

A Case Study on Proving Transformations Correct: Data-Parallel Conversion*

Stephen Fitzpatrick[†], M. Clint[‡] and P. Kilpatrick[‡]

[†] Kestrel Institute
3260 Hillview Avenue
Palo Alto
California 94304
U.S.A.

[‡] Department of Computer Science
The Queen's University of Belfast,
Belfast BT7 1NN
Northern Ireland

Fitzpatrick@kestrel.edu, M.Clint@qub.ac.uk, P.Kilpatrick@qub.ac.uk

Abstract

The issue of correctness in the context of a certain style of program transformation is investigated. This style is characterised by the fully automated application of large numbers of simple transformation rules to a representation of a functional program (serving as a specification) to produce an equivalent efficient imperative program. The simplicity of the transformation rules ensures that the proofs of their correctness are straightforward.

A selection of transformations appropriate for use in a particular context are shown to preserve program meaning. The transformations convert array operations expressed as the application of a small number of general-purpose functions into applications of a large number of functions which are amenable to efficient implementation on an array processor.

Keywords: Correctness proofs, program transformation, functional programming, array processor

1 Introduction

In [1, 2], and elsewhere, a style of program derivation is advocated in which a complex change is wrought on a program through a number of simpler changes brought about by the automated application of a sequence of sets of simple transformations. For example, in [3] an imperative, explicitly data-parallel program is derived from a pure, functional program (which is used as an abstract specification) in six main steps.

SML → λ -calculus → Unfolded → Simplified
→ Array Form → Common Sub-expressions Eliminated
→ Fortran Plus Enhanced (DAP)

The functional program (expressed in SML [4]) is translated into the λ -calculus; definitions are unfolded; expressions are simplified using algebraic rules; array expressions are translated into explicitly data-parallel (but still pure) forms; expressions are optimized; state and imperative control constructs are introduced to produce an imperative program (expressed in a variant of Fortran for execution on the AMT DAP array processor [5]).

One of the advantages claimed for this style of derivation is that establishing that program correctness is preserved at all stages is simplified by (i) the representation of a complex change as a sequence of

*The work reported in this paper was supported by SERC grant GR/G 57970.

simpler, conceptually independent changes; (ii) the simplicity of the individual transformations; and (iii) regarding transformations as abstract rewrite rules rather than concentrating on their concrete effect on given programs.

This paper presents evidence to support this claim. It contains, for the stage in the derivation that produces the explicitly data-parallel form, a proof of correctness for each of the more complex transformations and a sample of proofs of correctness for the simple transformations. This stage is conceptually and technically significant. It is the first of the two pivotal stages in the derivation, the other being the translation from functional to imperative form. Note that the proofs are of themselves not of great interest — none of them contains novel proof techniques. Indeed, the interest lies in the fact that the proofs *are* mundane.

The structure of the paper is as follows: the basic notation used is explained; the functions used in the initial form and in the Array Form are defined; useful lemmas relating to the functions are stated (their proofs are relegated to the appendix); the transformations are defined; and their correctness is proved.

2 Basic Notation

The basic notation used in this paper is that of the λ -calculus with optional type information and a set of primitive functions. An array is considered to be a mapping from a (finite) set of indices onto some range of values of type α . A 1-dimensional array may be referred to as a *vector* and a 2-dimensional array as a *matrix*. For a set, the operator $+$ indicates element insertion: $S+i \equiv S \cup \{i\}$.

In the following sections, four *primitive* array functions and the Array Form functions are defined. The primitive functions are taken to define the semantics of arrays. Most array operations commonly encountered in numerical mathematics can be compactly expressed using these primitives. However, it is not intended that a programmer be restricted to using only these primitives: other, perhaps more convenient, functions can be defined in terms of these primitive functions. Transformations can be employed to eliminate such *derived* functions using techniques such as unfolding and algebraic simplification.

The Array Form functions are much more restricted than the primitive functions, but are also much simpler to implement efficiently on an array processor such as the AMT DAP. In effect, the Array Form may be considered as a functional abstraction of an array processor.

2.1 Primitive Array Functions

The four primitive array functions are: `shape`, `element`, `generate` and `reduce`.

`shape(A: α array) \rightarrow Shape`

Given an array, the function `shape` can be used to obtain its index set. The extent of an array in a particular dimension can be obtained by specifying the dimension:

$$\text{shape}(A:\alpha \text{ array}, n:\text{int}) \rightarrow \text{int}$$

`element(A: α array, i:index) \rightarrow α`

The `element` function returns the value of the element of `A` at position `i`. For convenience, the infix operator `@` is defined to be equivalent to `element`:

$$A@i \equiv \text{element}(A, i)$$

`generate(S:Shape, g:index \rightarrow α) \rightarrow α array`

The basic function for constructing arrays is `generate`. The first argument, `S`, specifies the index set of the constructed array. The second argument, `g`, is a function, called the *generating function*, which determines the values of the elements: the value of element `i` is `g(i)`.

The following are some examples of arrays constructed using `generate`:

- the elementwise addition of two arrays, of arbitrary dimensionality, having the same shape:

`generate(shape(A), λi·A@i+B@i)`

- the transpose of a 2-dimensional array A of shape [m, n]:

`generate([n, m], λ[i, j]·A@[j, i])`

An argument, such as *i*, of a generating function is called a *generating index*. For multi-dimensional arrays, the term may also be used of the *components* of an index argument; for example, in `λ[i, j]·e`, the generating indices are *i* and *j*. It should be clear from context whether a whole index or a component is being considered.

