

# Program Transformation to Identify List-Based Parallel Skeletons

Venkatesh Kannan

G. W. Hamilton

School of Computing, Dublin City University, Ireland

{vkannan, hamilton}@computing.dcu.ie

Algorithmic skeletons are used as building-blocks to ease the task of parallel programming by abstracting the details of parallel implementation from the developer. Most existing libraries provide implementations of skeletons that are defined over flat data types such as lists or arrays. However, skeleton-based parallel programming is still very challenging as it requires intricate analysis of the underlying algorithm and often uses inefficient intermediate data structures. Further, the algorithmic structure of a given program may not match those of list-based skeletons. In this paper, we present a method to automatically transform any given program to one that is defined over a list and is more likely to contain instances of list-based skeletons. This facilitates the parallel execution of a transformed program using existing implementations of list-based parallel skeletons. Further, by using an existing transformation called *distillation* in conjunction with our method, we produce transformed programs that contain fewer inefficient intermediate data structures.

## 1 Introduction

In today's computing systems, parallel hardware architectures that use multi-core CPUs and GPUs (Graphics Processor Units) are ubiquitous. On such hardware, it is essential that the programs developed be executed in parallel in order to effectively utilise the computing power that is available. To enable this, the parallelism that is inherent in a given program needs to be identified and exploited. However, parallel programming is tedious and error-prone when done by hand and is very difficult for a compiler to do automatically to the desired level.

To ease the task of parallel programming, a collection of algorithmic skeletons [5, 6] are often used for program development to abstract away from the complexity of implementing the parallelism. In particular, *map*, *reduce* and *zipWith* are primitive parallel skeletons that are often used for parallel programming [20, 8]. Most libraries such as Eden [18], SkeTo [22], Data Parallel Haskell (DPH) [3], and Accelerate [2] provide parallel implementations for these skeletons defined over flat data types such as lists or arrays. However, there are two main challenges in skeleton-based programming:

1. Using multiple skeletons in a program often introduces inefficient intermediate data structures [21, 23].
2. There may be a mismatch in data structures and algorithms used by the skeletons and the program [27, 12].

For example, consider the matrix multiplication program shown in Example 1.1, where *mMul* computes the product of two matrices *xss* and *yss*. The function *map* is used to compute the dot-product (*dotp*) of each row in *xss* and those in the transpose of *yss*, which is computed by the function *transpose*. Note that this definition uses multiple intermediate data structures, which is inefficient.

**Example 1.1** (Matrix Multiplication – Original Program)
$$mMul :: [[a]] \rightarrow [[a]] \rightarrow [[a]]$$

$$mMul\ xss\ yss$$

**where**

$$mMul\ []\ yss$$

$$= []$$

$$mMul\ (xs : xss)\ yss = (map\ (transpose\ yss)\ (dotp\ xs)) : (mMul\ xss\ yss)$$

$$dotp\ xs\ ys = foldr\ (+)\ 0\ (zipWith\ (*)\ xs\ ys)$$

$$transpose\ xss = transpose'\ xss\ []$$

$$transpose'\ []\ yss = yss$$

$$transpose'\ (xs : xss)\ yss = transpose'\ xss\ (rotate\ xs\ yss)$$

$$rotate\ []\ yss = yss$$

$$rotate\ (x : xs)\ [] = [x] : (rotate\ xs\ yss)$$

$$rotate\ (x : xs)\ (ys : yss) = (ys ++ [x]) : (rotate\ xs\ yss)$$

A version of this program defined using the built-in *map*, *reduce* and *zipWith* skeletons is shown in Example 1.2.

**Example 1.2** (Hand-Parallelised Matrix Multiplication)
$$mMul\ xss\ yss$$

**where**

$$mMul\ []\ yss$$

$$= []$$

$$mMul\ (xs : xss)\ yss = (map\ (dotp\ xs)\ (transpose\ yss)) : (mMul\ xss\ yss)$$

$$dotp\ xs\ ys = reduce\ (+)\ 0\ (zipWith\ (*)\ xs\ ys)$$

$$transpose\ xss = transpose'\ xss\ []$$

$$transpose'\ []\ yss = yss$$

$$transpose'\ (xs : xss)\ yss = transpose'\ xss\ (rotate\ xs\ yss)$$

$$rotate\ xs\ yss = zipWith\ (\lambda x.\lambda ys.(ys ++ [x]))\ xs\ yss$$

As we can observe, though defined using parallel skeletons, this implementation still employs multiple intermediate data structures. For instance, the matrix constructed by the *transpose* function is subsequently decomposed by *map*. It is challenging to obtain a program that uses skeletons for parallel evaluation and contains very few intermediate data structures.

Therefore, it is desirable to have a method to automatically identify potential parallel computations in a given program, transform them to operate over flat data types to facilitate their execution using parallel skeletons provided in existing libraries, and reduce the number of inefficient intermediate data structures used.

In this paper, we present a transformation method with the following aspects:

1. Reduces inefficient intermediate data structures in a given program using an existing transformation technique called *distillation* [11]. (Section 3)
2. Automatically transforms the distilled program by *encoding* its inputs into a single *cons*-list, referred to as the *encoded list*. (Section 4)
3. Allows for parallel execution of the encoded program using efficient implementations of *map* and *map-reduce* skeletons that operate over lists. (Section 5)

In Section 6, we discuss the results of evaluating our proposed transformation method using two example programs. In Section 7, we present concluding remarks on possible improvements to our transformation method and discuss related work.

## 2 Language

We focus on the automated parallelisation of functional programs because pure functional programs are free of side-effects, which makes them easier to analyse, reason about, and manipulate using program transformation techniques. This facilitates parallel evaluation of independent sub-expressions in a program. The higher-order functional language used in this work is shown in Definition 2.1.

**Definition 2.1** (Language Grammar)

<b>data</b> $T \alpha_1 \dots \alpha_M ::= c_1 t_1^1 \dots t_N^1 \mid \dots \mid c_K t_1^K \dots t_N^K$	Type Declaration
$t ::= \alpha_m \mid T t_1 \dots t_M$	Type Component
$e ::= x$	Variable
$\mid c e_1 \dots e_N$	Constructor Application
$\mid e_0$	Function Definition
<b>where</b>	
$f p_1^1 \dots p_M^1 x_{(M+1)}^1 \dots x_N^1 = e_1 \dots f p_1^K \dots p_M^K x_{(M+1)}^K \dots x_N^K = e_K$	
$\mid f$	Function Call
$\mid e_0 e_1$	Application
$\mid \mathbf{let} x_1 = e_1 \dots x_N = e_N \mathbf{in} e_0$	<b>let</b> -expression
$\mid \lambda x. e$	$\lambda$ -Abstraction
$p ::= x \mid c p_1 \dots p_N$	Pattern

A program can contain data type declarations of the form shown in Definition 2.1. Here,  $T$  is the name of the data type, which can be polymorphic, with type parameters  $\alpha_1, \dots, \alpha_M$ . A data constructor  $c_k$  may have zero or more components, each of which may be a type parameter or a type application. An expression  $e$  of type  $T$  is denoted by  $e :: T$ .

A program in this language can also contain an expression which can be a variable, constructor application, function definition, function call, application, **let**-expression or  $\lambda$ -expression. Variables introduced in a  $\lambda$ -expression, **let**-expression, or function definition are *bound*, while all other variables are *free*. Each constructor has a fixed arity. In an expression  $c e_1 \dots e_N$ ,  $N$  must be equal to the arity of the constructor  $c$ . Patterns in a function definition header are grouped into two –  $p_1^k \dots p_M^k$  are inputs that are pattern-matched, and  $x_{(M+1)}^k \dots x_N^k$  are inputs that are not pattern-matched. The series of patterns  $p_1^1 \dots p_M^1, \dots, p_1^K \dots p_M^K$  in a function definition must be non-overlapping and exhaustive. We use  $[]$  and  $(:)$  as short notations for the *Nil* and *Cons* constructors of a *cons*-list and  $++$  for list concatenation. The set of free variables in an expression  $e$  is denoted as  $fv(e)$ .

**Definition 2.2** (Context)

A context  $E$  is an expression with *holes* in place of sub-expressions.  $E[e_1, \dots, e_N]$  is the expression obtained by filling holes in context  $E$  with the expressions  $e_1, \dots, e_N$ .

The call-by-name operational semantics of our language is defined using the one-step reduction relation shown in Definition 2.3.

