherche INRIA Rennes, Irista, Campus universitaire de Beaulieu, 35042 RENNES Cedex
Unité de recherche INRIA Rhône-Alpes, 655, avenue de l'Europe, 38330 MONTBONNOT ST MARTIN
Unité de recherche INRIA Rocquencourt, Domaine de Voluceau, Rocquencourt, BP 105, 78153 LE CHESNAY Cedex
Unité de recherche INRIA Sophia-Antipolis, 2004 route des Lucioles, BP 93, 06902 SOPHIA-ANTIPOLIS Cedex

Éditeur
INRIA, Domaine de Voluceau, Rocquencourt, BP 105, 78153 LE CHESNAY Cedex (France)
<http://www.inria.fr>
ISSN 0249-6399