`reduce(r:α × α → α, r0:α, S:Shape, g:index → α) → α`

Many array operations require the elements of an array to be accumulated — or *reduced* — into a single value by the repeated application of a binary *reducing function*. For example, the sum of the elements of a numeric array is a reduction using the addition function. Reductions are denoted using the `reduce` function.

The argument *r* is the reducing function. The argument *r0* is the *initial value* which is used to instantiate the accumulation (it is usually an identity of the reducing function and so does not alter the value of the reduction; its inclusion helps simplify the semantics of reductions by ensuring that a reduction is well defined even if an array contains only one element, or even no elements).

The arguments *S* and *g* (*g* is a generating function) can be used to specify the elements of the array which are to be reduced. For example, `reduce(+, 0, shape(A), λi·A@i)` produces the sum of the elements of array *A*. However, the generating function need not *necessarily* be an application of `element`: a reduction can involve any set of values which can be specified by applying a generating function over an index set. For example, the inner-product of two vectors *U* and *V*, of shape [n], can be expressed as `reduce(+, 0, [n], λi·U@i * V@i)`

The four functions `shape`, `element`, `generate` and `reduce` are the basic array functions; most common vector and matrix operations can be readily expressed using them. Some further examples are given below:

- Row *i* of a matrix *A*: `generate(shape(A, 2), λ[j]·A@[i, j])`
- Product of two matrices *A* and *B* (which are assumed to be conformant; i.e. the number of columns of *A* equals the number of rows of *B*):

`generate([shape(A, 1), shape(B, 2)],
λ[i, j]·reduce(+, 0, shape(A, 2), λ[k]·A@[i, k]*B@[k, j]))`

- Boolean matrix, of shape [n, n], having leading diagonal elements `true`, and all other elements `false`:

`generate([n, n], λ[i, j]·i=j)`

2.1.1 Formal Definition of generate and reduce

Definition 1: generate

The `generate` function is defined by the two identities

$$\text{shape}(\text{generate}(S, \lambda i \cdot g)) \equiv S \quad (\text{G1})$$

$$\forall i' \in S: \text{element}(\text{generate}(S, \lambda i \cdot g), i') \equiv \lambda i \cdot g(i') \quad (\text{G2})$$

□

That is, the shape of an array constructed by an application of `generate` is the shape specified by the shape argument; and the value of each element of the constructed array is found by applying the generating function to the element's index.

When the shape of a generation is a manifest list, axiom G1 may be used in the form

$$\text{shape}(\text{generate}([m, n], \lambda i \cdot g), 1) \equiv m$$

In the proofs presented in this paper, the condition $i' \in S$ in axiom G2 — that an index be a member of an array's shape — is usually ignored. It could be verified separately from the main proofs, or it could be included in the form

$$\text{element}(\text{generate}(S, \lambda i.g), i') \equiv \text{if } (i' \in S) \text{ then } \lambda i.g(i') \text{ else } \perp \quad (\text{G2}')$$

where \perp is the undefined value, bottom. Context could then be used to establish that the condition $i' \in S$ is true, and so the conditional expression can be reduced into just the true limb. This technique is illustrated in one proof (of lemma 5).

Definition 2: reduce

The **reduce** function is defined recursively on the index set over which the reduction is to be performed:

$$\begin{aligned} \text{reduce}(r, r0, \emptyset, \lambda i.g) &\equiv r0 && (\text{R1}) \\ \text{reduce}(r, r0, S+i', \lambda i.g) &\equiv r(\lambda i.g(i'), \text{reduce}(r, r0, S, \lambda i.g)) && (\text{R2}) \end{aligned}$$

where $i' \notin S$ and \emptyset denotes the empty set (of indices)

□

Note that no order is defined for performing reductions, so reducing functions must be associative and commutative.

2.2 Array Form Functions

The main Array Form functions are now defined. These functions are intended to capture the sorts of operations that any array processor could be expected to implement efficiently (for example, simultaneously adding corresponding elements of two arrays). Some of the functions, and some of the transformations discussed later, are perhaps peculiar to array processors in which the processing elements are arranged in a *two-dimensional* array (the AMT DAP is one such processor). It should be emphasized that what is under consideration in this paper is the correctness of the transformations that create the Array Form, rather than how well suited the Array Form is for use on a particular processor.

2.2.1 Main Array Form Functions

See Figure 1 for examples. Probably the most important functions of the Array Form are the *mapping* functions, which apply a scalar function to each element of an array, or to corresponding elements of a pair of arrays. Mappings are supported for only certain scalar functions — for example, the basic arithmetic and logical functions.

Definition 3: map

$$\begin{aligned} \text{map}(A:\alpha \text{ array}, f:\alpha \rightarrow \beta) &\rightarrow \beta \text{ array} \\ \stackrel{\text{def}}{=} \text{generate}(\text{shape}(A), \lambda i.f(A@i)) \end{aligned}$$

$$\begin{aligned} \text{map}(A:\alpha \text{ array}, B:\beta \text{ array}, f:\alpha \times \beta \rightarrow \gamma) &\rightarrow \gamma \text{ array} \\ \stackrel{\text{def}}{=} \text{generate}(\text{shape}(A), \lambda i.f(A@i, B@i)) \end{aligned}$$

where A and B have the same shape

□

The **fold** function performs restricted forms of reductions, in which the values reduced are elements of an array and where the reducing function is one of a limited set (and typically evaluating sum, product, logical and, logical or, minimum or maximum). A variant of **fold** is also defined which reduces a matrix along its rows, thereby forming a vector of (partial) cumulative values — this corresponds to operations on the DAP which compute multiple vector reductions simultaneously.