**Definition 2.3** (One-Step Reduction Relation)

$$\begin{array}{c}
 ((\lambda x. e_0) e_1) \xrightarrow{\beta} (e_0 \{x \mapsto e_1\}) \qquad \frac{e_0 \xrightarrow{r} e'_0}{(e_0 e_1) \xrightarrow{r} (e'_0 e_1)} \qquad \frac{e_1 \xrightarrow{r} e'_1}{(e_0 e_1) \xrightarrow{r} (e_0 e'_1)} \\
 \\
 \frac{(f p_1 \dots p_N = e) \wedge (\exists \theta \cdot \forall n \in \{1, \dots, N\} \cdot e_n = p_n \theta)}{(f e_1 \dots e_N) \xrightarrow{f} e \theta} \\
 \\
 (\mathbf{let} x_1 = e_1 \dots x_N = e_N \mathbf{in} e_0) \xrightarrow{\beta} (e_0 \{x \mapsto e_1, \dots, x_N \mapsto e_N\})
 \end{array}$$

### 3 Distillation

**Objective:** A given program may contain a number of inefficient intermediate data structures. In order to reduce them, we use an existing transformation technique called *distillation*.

*Distillation* [11] is a technique that transforms a program to remove intermediate data structures and yields a *distilled program*. It is an unfold/fold-based transformation that makes use of well-known transformation steps – unfold, generalise and fold [25] – and can potentially provide super-linear speedups to programs. The syntax of a distilled program  $de^{\{\}}$  is shown in Definition 3.1. Here,  $\rho$  is the set of variables introduced by **let**-expressions; these are not decomposed by pattern-matching. Consequently,  $de^{\{\}}$  is an expression that has fewer intermediate data structures.

**Definition 3.1** (Distilled Form Grammar)

$de^{\rho} ::= x de_1^{\rho} \dots de_N^{\rho}$	Variable Application
$c de_1^{\rho} \dots de_N^{\rho}$	Constructor Application
$de_0^{\rho}$	Function Definition
<b>where</b>	
$f p_1^1 \dots p_M^1 x_{(M+1)}^1 \dots x_N^1 = de_1^{\rho} \dots f p_1^K \dots p_M^K x_{(M+1)}^K \dots x_N^K = de_K^{\rho}$	
$f x_1 \dots x_M x_{(M+1)} \dots x_N$	Function Application
s.t. $\forall x \in \{x_1, \dots, x_M\} \cdot x \notin \rho$	
$\forall n \in \{1, \dots, N\} \cdot (x_n \in \rho \Rightarrow \forall k \in \{1, \dots, K\} \cdot p_n^k = x_n^k)$	
<b>let</b> $x_1 = de_1^{\rho} \dots x_N = de_N^{\rho}$ <b>in</b> $de_1^{\rho \cup \{x_1, \dots, x_N\}}$	<b>let</b> -expression
$\lambda x. de^{\rho}$	$\lambda$ -Abstraction
$p ::= x \mid c \mid p_1 \dots p_N$	Pattern

Example 3.1 shows the distilled form of the example matrix multiplication program in Example 1.1. Here, we have lifted the definitions of functions  $mMul_2$  and  $mMul_3$  to the top level using lambda lifting for ease of presentation.

**Example 3.1** (Matrix Multiplication – Distilled Program)

$mMul \ xss \ yss$	
<b>where</b>	
$mMul \ xss \ yss$	$= mMul_1 \ xss \ yss \ yss$
$mMul_1 \ [] \ zss \ yss$	$= []$
$mMul_1 \ xss \ [] \ yss$	$= []$
$mMul_1 \ (xs : xss) \ (zs : zss) \ yss$	$= \mathbf{let} \ v = \lambda xs.g \ xs$
<b>where</b>	
$g \ [] = 0$	
$g \ (x : xs) = x$	
<b>in</b> $(mMul_2 \ zs \ xs \ yss \ v) : (mMul_1 \ xss \ zss \ yss)$	
$mMul_2 \ [] \ xs \ yss \ v$	$= []$
$mMul_2 \ (z : zs) \ xs \ yss \ v$	$= \mathbf{let} \ v' = \lambda xs.g \ xs$
<b>where</b>	
$g \ [] = 0$	
$g \ (x : xs) = v \ xs$	
<b>in</b> $(mMul_3 \ xs \ yss \ v) : (mMul_2 \ zs \ xs \ yss \ v')$	
$mMul_3 \ [] \ yss \ v$	$= 0$
$mMul_3 \ (x : xs) \ [] \ v$	$= 0$
$mMul_3 \ (x : xs) \ (ys : yss) \ v$	$= (x * (v \ ys)) + (mMul_3 \ xs \ yss \ v)$

In this distilled program, function  $mMul_1$  computes the product of matrices  $xss$  and  $yss$ , and functions  $mMul_2$  and  $mMul_3$  compute the dot-product of a row in  $xss$  and those in the transpose of  $yss$ . This version of matrix multiplication is free from intermediate data structures. In particular, distillation removes data structures that are constructed and subsequently decomposed as a part of the algorithm that is implemented in a given program.

**Consequence:** Using the distillation transformation, we obtain a semantically equivalent version of the original program that has fewer intermediate data structures.

## 4 Encoding Transformation

**Objective:** The data types and the algorithm of a distilled program, which we want to parallelise, may not match with those of the skeletons defined over lists. This would inhibit the potential identification of parallel computations that could be encapsulated using the map or map-reduce skeletons. To resolve this, we define a transformation that encodes the inputs of a distilled program into a single *cons*-list. The resulting encoded program is defined in a form that facilitates identification of list-based parallel skeleton instances.

To perform the encoding transformation, we first lift the definitions of all functions in a distilled program to the top-level using lambda lifting. Following this, for each recursive function  $f$  defined in the top-level **where**-expression of the distilled program, we encode the inputs  $p_1, \dots, p_M$  that are pattern-matched in the definition of  $f$ . Other inputs  $x_{(M+1)}, \dots, x_N$  that are never pattern-matched in the definition of  $f$  are not encoded. Further, we perform this encoding only for the recursive functions in a distilled program because they are potential instances of parallel skeletons, which are also defined recursively. The three steps to encode inputs  $x_1, \dots, x_M$  of function  $f$  into a *cons*-list, referred to as the *encoded list*, are illustrated in Figure 1 and described below. Here, we encode the pattern-matched inputs  $x_1, \dots, x_M$  into a *cons*-list of type  $[T_f]$ , where  $T_f$  is a new type created to contain the pattern-matched variables in  $x_1, \dots, x_M$ .

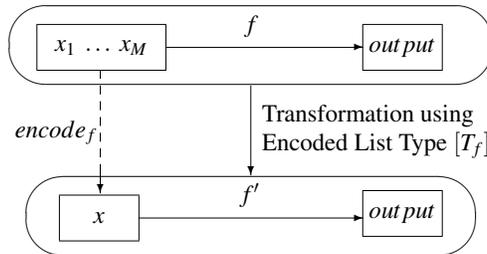


Figure 1: Steps to Encode Inputs of Function  $f$

Consider the definition of a recursive function  $f$ , with inputs  $x_1, \dots, x_M, x_{(M+1)}, \dots, x_N$ , of the form shown in Definition 4.1 in a distilled program. Here, for each body  $e_k$  corresponding to function header  $f p_1^k \dots p_M^k x_{(M+1)}^k \dots x_N^k$  in the definition of  $f$ , we use one of the recursive calls to function  $f$  that may appear in  $e_k$ . All other recursive calls to  $f$  in  $e_k$  are a part of the context  $E_k$ .

**Definition 4.1** (General Form of Recursive Function in Distilled Program)

$f x_1 \dots x_M x_{(M+1)} \dots x_N$

**where**

$f p_1^1 \dots p_M^1 x_{(M+1)}^1 \dots x_N^1 = e_1$

$\vdots$

$f p_1^K \dots p_M^K x_{(M+1)}^K \dots x_N^K = e_K$

where  $\exists k \in \{1, \dots, K\} \cdot e_k = E_k [f x_1^k \dots x_M^k x_{(M+1)}^k \dots x_N^k]$

The three steps to encode the pattern-matched inputs are as follows:

### 1. Declare a new encoded data type $T_f$ :

First, we declare a new data type  $T_f$  for elements of the encoded list. This new data type corresponds to the data types of the pattern-matched inputs of function  $f$  that are encoded. The rules to declare type  $T_f$  are shown in Definition 4.2.

**Definition 4.2** (Rules to Declare Encoded Data Type for List)

**data**  $T_f \alpha_1 \dots \alpha_G ::= c_1 T_1^1 \dots T_L^1 \mid \dots \mid c_K T_1^K \dots T_L^K$

where

$\alpha_1, \dots, \alpha_G$  are the type variables of the data types of the pattern-matched inputs

$\forall k \in \{1, \dots, K\}$ .

$c_k$  is a fresh constructor for  $T_f$  corresponding to  $p_1^k \dots p_M^k$  of the pattern-matched inputs

$$\{(z_1 :: T_1^k), \dots, (z_L :: T_L^k)\} = \begin{cases} fv(E_k) \setminus \{x_{(M+1)}, \dots, x_N\}, & \text{if } e_k = E_k [f x_1^k \dots x_M^k x_{(M+1)}^k \dots x_N^k] \\ fv(e_k) \setminus \{x_{(M+1)}, \dots, x_N\}, & \text{otherwise} \end{cases}$$

$$\text{where } f p_1^k \dots p_M^k x_{(M+1)} \dots x_N = e_k$$

Here, a new constructor  $c_k$  of the type  $T_f$  is created for each set  $p_1^k \dots p_M^k$  of the pattern-matched inputs  $x_1 \dots x_M$  of function  $f$  that are encoded. As stated above, our objective is to encode the inputs of a recursive function  $f$  into a list, where each element contains the pattern-matched variables consumed in an iteration of  $f$ . To achieve this, the variables bound by constructor  $c_k$  correspond to the variables  $z_1, \dots, z_L$  in  $p_1^k \dots p_M^k$  that occur in the context  $E_k$  (if  $e_k$  contains a recursive call to  $f$ ) or the expression  $e_k$  (otherwise). Consequently, the type components of constructor  $c_k$  are the data types of the variables  $z_1, \dots, z_L$ .

### 2. Define a function $encode_f$ :

For a recursive function  $f$  of the form shown in Definition 4.1, we use the rules in Definition 4.3 to define function  $encode_f$  to build the encoded list, in which each element is of type  $T_f$ .

**Definition 4.3** (Rules to Define Function  $encode_f$ )

$encode_f x_1 \dots x_M$

**where**

$$encode_f p_1^1 \dots p_M^1 = e'_1$$

$$\vdots$$

$$encode_f p_1^K \dots p_M^K = e'_K$$

where

$$\forall k \in \{1, \dots, K\} \cdot e'_k = \begin{cases} [c_k z_1^k \dots z_L^k] ++ (encode_f x_1^k \dots x_M^k), & \text{if } e_k = E_k [f x_1^k \dots x_M^k x_{(M+1)}^k \dots x_N^k] \\ \text{where } \{z_1^k, \dots, z_L^k\} = fv(E_k) \setminus \{x_{(M+1)}, \dots, x_N\} \\ [c_k z_1^k \dots z_L^k], & \text{otherwise} \\ \text{where } \{z_1^k, \dots, z_L^k\} = fv(e_k) \setminus \{x_{(M+1)}, \dots, x_N\} \\ \text{where } f p_1^k \dots p_M^k x_{(M+1)} \dots x_N = e_k \end{cases}$$

Here, for each pattern  $p_1^k \dots p_M^k$  of the pattern-matched inputs, the  $encode_f$  function creates a list element. This element is composed of a fresh constructor  $c_k$  of type  $T_f$  that binds  $z_1^k, \dots, z_L^k$ , which are the variables in  $p_1^k \dots p_M^k$  that occur in the context  $E_k$  (if  $e_k$  contains a recursive call to  $f$ ) or the expression  $e_k$  (otherwise). The encoded input of the recursive call  $f x_1^k \dots x_M^k x_{(M+1)}^k \dots x_N^k$  is then computed by  $encode_f x_1^k \dots x_M^k$  and appended to the element to build the complete encoded list for function  $f$ .

### 3. Transform the distilled program :

After creating the data type  $T_f$  for the encoded list and the  $encode_f$  function for each recursive function  $f$ , we transform the distilled program using the rules in Definition 4.4 by defining a recursive function  $f'$ , which operates over the encoded list, corresponding to function  $f$ .

**Definition 4.4** (Rules to Define Encoded Function Over Encoded List)
$$f' x x_{(M+1)} \dots x_N$$
**where**

$$f' ((c_1 z_1^1 \dots z_L^1) : x^1) x_{(M+1)} \dots x_N = e'_1$$

$$\vdots$$

$$\vdots$$

$$f' ((c_K z_1^K \dots z_L^K) : x^K) x_{(M+1)} \dots x_N = e'_K$$

**where**

$$\forall k \in \{1, \dots, K\} \cdot e'_k = \begin{cases} E_k [f' x^k x_{(M+1)}^k \dots x_N^k], & \text{if } e_k = E_k [f x_1^k \dots x_M^k x_{(M+1)}^k \dots x_N^k] \\ e_k, & \text{otherwise} \end{cases}$$

where  $f p_1^k \dots p_M^k x_{(M+1)}^k \dots x_N^k = e_k$

Here,

- In each function definition header of  $f$ , replace the pattern-matched inputs with a pattern to decompose the encoded list, such that the first element in the encoded list is matched with the corresponding pattern of the encoded type. For instance, a function header  $f p_1 \dots p_M x_{(M+1)} \dots x_N$  is transformed to  $f' p x_{(M+1)} \dots x_N$ , where  $p$  is a pattern to match the first element in the encoded list with a pattern of the type  $T_f$ .
- In each call to function  $f$ , replace the pattern-matched inputs with their encoding. For instance, a call  $f x_1 \dots x_M x_{(M+1)} \dots x_N$  is transformed to  $f' x x_{(M+1)} \dots x_N$ , where  $x$  is the encoding of the pattern-matched inputs  $x_1, \dots, x_M$ .

The encoded data types, encode functions and encoded program obtained for the distilled matrix multiplication program from Example 3.1 are shown in Example 4.1.

**Example 4.1** (Matrix Multiplication – Encoded Program)
$$\mathbf{data} T_{mMul_1} a ::= c_1 \mid c_2 \mid c_3 [a] [a]$$

$$\mathbf{data} T_{mMul_2} a ::= c_4 \mid c_5$$

$$\mathbf{data} T_{mMul_3} a ::= c_6 \mid c_7 \mid c_8 a [a]$$

$$\begin{aligned} \mathit{encode}_{mMul_1} [] zss &= [c_1] \\ \mathit{encode}_{mMul_1} xss [] &= [c_2] \\ \mathit{encode}_{mMul_1} (xs : xss) (zs : zss) &= [c_3 xs zs] ++ (\mathit{encode}_{mMul_1} xss zss) \end{aligned}$$

$$\begin{aligned} \mathit{encode}_{mMul_2} [] &= [c_4] \\ \mathit{encode}_{mMul_2} (z : zs) &= [c_5] ++ (\mathit{encode}_{mMul_2} xs yss zs) \end{aligned}$$

$$\begin{aligned} \mathit{encode}_{mMul_3} [] yss &= [c_6] \\ \mathit{encode}_{mMul_3} (x : xs) [] &= [c_7] \\ \mathit{encode}_{mMul_3} (x : xs) (ys : yss) &= [c_8 x ys] ++ (\mathit{encode}_{mMul_3} xs yss) \end{aligned}$$

$$mMul' xss yss$$
**where**

$$mMul' xss yss = mMul'_1 (\mathit{encode}_{mMul_1} xss yss) yss$$

$$mMul'_1 (c_1 : w) yss = []$$

$$mMul'_1 (c_2 : w) yss = []$$

$$mMul'_1 ((c_3 xs zs) : w) yss = \mathbf{let} v = \lambda xs.g xs$$

**where**

$$g [] = 0$$

$$g (x : xs) = x$$

$$\mathbf{in} (mMul'_2 (\mathit{encode}_{mMul_2} zs) xs yss v) : (mMul'_1 w yss)$$

$$mMul'_2 (c_4 : w) xs yss v = []$$

$$mMul'_2 (c_5 : w) xs yss v = \mathbf{let} v' = \lambda xs.g xs$$

**where**

$$g [] = 0$$

$$g (x : xs) = v xs$$

$$\mathbf{in} (mMul'_3 (\mathit{encode}_{mMul_3} xs yss) v) : (mMul'_2 w xs yss v')$$

$$\begin{aligned}
mMul'_3 (c_6 : w) v &= 0 \\
mMul'_3 (c_7 : w) v &= 0 \\
mMul'_3 ((c_8 x ys) : w) v &= (x * (v ys)) + (mMul'_3 w v)
\end{aligned}$$

## 4.1 Correctness

The correctness of the encoding transformation can be established by proving that the result computed by each recursive function  $f$  in the distilled program is the same as the result computed by the corresponding recursive function  $f'$  in the encoded program. That is,

$$\begin{aligned}
(f x_1 \dots x_M x_{(M+1)} \dots x_N) &= (f' x x_{(M+1)} \dots x_N) \\
\text{where } x &= \text{encode}_f x_1 \dots x_M
\end{aligned}$$

### Proof:

The proof is by structural induction over the encoded list type  $[T_f]$ .