Definition 4: fold

$$\begin{aligned} \text{fold}(r:\alpha \times \alpha \rightarrow \alpha, r0:\alpha, A:\alpha \text{ array}) &\rightarrow \alpha \\ \stackrel{\text{def}}{=} \text{reduce}(r, r0, \text{shape}(A), \lambda i.A@i) \end{aligned}$$

□

Definition 5: fold.rows

$$\text{map} \left(\left[\begin{array}{ccc} 1 & 2 & 3 \\ 4 & 5 & 6 \\ 7 & 8 & 9 \end{array} \right], \left[\begin{array}{ccc} 1 & 2 & 3 \\ 4 & 5 & 6 \\ 7 & 8 & 9 \end{array} \right], + \right) = \left[\begin{array}{ccc} 2 & 4 & 6 \\ 8 & 10 & 12 \\ 14 & 16 & 18 \end{array} \right]$$

$$\text{fold} \left(+, 0, \left[\begin{array}{ccc} 1 & 2 & 3 \\ 4 & 5 & 6 \\ 7 & 8 & 9 \end{array} \right] \right) = 45$$

$$\text{fold.rows} \left(+, \left[\begin{array}{c} 0 \\ 0 \\ 0 \end{array} \right], \left[\begin{array}{ccc} 1 & 2 & 3 \\ 4 & 5 & 6 \\ 7 & 8 & 9 \end{array} \right] \right) = \left[\begin{array}{c} 6 \\ 15 \\ 24 \end{array} \right]$$

$$\text{join} \left(\left[\begin{array}{ccc} T & T & F \\ F & F & F \\ T & F & T \end{array} \right], \left[\begin{array}{ccc} 11 & 12 & 13 \\ 14 & 15 & 16 \\ 17 & 18 & 19 \end{array} \right], \left[\begin{array}{ccc} 21 & 22 & 23 \\ 24 & 25 & 26 \\ 27 & 28 & 29 \end{array} \right], \right) = \left[\begin{array}{ccc} 11 & 12 & 23 \\ 24 & 25 & 26 \\ 17 & 28 & 19 \end{array} \right]$$

Figure 1: Examples of Array Form functions

$$\text{fold.rows}(r:\alpha \times \alpha \rightarrow \alpha, R0:\alpha \text{ vector}, A:\alpha \text{ matrix}) \rightarrow \alpha \text{ vector}$$

$$\stackrel{\text{def}}{=} \text{generate}([\text{shape}(A, 1)], \lambda[i].\text{reduce}(r, R0@[i], [\text{shape}(A, 2)], \lambda[j].A@[i, j]))$$

□

The join function is a data-parallel conditional. The result of applying join is an array with elements merged from two arrays according to whether the corresponding element of a *mask* array is true or false.

Definition 6: join

$$\text{join}(M:\text{boolean array}, T:\alpha \text{ array}, F:\alpha \text{ array}) \rightarrow \alpha \text{ array}$$

$$\stackrel{\text{def}}{=} \text{generate}(\text{shape}(M), \lambda i.\text{if } M@i \text{ then } T@i \text{ else } F@i)$$

where M, T and F have the same Shape

□

2.2.2 Miscellaneous Functions

The Array Form defines many functions for performing miscellaneous array operations such as transposing a matrix, ‘shifting’ the elements of a matrix in a specified direction and constructing logical matrices having true values in certain patterns (such as along their main diagonals or in their upper triangles). Here, only a few examples of such functions are considered, since the correctness of transformations involving such functions is usually trivial to establish from the function definitions.

Definition 7: matrix.transpose

$$\text{matrix.transpose}(A:\alpha \text{ array}) \rightarrow \alpha \text{ array}$$

$$\stackrel{\text{def}}{=} \text{generate}([\text{shape}(A, 2), \text{shape}(A, 1)], \lambda[i, j].A@[j, i])$$

□

Definition 8: row

$$\text{row}(A:\alpha \text{ array}, i:\text{int}) \rightarrow \alpha \text{ array}$$

$$\stackrel{\text{def}}{=} \text{generate}(\text{shape}(A, 2), \lambda[j].A@[i, j])$$

□

3 Preliminary Results

Some properties of `generate` and `reduce` are presented below — proofs of these properties are presented in the appendix. In addition, some elementary identities of the λ -calculus are noted (without proof).

In the following, if an identifier is introduced on the right of an identity, then it should be assumed that it is a ‘new’ identifier, i.e. one that does not occur free in the expression on the left. For example, in the identity

$$B \equiv \lambda x. B[e \rightarrow x] (e)$$

it is to be assumed that x does not occur free in B , so that no problem arises with name clashes.

Lemma 1: Identity λ -bindings

A λ -binding in which the bound identifier and the bound value are the same is redundant.

$$\lambda x. B (x) \equiv B$$

□

Lemma 2: Propagation of λ -binding out of abstraction

A λ -binding can be moved out of an immediately enclosing λ -abstraction if the bound value does not depend on the identifier of the abstraction.

$$\lambda i. (\lambda x. B (e)) \equiv \lambda x. (\lambda i. B) (e)$$

where e does not contain i

(Note that x is still bound to e .)