#### Base Case:

For the encoded list  $x^k = ((c_k z_1^k \dots z_L^k) : [])$  computed by  $\text{encode}_f p_1^k \dots p_M^k$ ,

1. By Definition 4.1, L.H.S. evaluates to  $e_k$ .
2. By Definition 4.4, R.H.S. evaluates to  $e_k$ .

#### Inductive Case:

For the encoded list  $x^k = ((c_k z_1^k \dots z_L^k) : x^k)$  computed by  $\text{encode}_f p_1^k \dots p_M^k$ ,

1. By Definition 4.1, L.H.S. evaluates to  $E_k [f x_1^k \dots x_M^k x_{(M+1)}^k \dots x_N^k]$ .
2. By Definition 4.4, R.H.S. evaluates to  $E_k [f' x^k x_{(M+1)}^k \dots x_N^k]$ .
3. By inductive hypothesis,  $(f x_1 \dots x_M x_{(M+1)} \dots x_N) = (f' x x_{(M+1)} \dots x_N)$ . □

**Consequence:** As a result of the encoding transformation, the pattern-matched inputs of a recursive function are encoded into a *cons*-list by following the recursive structure of the function. Parallelisation of the encoded program produced by this transformation by identifying potential instances of *map* and *map-reduce* skeletons is discussed in Section 5.

## 5 Parallel Execution of Encoded Programs

**Objective:** An encoded program defined over an encoded list is more likely to contain recursive functions that resemble the structure of *map* or *map-reduce* skeletons. This is because the  $\text{encode}_f$  function constructs the encoded list in such a way that it reflects the recursive structure of the *map* and *map-reduce* skeletons defined over a *cons*-list. Therefore, we look for instances of these skeletons in our encoded program.

In this work, we identify instances of only *map* and *map-reduce* skeletons in an encoded program. This is because, as shown in Property 5.1, any function that is an instance of a reduce skeleton in an encoded program that operates over an encoded list cannot be efficiently evaluated in parallel because the reduction operator will not be associative.

**Property 5.1** (Non-Associative Reduction Operator for Encoded List)

Given an encoded program defined over an encoded list, the reduction operator  $\oplus$  in any instance of a reduce skeleton is not associative, that is  $\forall x, y, z. (x \oplus (y \oplus z)) \neq ((x \oplus y) \oplus z)$ .

**Proof:**

1. From Definition 4.4, given an encoded function  $f'$ ,

$$f' :: [T_f] \rightarrow T_{(M+1)} \dots T_N \rightarrow b$$

where  $[T_f]$  is the encoded list data type.

$T_{(M+1)}, \dots, T_N$  are data types for inputs that are not encoded.

$b$  is the output data type.

2. If  $f'$  is an instance of a reduce skeleton, then the type of the binary reduction operator is given by  $\oplus :: T_f \rightarrow b \rightarrow b$ .
3. Given that  $T_f$  is a newly created data type, it follows from (2) that the binary operator  $\oplus$  is not associative because the two input data types  $T_f$  and  $b$  cannot not be equal.  $\square$

**5.1 Identification of Skeletons**

To identify skeleton instances in a given program, we use a framework of *labelled transition systems* (LTSs), presented in Definition 5.1, to represent and analyse the encoded programs and skeletons. This is because LTS representations enable matching the recursive structure of the encoded program with that of the skeletons rather than finding instances by matching expressions.

**Definition 5.1** (Labelled Transition System (LTS))

A LTS for a given program is given by  $l = (\mathcal{S}, s_0, Act, \rightarrow)$  where:

- $\mathcal{S}$  is the set of *states* of the LTS, where each state has a unique label  $s$ .
- $s_0 \in \mathcal{S}$  is the start state denoted by  $start(l)$ .
- $Act$  is one of the following actions:
  - $x$ , a free variable or **let**-expression variable,
  - $c$ , a constructor in an application,
  - $\lambda$ , a  $\lambda$ -abstraction,
  - $@$ , an expression application,
  - $\#i$ , the  $i^{th}$  argument in an application,
  - $p$ , the set of patterns in a function definition header,
  - **let**, a **let**-expression body.
- $\rightarrow \subseteq \mathcal{S} \times Act \times \mathcal{S}$  relates pairs of states by actions in  $Act$  such that if  $s \in \mathcal{S}$  and  $s \xrightarrow{\alpha} s'$  then  $s' \in \mathcal{S}$  where  $\alpha \in Act$ .

The LTS corresponding to a given program  $e$  can be constructed by  $\mathcal{L}[[e]] s_0 \emptyset \emptyset$  using the rules  $\mathcal{L}$  shown in Definition 5.2. Here,  $s_0$  is the start state,  $\emptyset$  is the set of previously encountered function calls mapped to their corresponding states, and  $\Delta$  is the set of function definitions. A LTS built using these rules is always finite because if a function call is re-encountered, then the corresponding state is reused.

**Definition 5.2** (LTS Representation of Program)

$$\begin{aligned}
\mathcal{L}[[x]] s \phi \Delta &= s \rightarrow (x, \mathbf{0}) \\
\mathcal{L}[[c \ e_1 \dots e_N]] s \phi \Delta &= s \rightarrow (c, \mathbf{0}), (\#1, \mathcal{L}[[e_1]] s_1 \phi \Delta), \dots, (\#N, \mathcal{L}[[e_N]] s_N \phi \Delta) \\
\mathcal{L}[[e_0 \ \mathbf{where} \ \delta_1 \dots \delta_J]] s \phi \Delta &= \mathcal{L}[[e_0]] s \phi (\Delta \cup \{f_1 \mapsto \delta_1, \dots, f_J \mapsto \delta_J\}) \\
&\quad \text{where } \forall j \in \{1, \dots, J\} \cdot \delta_j = f_j \ p_1^1 \dots p_M^1 \ x_{(M+1)}^1 \dots x_N^1 = e_1 \dots f_j \ p_1^K \dots p_M^K \ x_{(M+1)}^K \dots x_N^K = e_K \\
\mathcal{L}[[f]] s \phi \Delta &= \begin{cases} l & \text{where } \phi(f) = \text{start}(l), \text{ if } f \in \text{dom}(\phi) \\ \mathcal{L}[[\Delta(f)]] s (\phi \cup \{f \mapsto s\}) \Delta, & \text{otherwise} \end{cases} \\
\mathcal{L} \left[ \left[ \begin{array}{c} f \ p_1^1 \dots p_M^1 \ x_{(M+1)}^1 \dots x_N^1 = e_1 \\ \vdots \\ f \ p_1^K \dots p_M^K \ x_{(M+1)}^K \dots x_N^K = e_K \end{array} \right] s \phi \Delta \right. &= \begin{cases} s \rightarrow (p_1^1 \dots p_M^1 \ x_{(M+1)}^1 \dots x_N^1, l_1), \dots, \\ (p_1^K \dots p_M^K \ x_{(M+1)}^K \dots x_N^K, l_K) \\ \text{where } \forall k \in \{1, \dots, K\} \cdot l_k = (\mathcal{L}[[e_k]] s_k \phi \Delta) \end{cases} \\
\mathcal{L}[[e_0 \ e_1]] s \phi \Delta &= s \rightarrow (@, \mathcal{L}[[e_0]] s_0 \phi \Delta), (\#1, \mathcal{L}[[e_1]] s_1 \phi \Delta) \\
\mathcal{L}[[\mathbf{let} \ x_1 = e_1 \dots x_N = e_N \ \mathbf{in} \ e_0]] s \phi \Delta &= s \rightarrow (\mathbf{let}, \mathcal{L}[[e_0]] s_0 \phi \Delta), \\
&\quad (x_1, \mathcal{L}[[e_1]] s_1 \phi \Delta), \dots, (x_N, \mathcal{L}[[e_N]] s_N \phi \Delta) \\
\mathcal{L}[[\lambda x. e]] s \phi \Delta &= s \rightarrow (\lambda, \mathcal{L}[[e]] s_1 \phi \Delta)
\end{aligned}$$

**Definition 5.3** (LTS Substitution)

A substitution is denoted by  $\theta = \{x_1 \mapsto l_1, \dots, x_N \mapsto l_N\}$ . If  $l$  is an LTS, then  $l\theta = l\{x_1 \mapsto l_1, \dots, x_N \mapsto l_N\}$  is the result of simultaneously replacing the LTSs  $s_n \rightarrow (x_n, \mathbf{0})$  with the corresponding LTS  $l_n$  in the LTS  $l$  while ensuring that bound variables are renamed appropriately to avoid name capture.

Potential instances of skeleton LTSs can be identified and replaced with suitable calls to corresponding skeletons in the LTS of an encoded program  $l$  by  $\mathcal{S}[[l]] \ \emptyset \ \langle \rangle \ \omega$  using the rules presented in Definition 5.4.

**Definition 5.4** (Extraction of Program from LTS with Skeletons)

$$\begin{aligned}
\mathcal{S}[[l]] \ \rho \ \sigma \ \omega &= \left. \begin{cases} f \ e_1 \dots e_N, & \text{if } \exists (f \ x_1 \dots x_N, l') \in \omega, \theta \cdot l'\theta = l \\ \text{where} \\ \theta = \{x_1 \mapsto l_1, \dots, x_N \mapsto l_N\} \\ \forall n \in \{1, \dots, N\} \cdot e_n = (\mathcal{S}[[l_n]] \ \rho \ \sigma \ \omega) \end{cases} \right\}, \text{ if } \exists s \in \text{states}(l), \alpha \cdot s \xrightarrow{\alpha} \text{start}(l) \\
&\quad \left. \begin{cases} f \ e_1 \dots e_N, & \text{if } \exists (s, f) \in \rho \cdot \text{start}(l) = s \\ \text{where } \sigma = \langle e_1, \dots, e_N \rangle \\ \mathcal{S}'[[l]] \ \rho \ \sigma \ \omega, & \text{otherwise} \\ \mathcal{S}'[[l]] \ \rho \ \sigma \ \omega, & \text{otherwise} \end{cases} \right\} \\
\mathcal{S}'[[s \rightarrow (x, \mathbf{0})]] \ \rho \ \sigma \ \omega &= x \ e_1 \dots e_N \quad \text{where } \sigma = \langle e_1, \dots, e_N \rangle \\
\mathcal{S}'[[s \rightarrow (c, \mathbf{0}), (\#1, l_1), \dots, (\#N, l_N)]] \ \rho \ \sigma \ \omega &= c \ (\mathcal{S}'[[l_1]] \ \rho \ \sigma \ \omega) \dots (\mathcal{S}'[[l_N]] \ \rho \ \sigma \ \omega) \\
\mathcal{S}' \left[ \left[ \begin{array}{c} s \rightarrow (p_1^1 \dots p_M^1 \ x_{(M+1)}^1 \dots x_N^1, l_1), \dots, \\ (p_1^K \dots p_M^K \ x_{(M+1)}^K \dots x_N^K, l_K) \end{array} \right] \rho \ \sigma \ \omega \right. &= \begin{cases} f \ e_1 \dots e_N \\ \text{where} \\ f \ p_1^1 \dots p_M^1 \ x_{(M+1)}^1 \dots x_N^1 = e'_1 \\ \vdots \\ f \ p_1^K \dots p_M^K \ x_{(M+1)}^K \dots x_N^K = e'_K \\ \text{where } f \text{ is fresh, } \sigma = \langle e_1, \dots, e_N \rangle \\ \forall k \in \{1, \dots, K\} \cdot e'_k = (\mathcal{S}'[[l_k]] \ \rho' \ \langle \rangle \ \omega) \\ \rho' = \rho \cup \{(s, f)\} \end{cases} \\
\mathcal{S}'[[s \rightarrow (@, l_0), (\#1, l_1)]] \ \rho \ \sigma \ \omega &= \mathcal{S}'[[l_0]] \ \rho \ \langle (\mathcal{S}'[[l_1]] \ \rho \ \omega \ \langle \rangle) : \sigma \rangle \ \omega \\
\mathcal{S}'[[s \rightarrow (\mathbf{let}, l_0), (x_1, l_1), \dots, (x_N, l_N)]] \ \rho \ \sigma \ \omega &= \mathbf{let} \ x_1 = (\mathcal{S}'[[l_1]] \ \rho \ \sigma \ \omega) \dots x_N = (\mathcal{S}'[[l_N]] \ \rho \ \sigma \ \omega) \\
&\quad \mathbf{in} \ (\mathcal{S}'[[l_0]] \ \rho \ \sigma \ \omega) \\
\mathcal{S}'[[s \rightarrow (\lambda, l)]] \ \rho \ \sigma \ \omega &= \lambda x. (\mathcal{S}'[[l]] \ \rho \ \sigma \ \omega) \quad \text{where } x \text{ is fresh}
\end{aligned}$$

Here, the parameter  $\rho$  contains the set of new functions that are created and associates them with their corresponding states in the LTS. The parameter  $\sigma$  contains the sequence of arguments of an application expression. The set  $\omega$  is initialised with pairs of application expression and corresponding LTS representation of each parallel skeleton to be identified in a given LTS; for example,  $(\text{map } xs \ f, l)$  is a pair in  $\omega$  where  $\text{map } xs \ f$  is the application expression for  $\text{map}$  and  $l$  is its LTS representation.



## 5.2 Parallel Implementation of Skeletons

In order to evaluate the parallel programs obtained by our method presented in this chapter, we require efficient parallel implementations of the map and map-reduce skeletons. For the work presented in this paper, we use the Eden library [18] that provides parallel implementations of the map and map-reduce skeletons in the following forms:

$$\begin{aligned} \mathit{farmB} &:: (\mathit{Trans} \ a, \mathit{Trans} \ b) \Rightarrow \mathit{Int} \rightarrow (a \rightarrow b) \rightarrow [a] \rightarrow [b] \\ \mathit{parMapRedr} &:: (\mathit{Trans} \ a, \mathit{Trans} \ b) \Rightarrow (b \rightarrow b \rightarrow b) \rightarrow b \rightarrow (a \rightarrow b) \rightarrow [a] \rightarrow b \\ \mathit{parMapRedl} &:: (\mathit{Trans} \ a, \mathit{Trans} \ b) \Rightarrow (b \rightarrow b \rightarrow b) \rightarrow b \rightarrow (a \rightarrow b) \rightarrow [a] \rightarrow b \end{aligned}$$

The *farmB* skeleton implemented in Eden divides a given list into  $N$  sub-lists and creates  $N$  parallel processes, each of which applies the map computation on a sub-list. The parallel map-reduce skeletons, *parMapRedr* and *parMapRedl*, are implemented using the *parMap* skeleton which applies the *map* operation in parallel on each element in a given list. The result of *parMap* is reduced sequentially using the conventional *foldr* and *foldl* functions, respectively.

Currently, the map-reduce skeletons in the Eden library are defined using the *foldr* and *foldl* functions that require a unit value for the reduction/fold operator to be provided as an input. However, it is evident from Property 5.2 that the skeletons that are potentially identified will always be applied on non-empty lists. Therefore, we augment the skeletons provided in Eden by adding the following parallel map-reduce skeletons that are defined using the *foldr1* and *foldl1* functions, which are defined for non-empty lists, thereby avoiding the need to obtain a unit value for the reduction operator.

$$\begin{aligned} \mathit{parMapRedr1} &:: (\mathit{Trans} \ a, \mathit{Trans} \ b) \Rightarrow (b \rightarrow b \rightarrow b) \rightarrow (a \rightarrow b) \rightarrow [a] \rightarrow b \\ \mathit{parMapRedl1} &:: (\mathit{Trans} \ a, \mathit{Trans} \ b) \Rightarrow (b \rightarrow b \rightarrow b) \rightarrow (a \rightarrow b) \rightarrow [a] \rightarrow b \end{aligned}$$

To execute the encoded program produced by our transformation in parallel, we replace the identified skeleton instances with suitable calls to the corresponding skeletons in the Eden library. For example, by replacing functions  $mMul'_1$  and  $mMul'_3$ , which are instances of *map* and *mapReduce* skeletons respectively, with suitable calls to *parMap* and *parMapRedr1*, we obtain the transformed matrix multiplication program  $mMul''$  shown in Example 5.1.