□

Lemma 3: Propagation of λ -binding through a function application

An applied λ -binding that is an argument in a function application can be moved outside the function application.

$$f(\lambda x. B (e)) \equiv \lambda x. f(B) (e)$$

where f does not contain x

□

This identity can be generalised for moving a λ -binding out of an argument in a function application which has more than one argument, provided all of the other arguments are free of the bound identifier. Note that this lemma can be applied in both directions, to move a λ -binding into, as well as out of, an argument position.

3.1 Properties of Elementwise Applications

The following lemmas pertain to elementwise applications of functions. Such applications can be denoted using the `map` function but, for convenience in later proofs, the operator ϵ is used to promote a binary function to a binary function on arrays. For example, $\epsilon(+)$ is a function that performs elementwise addition of two arrays.

Definition 9: Elementwise Operator, ϵ

$$\begin{aligned} & \epsilon(f: \alpha \times \beta \rightarrow \gamma) \rightarrow (\alpha \text{ array} \times \beta \text{ array} \rightarrow \gamma \text{ array}) \\ \stackrel{\text{def}}{=} & \lambda X: \alpha \text{ array}, Y: \beta \text{ array}. \text{generate}(\text{shape}(Y), \lambda i. f(X@i, Y@i)) \end{aligned}$$

□

Lemma 4: Shape of an elementwise application

The shape of an application of an elementwise function to two arrays is the same as the shape of the second argument array (which is required to be of the same shape as the first argument array).

$$\text{shape}(\epsilon(f)(A, B)) \equiv \text{shape}(B)$$

□

Lemma 5: Element of an elementwise application

An element of an elementwise application of a function to two arrays is the value of that function applied to the corresponding elements of the arrays; that is, `element` propagates through ϵ :

$$\begin{aligned} \text{element}(\epsilon(f)(A, B), i) &\equiv f(\text{element}(A, i), \text{element}(B, i)) \equiv f(A@i, B@i) \\ &\text{for } i \in \text{shape}(\epsilon(f)(A, B)) \end{aligned}$$

□

3.2 Properties of Reductions

If the reducing function of a reduction is an elementwise function, then the result of the reduction is an array, so it is valid to apply the `shape` and `element` functions to the result. Such a reduction is called an ϵ -reduction:

Definition 10: ϵ -reduction

A reduction of the form

$$\text{reduce}(\epsilon(r), R0, S, \lambda i \cdot g)$$

is called an ϵ -reduction.

□

The following two lemmas pertain to ϵ -reductions.

Lemma 6: Shape of an ϵ -reduction

The result of an ϵ -reduction has the same shape as the reduction's initial value:

$$\text{shape}(\text{reduce}(\epsilon(r), R0, S, \lambda i \cdot g)) \equiv \text{shape}(R0)$$

□

Lemma 7: Element of an ϵ -reduction

An element of a reduction using $\epsilon(r)$ is a reduction using r : that is, `element` can be propagated through an ϵ -reduction.

$$\begin{aligned} &\text{element}(\text{reduce}(\epsilon(r), R0, S, \lambda i \cdot g), j) \\ &\equiv \text{reduce}(r, \text{element}(R0, j), S, \lambda i \cdot \text{element}(g, j)) \\ &\text{where } S \text{ is independent of } j \end{aligned}$$

□

Lemma 8: Reduction over a union

Since no order is specified for performing reductions (and since reducing functions are required to be associative and commutative), a reduction over an index set that is the union of two sets can be split into a pair of reductions.

$$\begin{aligned} \text{reduce}(r, r0, S \cup T, \lambda i \cdot e) &\equiv r(\text{reduce}(r, r0, S, \lambda i \cdot e), \text{reduce}(r, r0, T, \lambda i \cdot e)) \\ &\text{where } r0 \text{ is an identity element of } r, \text{ and } S \text{ and } T \text{ are disjoint} \end{aligned}$$

□

Lemma 9: Collapsing singleton dimensions

If one of the dimensions of a multi-dimensional reduction contains only a single member, that dimension can be collapsed — that is, removed from the reduction's index set. For this paper, collapsing is required for only leading dimensions (i is a leading dimension in $-i' \times S$).

$$\begin{aligned} \text{reduce}(r, r0, -i' \times S, \lambda ij \cdot e) &\equiv \text{reduce}(r, r0, S, \lambda j \cdot (\lambda i \cdot e(i'))) \\ &\text{where } r0 \text{ is an identity element of } r \\ &\text{and where } i' \text{ and } i \text{ have the same dimensionality} \end{aligned}$$

□

4 The Transformations

In this section, the main transformations (actually, identities) for converting from expressions that use the basic array functions (`generate`, `reduce`, etc.) into expressions that use Array Form functions are listed; their correctness is established in the following section. The basic strategy for converting into Array Form is to propagate applications of `generate` into generating functions; for example, a generation of a suitable scalar function becomes an application of `map`; and a generation of a conditional expression becomes an application of `join` — this strategy is discussed at length in [3]. In addition, several transformations optimize combinations of operations to make best use of the DAP hardware.

4.1 General Transformations

Constructing an array by applying a scalar function to the elements of an array (or to corresponding elements of a pair of arrays) is equivalent to applying the elementwise version of the function (expressed using `map`).