### Example 5.1 (Matrix Multiplication – Encoded Parallel Program)

$mMul'' \ xss \ yss$

**where**

$mMul'' \ xss \ yss = mMul''_1 (\mathit{encode}_{mMul_1} \ xss \ yss) \ yss$

$mMul''_1 \ w \ yss = \mathit{farmB} \ \mathit{noPe} \ f \ w$

**where**

$f \ c_1 = []$

$f \ c_2 = []$

$f \ (c_3 \ xs \ zs) = \mathbf{let} \ v = \lambda xs.g \ xs$

**where**

$g \ [] = 0$

$g \ (x : xs) = x$

**in**  $mMul''_2 (\mathit{encode}_{mMul_2} \ zs) \ xs \ yss \ v$

$mMul''_2 \ (c_4 : w) \ xs \ yss \ v = []$

$mMul''_2 \ (c_5 : w) \ xs \ yss \ v = \mathbf{let} \ v' = \lambda xs.g \ xs$

**where**

$g \ [] = 0$

$g \ (x : xs) = v \ xs$

**in**  $(mMul''_3 (\mathit{encode}_{mMul_3} \ xs \ yss) \ v) : (mMul''_2 \ w \ xs \ yss \ v')$

```

mMul3 w v = parMapRedr1 g f w
  where
    g x y      = x + y
    f c6      = 0
    f c7      = 0
    f (c8 x ys) = x * (v ys)

```

**Consequence:** By automatically identifying instances of list-based map and map-reduce skeletons, we produce a program that is defined using these parallelisable skeletons. Using parallel implementations for these skeletons that are available in existing libraries such as Eden, it is possible to execute the transformed program on parallel hardware.

## 6 Evaluation

In this paper, we present the evaluation of two benchmark programs – matrix multiplication and dot-product of binary trees – to illustrate interesting aspects of our transformation. The programs are evaluated on a Mac Pro computer with a 12-core Intel Xeon E5 processor each clocked at 2.7 GHz and 64 GB of main memory clocked at 1866 MHz. GHC version 7.10.2 is used for the sequential versions of the benchmark programs and the latest Eden compiler based on GHC version 7.8.2 for the parallel versions.

For all parallel versions of a benchmark program, only those skeletons that are present in the top-level expression are executed using their parallel implementations. That is, nesting of parallel skeletons is avoided. The nested skeletons that are present inside top-level skeletons are executed using their sequential versions. The objective of this approach is to avoid uncontrolled creation of too many threads which we observe to result in inefficient parallel execution where the cost of thread creation and management is greater than the cost of parallel execution.

### 6.1 Example – Matrix Multiplication

The original sequential version, distilled version, encoded version and encoded parallel version of the matrix multiplication program are presented in Examples 1.1, 3.1, 4.1 and 5.1 respectively.

A hand-parallel version of the original matrix multiplication program in Example 1.1 is presented in Example 6.1. We identify that function *mMul* is an instance of the *map* skeleton and therefore define it using a suitable call to the *farmB* skeleton available in the Eden library.

#### Example 6.1 (Matrix Multiplication – Hand-Parallel Program)

```

mMul xss yss
  where
    mMul xss yss      = farmB noPe f xss
                      where
                        f xs = map (dotp xs) (transpose yss)
    dotp xs ys        = foldr (+) 0 (zipWith (*) xs ys)
    transpose xss     = transpose' xss []
    transpose' [] yss = yss
    transpose' (xs : xss) yss = transpose' xss (rotate xs yss)
    rotate [] yss     = yss
    rotate (x : xs) [] = [x] : (rotate xs yss)
    rotate (x : xs) (ys : yss) = (ys ++ [x]) : (rotate xs yss)

```

Figure 4 presents the speedups achieved by the encoded parallel version of the matrix multiplication program in comparison with the original, distilled and hand-parallel versions. Since we avoid nested

parallel skeletons as explained earlier, the encoded parallel program used in this evaluation contains only the  $mMul_1''$  function defined using the *farmB* skeleton and uses the sequential definition for  $mMul_3''$ . An input size indicated by  $N \times M$  denotes the multiplication of matrices of sizes  $N \times M$  and  $M \times N$ .

When compared to the original program, we observe that the encoded parallel version achieves a positive speedup of 3x–8x for all input sizes except for 100x1000. In the case with input size 100x1000, the speedup achieved is 6x–20x more than the speedups achieved for the other input sizes. This is due to the intermediate data structure *transpose yss*, which is of the order of 1000 elements for input size 100x1000 and of the order of 100 elements for the other inputs, that is absent in the encoded parallel program. This can be verified from the comparison with the distilled version, which is also free of intermediate data structures. Hence, the encoded parallel program has a linear speedup compared to the distilled version. From Examples 6.1 and 5.1, we observe that both the hand-parallel and encoded parallel

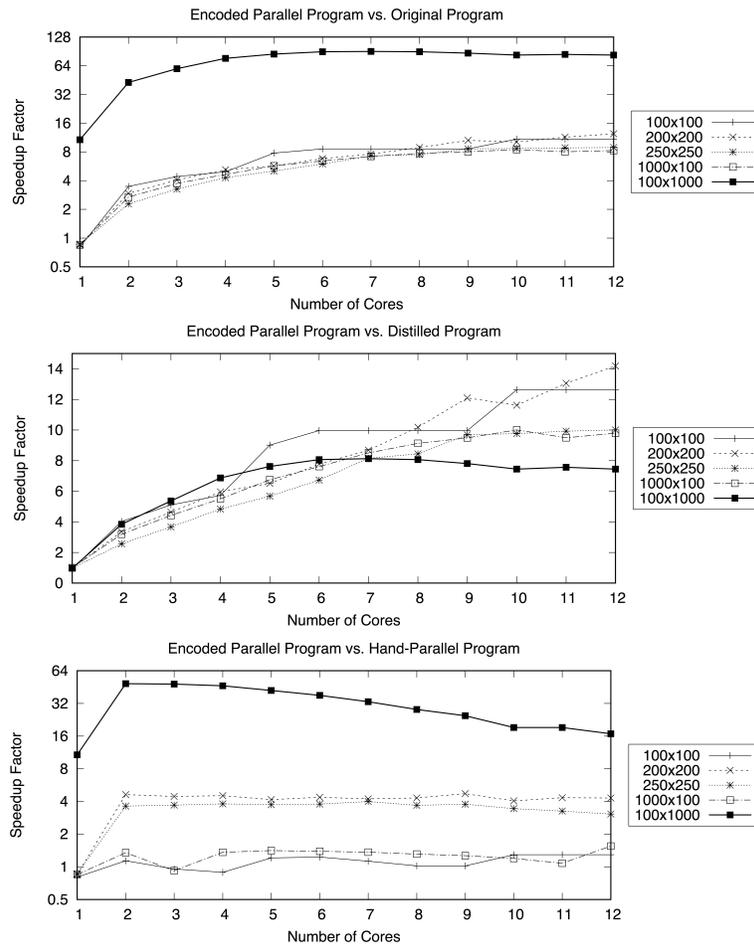


Figure 4: Evaluation of Speedup for Matrix Multiplication

versions parallelise the equivalent computations that multiply rows in the first matrix with columns in the second matrix using the *farmB* skeleton. However, the encoded parallel version is marginally faster than the hand-parallel version for input sizes 100x100 and 1000x100, and 4x faster for other input sizes except 100x1000. This is due to the use of intermediate data structures in the hand-parallel version which is of the order of 100 for all input sizes except 100x1000 for which the speedup achieved is 48x–18x

more than the hand-parallel version. Also, the hand-parallel version scales better with a higher number of cores than the encoded parallel version for the input size 100x1000. This is because the encoded parallel version achieves better speedup even with fewer cores due to the elimination of intermediate data structures, and hence does not scale as impressively as the hand-parallel version.

## 6.2 Example – Dot-Product of Binary Trees

Example 6.2 presents a sequential program to compute the dot-product of binary trees, where *dotP* computes the product of values at the corresponding branch nodes of trees *xt* and *yt*, and adds the dot-products of the left and right sub-trees. The distilled version of this program remains the same as there are no intermediate data structures and a hand-parallel version cannot be defined using list-based parallel skeletons.

### Example 6.2 (Dot-Product of Binary Trees – Original/Distilled Program)

```

data BTree a ::= E | B a (BTree a) (BTree a)
dotP :: (BTree a) → (BTree a) → (BTree a)
dotP xt yt
where
dotP E yt = 0
dotP (B x xt1 xt2) E = 0
dotP (B x xt1 xt2) (B y yt1 yt2) = (x*y) + (dotP xt1 yt1) + (dotP xt2 yt2)

```

By applying the encoding transformation, we obtain the encoded version for the dot-product program as shown in Example 6.3.

### Example 6.3 (Dot-Product of Binary Trees – Encoded Program)

```

data TdotP a ::= c1 | c2 | c3 a a (BTree a) (BTree a)
encodedotP E yt = [c1]
encodedotP (B x xt1 xt2) E = [c2]
encodedotP (B x xt1 xt2) (B y yt1 yt2) = [c3 x y xt1 yt1] ++ (encodedotP xt2 yt2)
dotP (encodedotP xt yt)
where
dotP (c1 : w) = 0
dotP (c2 : w) = 0
dotP ((c3 x y xt yt) : w) = (x*y) + (dotP (encodedotP xt yt)) + (dotP w)

```

By applying the skeleton identification rule to this encoded program, we identify that the encoded version of function *dotP* is an instance of the *mapRedr* skeleton. Example 6.4 shows the encoded parallel version defined using a suitable call to the *parMapRedr1* skeleton in the Eden library. As explained before, we define only the top-level call to *dotP''* using the parallel skeleton and use the sequential for the nested call to *dotP* because we avoid nested parallel skeletons in this evaluation.

### Example 6.4 (Dot Product of Binary Trees – Encoded Parallel Program)

```

data TdotP a ::= c1 | c2 | c3 a a (BTree a) (BTree a)
encodedotP E yt = [c1]
encodedotP (B x xt1 xt2) E = [c2]
encodedotP (B x xt1 xt2) (B y yt1 yt2) = [c3 x y xt1 yt1] ++ (encodedotP xt2 yt2)

```

$$\text{dotP}'' (\text{encode}_{\text{dotP}} \text{xt yt})$$

**where**

$$\text{dotP}'' w = \text{parMapRedr1 } g \ f \ w$$

**where**

$$g = (+)$$

$$f \ c_1 = 0$$

$$f \ c_2 = 0$$

$$f \ (c_3 \ x \ y \ \text{xt} \ \text{yt}) = (x * y) + (\text{dotP} \ \text{xt} \ \text{yt})$$

$$\text{dotP} \ E \ \text{yt} = 0$$

$$\text{dotP} \ (B \ x \ \text{xt}_1 \ \text{xt}_2) \ E = 0$$

$$\text{dotP} \ (B \ x \ \text{xt}_1 \ \text{xt}_2) \ (B \ y \ \text{yt}_1 \ \text{yt}_2) = (x * y) + (\text{dotP} \ \text{xt}_1 \ \text{yt}_1) + (\text{dotP} \ \text{xt}_2 \ \text{yt}_2)$$

Figure 5 presents the speedups of the encoded parallel version compared to the original version. An input size indicated by  $N$  denotes the dot-product of two identical balanced binary trees with  $N$  nodes each. We observe that the encoded parallel version achieves a positive speedup only upon using more than 4 cores, resulting in a maximum speedup of 2.4x for input size 1,000,000 and 1.4x for input size 100,000,000. For all input sizes, the speedup factor does not improve when using more than 6 cores. Further, we also observe that the speedup achieved scales negatively as the input size increases. The

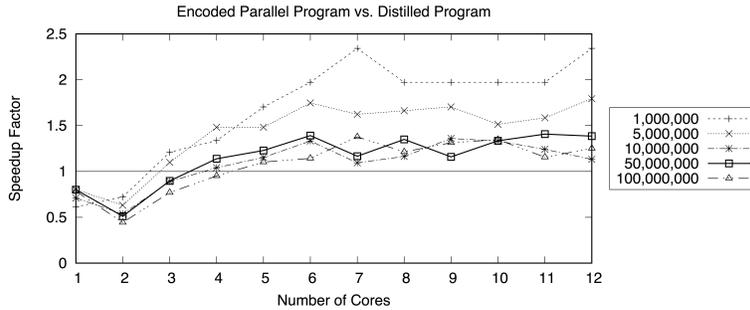


Figure 5: Evaluation of Speedup for Dot-Product of Binary Trees

reason for this performance of the encoded parallel version is as follows: From Example 6.4, we observe that each element in the encoded list contains the values at the branch nodes ( $x$  and  $y$ ) and the sub-trees ( $\text{xt}_1$  and  $\text{yt}_1$ ), which are arguments of the first recursive call  $\text{dotP} \ \text{xt}_1 \ \text{yt}_1$  in Example 6.2. Consequently, the sizes of the elements in the encoded list progressively decrease from the first to the last element if the input trees are balanced. The encoded list is then split into sub-lists in round-robin fashion by the  $\text{parMapRedr1}$  skeleton and each thread in the parallel execution applies the sequential dot-product computation  $\text{dotP}$  over a sub-list. As a result, the workloads of the parallel threads are not well-balanced and this results in significant idle times for threads that process smaller sub-lists. We also note that, left-skewed input binary trees would result in poorer performance, while right-skewed input binary trees result in better performance of the encoded parallel versions.

## 7 Conclusion

### 7.1 Summary

We have presented a transformation method to efficiently parallelise a given program by automatically identifying parallel skeletons and reducing the number of intermediate data structures used. By encoding the inputs of a program into a single input, we facilitate the identification of parallel skeletons,

particularly for map- and reduce-based computations. Additionally, we can automatically check skeleton operators for desired properties thereby allowing complete automation of the parallelisation process. Importantly, our transformation does not place any restriction on the programs that can be transformed using our method.

To evaluate our transformation method, we presented two interesting benchmark programs whose inputs are encoded into a *cons*-list. From the results, we observe two possible extreme performances. In one case, linear to super-linear speedups are achieved due to the distillation transformation, which reduces the use of intermediate data structures, as well as our parallelisation transformation. In another case, despite parallelising a program that cannot be defined using existing skeleton implementations in libraries, the positive speedups achieved are limited and not as desired. Despite not being discussed here, by employing additional skeletons such as *accumulate* [15], we are able to automatically parallelise interesting programs such as maximum prefix sum.

The primary challenge lies in the efficient execution of the parallel programs produced that are defined using skeletons. It is important to have efficient implementations of these parallel skeletons that incorporate intelligent data-partitioning and load-balancing methods across the parallel threads created to execute the skeletons. We believe that better load-balancing across threads can be facilitated by polytypic parallel skeletons, list or array data structures that support nested parallelism, or dynamic load-balancing at run-time.

## 7.2 Related Work

Previously, following the seminal works by Cole [5] and Darlington et. al. [6] on skeleton-based program development, a majority of the work that followed [20, 21, 19, 2] catered to manual parallel programming. To address the difficulties in choosing appropriate skeletons for a given algorithm, Hu et. al. proposed the *diffusion* transformation [13], which is capable of decomposing recursive functions of a certain form into several functions, each of which can be described by a skeleton. Even though diffusion can transform a wider range of functions to the required form, this method is only applicable to functions with one recursive input. Further they proposed the *accumulate* skeleton [15] that encapsulates the computational forms of map and reduce skeletons that use an accumulating parameter to build the result. However, the associative property of the reduce and scan operators used in the *accumulate* skeleton have to be verified and their unit values derived manually.

The calculational approaches to program parallelisation are based on list-homomorphisms [26] and propose systematic ways to derive parallel programs. However, most methods are restricted to programs that are defined over lists [9, 7, 10, 14]. Further, they require manual derivation of operators or their verification for certain algebraic properties to enable parallel evaluation of the programs obtained. Morihata et. al. [24] extended this approach for trees by decomposing a binary tree into a list of sub-trees called *zipper*, and defining upward and downward computations on the zipper structure. However, such calculational methods are often limited by the range of programs and data types they can transform. Also, a common aspect of these calculational approaches is the need to manually derive operators that satisfy certain properties, such as associativity to guarantee parallel evaluation. To address this, Chin et. al. [4] proposed a method that systematically derives parallel programs from sequential definitions and automatically creates auxiliary functions that can be used to define associative operators needed for parallel evaluation. However, their method is restricted to a first-order language and applicable to functions defined over a single recursive linear data type, such as lists, that has an associative decomposition operator, such as  $++$ .