Transformation 1: Propagation through scalar functions

$$\begin{aligned}\text{generate}(S, \lambda i \cdot f(a)) &\equiv \text{map}(\text{generate}(S, \lambda i \cdot a), f) \\ \text{generate}(S, \lambda i \cdot f(a, b)) &\equiv \text{map}(\text{generate}(S, \lambda i \cdot a), \text{generate}(S, \lambda i \cdot b), f) \\ &\text{where } f \text{ is a scalar function for which elementwise application is supported by } \text{map}\end{aligned}$$

□

(The function f is necessarily independent of i .)

Constructing an array by evaluating a conditional expression for each element is equivalent to forming a data-parallel conditional, in which arrays are constructed from the true and false limbs of the conditional and are merged according to a mask generated from the predicate.

Transformation 2: Propagation through conditional

$$\begin{aligned}\text{generate}(S, \lambda i \cdot \text{if } p \text{ then } t \text{ else } f) \\ \equiv \text{join}(\text{generate}(S, \lambda i \cdot p), \text{generate}(S, \lambda i \cdot t), \text{generate}(S, \lambda i \cdot f))\end{aligned}$$

□

As mentioned previously, only a few examples of miscellaneous Array Form functions will be considered.

Transformation 3: `matrix.transpose`

$$\begin{aligned}\text{generate}([m, n], \lambda [i, j] \cdot A@[j, i]) &\equiv \text{matrix.transpose}(A) \\ &\text{where } A \text{ has shape } [n, m]\end{aligned}$$

□

Transformation 4: `row`

$$\begin{aligned}\text{generate}([m], \lambda [j] \cdot A@[r, j]) &\equiv \text{row}(A, r) \\ &\text{where } A \text{ has shape } [n, m] \text{ and } A \text{ and } r \text{ are independent of } j\end{aligned}$$

□

4.2 Propagation through λ -expressions

Consider a generation in which the body of the generating function is a λ -binding:

$$\text{generate}(S, \lambda i \cdot (\lambda x \cdot B(e)))$$

The task of the Array Form transformations is to convert this generation into a form that can be efficiently implemented on an array processor: in the general case, all of the bindings are evaluated in parallel, then all of the body expressions are evaluated in parallel:

$\forall i \in S$: evaluate e ;

$\forall i \in S$: evaluate B

If parallelism of unlimited dimensionality were permitted, it would be a simple matter to create this parallel form. However, because the DAP is limited to 2-dimensional parallelism, it is incapable of efficiently implementing the general case in the above manner. For example, if S were 2-dimensional and e 1-dimensional, then the above scheme would require the creation of a 2-dimensional array of 1-dimensional arrays, a structure which the DAP can manipulate, *but not in a completely parallel manner*.

Nevertheless, the above scheme can be used for certain cases where the *effective* or useful parallelism is at most 2-dimensional: for example, if S is 1- or 2-dimensional and e is a scalar value; or if S is 2-dimensional and e is a vector which is independent of one of the dimensions of S . The transformations below pertain to such cases; if none of these transformations applies to a given binding, then, as a last resort, the binding can be β -reduced in the hope that the resulting expression can be parallelised. In the following, it is assumed that all shapes in generations and reductions are at most 2-dimensional.

If the bound value is independent of the generating index, then the generation can be propagated into the binding.

Transformation 5: Invariant binding

$$\text{generate}(S, \lambda i. \lambda x. B(e)) \equiv \lambda x. \text{generate}(S, \lambda i. B(e))$$

where e is independent of i and S is independent of x

□

If the bound value is a scalar, then the generation can be propagated into the binding by creating an array of bound values.

Transformation 6: Scalar binding

$$\begin{aligned} &\text{generate}(S, \lambda i. (\lambda x. B(e))) \\ &\equiv \lambda X. \text{generate}(S, \lambda i. (\lambda x. B(X@[i]))) (\text{generate}(S, \lambda i. e)) \end{aligned}$$

where e is a scalar expression

□

The binding for x is removed by β -reduction after application of this transformation.

Suppose that a 2-dimensional array of shape $[l, m]$ is being generated over indices i and j , and that for each element a λ -binding is formed in which the bound value is a vector generation over index k . In the general case, $l \times m$ vectors must be created. However, if the bound vector is independent of j , then only l vectors need be created (one for each value of i). These l vectors can be created simultaneously as the rows of a matrix generation over indices i and k .

Transformation 7: Matrix generation, vector binding independent of one generating index

$$\begin{aligned} &\text{generate}([l, m], \lambda [i, j]. (\lambda x. B(\text{generate}([n], \lambda [k]. e)))) \\ &\equiv \lambda X. \text{generate}([l, m], \lambda [i, j]. (\lambda x. B(\text{generate}([n], \lambda [k]. X@[i, k]))) \\ &\quad (\text{generate}([l, n], \lambda [i, k]. e)) \end{aligned}$$

where e is a function of i and k but not of j , and is not an application of element

□

As it stands, the expression produced by this transformation does not appear to be an improvement over the initial expression, as the generation still contains a vector binding. However, it is usually the case that B requires only individual elements of x , and not the vector as a whole.¹ Then the binding for x can be reduced and parallelisation of the resulting expression can proceed.

A similar transformation can be applied when e is independent of i rather than j . (It is assumed that the case of e independent of k is optimized by converting the vector binding into a scalar binding.)

If the array being constructed is a matrix and the bound value is a vector which is dependent on both matrix indices, and if the generating function of the vector is a function application, then the function application can be moved outside the vector. This transformation is expressed below for arrays of arbitrary dimensionality, but it is applied in practice only to a matrix generation/vector binding combination.

¹Function unfolding and algebraic simplification normally result in array expressions occurring solely as arguments of element. The exceptions are arrays which occur as arguments in or results of applications of recursive functions.

Transformation 8: Binding is a function application

$$\begin{aligned} & \text{generate}(\mathbf{S}, \lambda i.(\lambda x.B(\text{generate}(\mathbf{T}, \lambda j.f(a)))))) \\ \equiv & \text{generate}(\mathbf{S}, \lambda i.(\lambda x'.(\lambda x.B(\text{generate}(\mathbf{T}, \lambda j.f(x'@j))))))(\text{generate}(\mathbf{T}, \lambda j.a)) \end{aligned}$$

□

As discussed above, it may be possible to remove the binding for x if only individual elements of x are required, and not the entire array. The repeated application of this transformation may reduce the binding into a form that can be parallelised by one of the preceding transformations. A similar transformation can be used if f is a binary function.

4.3 Transformations for Reductions

Consider a generation having a generating function that is a reduction:

$$\text{generate}(\mathbf{S}, \lambda i.\text{reduce}(r, r0, \mathbf{T}, \lambda j.g))$$

There are several ways that this expression could be converted into Array Form:

- Each of the reductions can be parallelised: that is, i is iterated over sequentially, and for each i , a parallel reduction is performed.
- The generation can be parallelised by exchanging the **generate** and the **reduce**; then j is iterated over sequentially, while for each j , the generation is evaluated in parallel.
- The generation and the reduction can be combined into a partial reduction (such as **fold.rows**), so that both are evaluated in parallel.

The third option, combination, is generally preferable when it is feasible, since it makes maximum use of parallelism. However, on a computer such as the DAP, which is limited to 2-dimensional parallelism, combination is possible only when both \mathbf{S} and \mathbf{T} are 1-dimensional.

Failing the third option, the second option is preferable, since generations generally make better use of parallelism than reductions. (For example, two arrays of arbitrary size can, theoretically, be added in a single step, whereas the reduction of an array of size $[n]$ requires $\log_2(n)$ steps.)²

The following two transformations enforce these preferences.

Transformation 9: generate-reduce combination

$$\begin{aligned} & \text{generate}([m], \lambda[i].\text{reduce}(r, r0, [n], \lambda[j].e)) \\ \equiv & \text{fold.rows}(r, \text{generate}([m], \lambda[i].r0), \text{generate}([m, n], \lambda[i, j].e)) \\ & \text{where } n \text{ and } r \text{ are independent of } i \end{aligned}$$

□

Transformation 10: generate-reduce swap

$$\begin{aligned} & \text{generate}(\mathbf{S}, \lambda i.\text{reduce}(r, r0, \mathbf{T}, \lambda j.e)) \\ \equiv & \text{reduce}(\epsilon(r), \text{generate}(\mathbf{S}, \lambda i.r0), \mathbf{T}, \lambda j.\text{generate}(\mathbf{S}, \lambda i.e)) \\ & \text{where } r, r0 \text{ and } \mathbf{T} \text{ are independent of } i \text{ and } \mathbf{S} \text{ is independent of } j \end{aligned}$$

□

This transformation can be applied for arrays of arbitrary dimensionality; for the DAP, however, it is used only for matrix generation and vector reduction.

If each component of a reduction is itself a reduction (which uses the same reducing function), then coalescing the reductions into a single reduction increases parallelism.

²Exchanging a matrix generation and a vector reduction is equivalent to the well known optimization for matrix product, in which the 'ijk' order (k parallel) is converted into the 'kij' order (ij parallel).

Transformation 11: reduce-reduce combination

$$\text{reduce}(r, r0, S, \lambda i.\text{reduce}(r, r0, T, \lambda j.e)) \equiv \text{reduce}(r, r0, S \times T, \lambda ij.e)$$

where $r0$ is an identity element of r , and r and T are independent of i

□

5 Proofs of Correctness

The correctness of the transformations discussed in the preceding section is established below. The proofs are presented in the same order as the transformations.

Proof 1: Propagation through scalar functions

Consider the case of a binary function.

$$\text{generate}(S, \lambda i.f(a, b)) \equiv \text{map}(\text{generate}(S, \lambda i.a), \text{generate}(S, \lambda i.b), f)$$

Proof follows directly from the definition of `map`:

$$\begin{aligned} & \text{map}(\text{generate}(S, \lambda i.a), \text{generate}(S, \lambda i.b), f) \\ &= \text{definition 3} \\ & \text{generate}(\text{shape}(\text{generate}(S, \lambda i.a)), \\ & \quad \lambda i.f(\text{element}(\text{generate}(S, \lambda i.a), i), \text{element}(\text{generate}(S, \lambda i.b), i))) \\ &= \text{by G1 and G2} \\ & \text{generate}(S, \lambda i.f(\lambda i.a(i), \lambda i.b(i))) \\ &= \text{remove identity bindings (lemma 1)} \\ & \text{generate}(S, \lambda i.f(a, b)) \end{aligned}$$

□

A similar proof applies for unary functions.