As an alternative to calculational approaches, Ahn et. al. [1] proposed an analytical method to trans-

form general recursive functions into a composition of polytypic data parallel skeletons. Even though their method is applicable to a wider range of problems and does not need associative operators, the transformed programs are defined by composing skeletons and employ multiple intermediate data structures.

Previously, the authors proposed a method to transform the input of a given program into a *cons*-list based on the recursive structure of the input [16]. Since this method does not use the recursive structure of the program to build the *cons*-list, the transformed programs do not lend themselves to be defined using list-based parallel skeletons. This observation led to creating a new encoded data type that matches the algorithmic structure of the program and hence enables identification of polytypic parallel map and reduce skeletons [17]. The new encoded data type is created by pattern-matching and recursively consuming inputs, where a recursive component is created in the new encoded input for each recursive call that occurs in a function body using the input arguments of the recursive call. Consequently, the data structure of the new encoded input reflects the recursive structure of the program. Even though this method leads to better identification of polytypic skeletons, it is not easy to evaluate the performance of the transformed programs defined using these skeletons because existing libraries do not offer implementations of skeletons that are defined over a generic data type. Consequently, the proposed method of encoding the inputs into a list respects the recursive structures of programs and allows evaluation of the transformed programs using existing implementations of list-based parallel skeletons.

## Acknowledgment

This work was supported, in part, by the Science Foundation Ireland grant 10/CE/I1855 to Lero - the Irish Software Research Centre ([www.lero.ie](http://www.lero.ie)).

## References

- [1] Joonseon Ahn & Taisook Han (2001): *An Analytical Method for Parallelization of Recursive Functions*. *Parallel Processing Letters*, doi:10.1142/S0129626400000330.
- [2] Manuel M. T. Chakravarty, Gabriele Keller, Sean Lee, Trevor L. McDonell & Vinod Grover (2011): *Accelerating Haskell Array Codes with Multicore GPUs*. *Proceedings of the Sixth ACM Workshop on Declarative Aspects of Multicore Programming*, doi:10.1145/1926354.1926358.
- [3] Manuel M. T. Chakravarty, Roman Leshchinskiy, Simon Peyton Jones, Gabriele Keller & Simon Marlow (2007): *Data Parallel Haskell: A Status Report*. *Proceedings of the 2007 Workshop on Declarative Aspects of Multicore Programming (DAMP)*, doi:10.1145/1248648.1248652.
- [4] Wei-Ngan Chin, A. Takano & Zhenjiang Hu (1998): *Parallelization via Context Preservation*. *International Conference on Computer Languages*, doi:10.1109/ICCL.1998.674166.
- [5] Murray Cole (1991): *Algorithmic Skeletons: Structured Management of Parallel Computation*. MIT Press, Cambridge, MA, USA.
- [6] John Darlington, A. J. Field, Peter G. Harrison, Paul Kelly, D. W. N. Sharp, Qiang Wu & R. Lyndon While (1993): *Parallel Programming Using Skeleton Functions*. *Lecture Notes in Computer Science, 5th International PARLE Conference on Parallel Architectures and Languages Europe*, doi:10.1007/3-540-56891-3\_12.
- [7] Jeremy Gibbons (1996): *The Third Homomorphism Theorem*. *Journal of Functional Programming* Vol. 6, No. 3, doi:10.1017/S0956796800001908.
- [8] Horacio González-Vélez & Mario Leyton (2010): *A Survey of Algorithmic Skeleton Frameworks: High-level Structured Parallel Programming Enablers*. *Software – Practice and Experience*, doi:10.1002/spe.v40:12.

- [9] Sergei Gorlatch (1995): *Constructing List Homomorphisms for Parallelism*. Fakultät für Mathematik und Informatik: MIP.
- [10] Sergei Gorlatch (1999): *Extracting and Implementing List Homomorphisms In Parallel Program Development*. *Science of Computer Programming*, doi:10.1016/S0167-6423(97)00014-2.
- [11] G. W. Hamilton & Neil D. Jones (2012): *Distillation with Labelled Transition Systems*. *Proceedings of the ACM SIGPLAN Workshop on Partial Evaluation and Program Manipulation*, doi:10.1145/2103746.2103753.
- [12] Zhenjiang Hu, Masato Takeichi & Wei-Ngan Chin (1998): *Parallelization in Calculational Forms*. *Proceedings of the 25th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL)*, doi:10.1145/268946.268972.
- [13] Zhenjiang Hu, Masato Takeichi & Hideya Iwasaki (1999): *Diffusion: Calculating Efficient Parallel Programs*. *ACM SIGPLAN Workshop on Partial Evaluation and Semantics-Based Program Manipulation (PEPM)*.
- [14] Zhenjiang Hu, Tetsuo Yokoyama & Masato Takeichi (2005): *Program Optimizations and Transformations an Calculation Form*. *GTTSE*, doi:10.1007/11877028\_5.
- [15] Hideya Iwasaki & Zhenjiang Hu (2004): *A New Parallel Skeleton for General Accumulative Computations*. *International Journal of Parallel Programming*, Kluwer Academic Publishers, doi:10.1023/B:IJPP.0000038069.80050.74.
- [16] Venkatesh Kannan & G. W. Hamilton (2014): *Extracting Data Parallel Computations from Distilled Programs*. *Fourth International Valentin Turchin Workshop on Metacomputation (META)*.
- [17] Venkatesh Kannan & G. W. Hamilton (2016): *Program Transformation To Identify Parallel Skeletons*. *24th Euromicro International Conference on Parallel, Distributed, and Network-Based Processing (PDP)*, doi:10.1109/PDP.2016.32.
- [18] Rita Loogen (2012): *Eden Parallel Functional Programming with Haskell*. *Lecture Notes in Computer Science, Central European Functional Programming School, Springer Berlin Heidelberg*, doi:10.1007/978-3-642-32096-5\_4.
- [19] K. Matsuzaki, Z. Hu & M. Takeichi (2006): *Parallel Skeletons for Manipulating General Trees*. *Parallel Computing*, doi:10.1016/j.parco.2006.06.002.
- [20] K. Matsuzaki, H. Iwasaki, K. Emoto & Z. Hu (2006): *A Library of Constructive Skeletons for Sequential Style of Parallel Programming*. *Proceedings of the 1st ACM International Conference on Scalable Information Systems, InfoScale*, doi:10.1145/1146847.1146860.
- [21] K. Matsuzaki, K. Kakehi, H. Iwasaki, Z. Hu & Y. Akashi (2004): *Fusion-Embedded Skeleton Library*. *Euro-Par 2004 Parallel Processing, Lecture Notes in Computer Science, Springer Berlin Heidelberg*, doi:10.1007/978-3-540-27866-5\_85.
- [22] Kiminori Matsuzaki, Hideya Iwasaki, Kento Emoto & Zhenjiang Hu (2006): *A Library of Constructive Skeletons for Sequential Style of Parallel Programming*. *Proceedings of the 1st International Conference on Scalable Information Systems*, doi:10.1145/1146847.1146860.
- [23] Trevor L. McDonnell, Manuel M.T. Chakravarty, Gabriele Keller & Ben Lippmeier (2013): *Optimising Purely Functional GPU Programs*. *ACM SIGPLAN Notices*, doi:10.1145/2500365.2500595.
- [24] Akimasa Morihata, Kiminori Matsuzaki, Zhenjiang Hu & Masato Takeichi (2009): *The Third Homomorphism Theorem on Trees: Downward & Upward Lead to Divide-and-Conquer*. *POPL*, doi:10.1145/1594834.1480905.
- [25] Alberto Pettorossi & Maurizio Proietti (1996): *Rules and Strategies for Transforming Functional and Logic Programs*. *ACM Computing Surveys*, doi:10.1145/234528.234529.
- [26] D. B. Skillicorn (1993): *The Bird-Meertens Formalism as a Parallel Model*. *Software for Parallel Computation, NATO ASI Series F, Springer-Verlag*, doi:10.1007/978-3-642-58049-9\_9.
- [27] D. B. Skillicorn & Domenico D. Talia (1998): *Models and Languages for Parallel Computation*. *ACM Computing Surveys*, doi:10.1145/280277.280278.