Note that the proof is simplified by choosing the appropriate generating index when the generation is introduced. Any other generating index, say j , could be used and would lead to an expression such as

$$\text{generate}(S, \lambda j.f(\lambda i.a(j), \lambda i.b(j)))$$

which, since f is independent of j , is equivalent to

$$\text{generate}(S, \lambda j.(\lambda i.f(a, b)(j)))$$

which is one way of expressing the process of α -converting a λ -abstraction from using identifier i to using identifier j ; that is, this expression is equivalent under α -conversion to $\text{generate}(S, \lambda i.f(a, b))$.

Proof 2: Propagation through conditional

$$\begin{aligned} & \text{generate}(S, \lambda i.\text{if } p \text{ then } t \text{ else } f) \\ & \equiv \text{join}(\text{generate}(S, \lambda i.p), \text{generate}(S, \lambda i.t), \text{generate}(S, \lambda i.f)) \end{aligned}$$

Proof follows directly from the definition of `join`:

$$\begin{aligned} & \text{join}(\text{generate}(S, \lambda i.p), \text{generate}(S, \lambda i.t), \text{generate}(S, \lambda i.f)) \\ &= \text{definition 6} \\ & \text{generate}(\text{shape}(\text{generate}(S, \lambda i.p)), \\ & \quad \lambda i.\text{if } \text{element}(\text{generate}(S, \lambda i.p), i) \\ & \quad \text{then } \text{element}(\text{generate}(S, \lambda i.t), i) \\ & \quad \text{else } \text{element}(\text{generate}(S, \lambda i.f), i)) \\ &= \text{by G1 and G2} \\ & \text{generate}(S, \lambda i.\text{if } \lambda i.p(i) \text{ then } \lambda i.t(i) \text{ else } \lambda i.f(i)) \\ &= \text{remove identity bindings (lemma 1)} \\ & \text{generate}(S, \lambda i.\text{if } p \text{ then } t \text{ else } f) \end{aligned}$$

□

Proof 3: matrix.transpose

$$\text{generate}([m, n], \lambda[i, j].A@[j, i]) \equiv \text{matrix.transpose}(A)$$

where A has shape [n, m]

Proof follows directly from the definition of matrix.transpose:

$$\begin{aligned} & \text{matrix.transpose}(A) \\ &= \textit{definition 7} \\ & \text{generate}([\text{shape}(A, 2), \text{shape}(A, 1)], \lambda[i, j].A@[j, i]) \\ &= \textit{substituting for the shape of A} \\ & \text{generate}([m, n], \lambda[i, j].A@[j, i]) \end{aligned}$$

□

Proof 4: row

$$\text{generate}([m], \lambda[j].A@[r, j]) \equiv \text{row}(A, r)$$

where A and r are independent of i and A has shape [n, m]

Proof follows directly from the definition of row:

$$\begin{aligned} & \text{row}(A, r) \\ &= \textit{definition 8} \\ & \text{generate}(\text{shape}(A, 2), \lambda[j].A@[r, j]) \\ &= \textit{substituting for the shape of A} \\ & \text{generate}([m], \lambda[j].A@[r, j]) \end{aligned}$$

□

5.1 Proofs of Transformations for λ -bindings

The following proofs pertain to transformations for generating functions whose bodies are λ -bindings.

Proof 5: Invariant binding

$$\text{generate}(S, \lambda i.(\lambda x.B(e))) \equiv \lambda x.\text{generate}(S, \lambda i.B(e))$$

where e is independent of i and S is independent of x

Proof involves only elementary properties of the λ -calculus:

$$\begin{aligned} & \text{generate}(S, \lambda i.(\lambda x.B(e))) \\ &= \textit{move binding of x out of abstraction of i (lemma 2)} \\ & \text{generate}(S, \lambda x.(\lambda i.B(e))) \\ &= \textit{move binding out of application of generate (lemma 3)} \\ & \lambda x.\text{generate}(S, \lambda i.B(e)) \end{aligned}$$

□

Proof 6: Scalar binding

$$\begin{aligned} & \text{generate}(S, \lambda i.(\lambda x.B(e))) \\ & \equiv \lambda X.\text{generate}(S, \lambda i.(\lambda x.B(X@[i]))) (\text{generate}(S, \lambda i.e)) \\ & \textit{where e is a scalar} \end{aligned}$$

Proof follows by β -reducing X:

$$\begin{aligned} & \lambda X.\text{generate}(S, \lambda i.(\lambda x.B(X@[i]))) (\text{generate}(S, \lambda i.e)) \\ &= \textit{\beta-reduce X; only occurrence of X is that shown} \\ & \text{generate}(S, \lambda i.(\lambda x.B(\text{element}(\text{generate}(S, \lambda i.e), i)))) \\ &= \textit{by G2} \\ & \text{generate}(S, \lambda i.(\lambda x.B(\lambda i.e(i)))) \\ &= \textit{remove identity binding (lemma 1)} \\ & \text{generate}(S, \lambda i.(\lambda x.B(e))) \end{aligned}$$

□

Proof 7: Matrix generation, vector binding independent of one generating index

$$\begin{aligned} & \text{generate}([l, m], \lambda[i, j].(\lambda x.B(\text{generate}([n], \lambda[k].e)))) \\ \equiv & \lambda X.\text{generate}([l, m], \lambda[i, j].(\lambda x.B(\text{generate}([n], \lambda[k].X@[i, k]))) \\ & (\text{generate}([l, n], \lambda[i, k].e))) \\ & \text{where } e \text{ is depednent on } i \text{ and } k, \text{ but not on } j \end{aligned}$$

Proof follows by β -reducing X :

$$\begin{aligned} & \lambda X.\text{generate}([l, m], \lambda[i, j].(\lambda x.B(\text{generate}([n], \lambda[k].X@[i, k]))) \\ & (\text{generate}([l, n], \lambda[i, k].e))) \\ = & \beta\text{-reduce } X; \text{ only occurrence of } X \text{ is that shown} \\ & \text{generate}([l, m], \\ & \lambda[i, j].(\lambda x.B(\text{generate}([n], \lambda[k].\text{element}(\text{generate}([l, n], \lambda[i, k].e), [i, k]))) \\ = & \text{by } G2 \\ & \text{generate}([l, m], \lambda[i, j].(\lambda x.B(\text{generate}([n], \lambda[k].(\lambda[i, k].e([i, k]))))) \\ = & \text{remove identity binding (lemma 1)} \\ & \text{generate}([l, m], \lambda[i, j].(\lambda x.B(\text{generate}([n], \lambda[k].e)))) \end{aligned}$$

□

Proof 8: Binding is a function application

$$\begin{aligned} & \text{generate}(S, \lambda i.(\lambda x.B(\text{generate}(T, \lambda j.f(a)))) \\ \equiv & \text{generate}(S, \lambda i.(\lambda x'.(\lambda x.B(\text{generate}(T, \lambda j.f(x'@j)))) (\text{generate}(T, \lambda j.a)))) \end{aligned}$$

Proof follows by β -reducing x' :

$$\begin{aligned} & \text{generate}(S, \lambda i.(\lambda x'.(\lambda x.B(\text{generate}(T, \lambda j.f(x'@j)))) (\text{generate}(T, \lambda j.a)))) \\ = & \beta\text{-reduce } x'; \text{ only instance of } x' \text{ is that shown} \\ & \text{generate}(S, \lambda i.(\lambda x.B(\text{generate}(T, \lambda j.f(\text{element}(\text{generate}(T, \lambda j.a), j)))))) \\ = & \text{by } G2 \\ & \text{generate}(S, \lambda i.(\lambda x.B(\text{generate}(T, \lambda j.f(\lambda j.a(j)))))) \\ = & \text{remove identity binding (lemma 1)} \\ & \text{generate}(S, \lambda i.(\lambda x.B(\text{generate}(T, \lambda j.f(a)))) \end{aligned}$$

□

5.2 Proofs of Transformations for Reductions

Proof 9: generate-reduce combination

$$\begin{aligned} & \text{generate}([m], \lambda[i].\text{reduce}(r, r0, [n], \lambda[j].e)) \\ \equiv & \text{fold.rows}(r, \text{generate}([m], \lambda[i].r0), \text{generate}([m, n], \lambda[i, j].e)) \\ & \text{where } n \text{ and } r \text{ are independent of } i \end{aligned}$$

Proof follows directly from the definition of `fold.rows`:

$$\begin{aligned}
& \text{fold.rows}(r, \text{generate}([m], \lambda[i].r0), \text{generate}([m, n], \lambda[i, j].e)) \\
& = \text{by definition 5} \\
& \text{generate}([\text{shape}(\text{generate}([m, n], \lambda[i, j].e), 1)], \\
& \quad \lambda[i].\text{reduce}(r, \\
& \quad \quad \text{element}(\text{generate}([m], \lambda[i].r0), [i]), \\
& \quad \quad [\text{shape}(\text{generate}([m, n], \lambda[i, j].e), 2)], \\
& \quad \quad \lambda[j].\text{element}(\text{generate}([m, n], \lambda[i, j].e), [i, j])) \\
& = \text{by G1 and G2} \\
& \text{generate}([m], \lambda[i].\text{reduce}(r, \lambda[i].r0 ([i]), [n], \lambda[j].\lambda[i, j].e ([i, j]))) \\
& = \text{remove identity bindings (lemma 1)} \\
& \text{generate}([m], \lambda[i].\text{reduce}(r, r0, [n], \lambda[j].e))
\end{aligned}$$

□

Proof 10: generate-reduce swap

$$\begin{aligned}
& \text{generate}(S, \lambda i.\text{reduce}(r, r0, T, \lambda j.e)) \\
& \equiv \text{reduce}(\epsilon(r), \text{generate}(S, \lambda i.r0), T, \lambda j.\text{generate}(S, \lambda i.e)) \\
& \quad \text{where } r, r0 \text{ and } T \text{ are independent of } i \text{ and } S \text{ is independent of } j
\end{aligned}$$

Since both sides of this identity evaluate to arrays, proof of this identity requires proof that the two arrays have the same shape and the same elements.

Same Shapes

Left side:

$$\begin{aligned}
& \text{shape}(\text{generate}(S, \lambda i.\text{reduce}(r, r0, T, \lambda j.e))) \\
& = \text{by G1} \\
& S
\end{aligned}$$

Right side:

$$\begin{aligned}
& \text{shape}(\text{reduce}(\epsilon(r), \text{generate}(S, \lambda i.r0), T, \lambda j.\text{generate}(S, \lambda i.e))) \\
& = \text{by lemma 6} \\
& \text{shape}(\text{generate}(S, \lambda i.r0)) \\
& = \text{by G1} \\
& S
\end{aligned}$$

Same Elements

Consider an arbitrary element i' .

Left side:

$$\begin{aligned}
& \text{element}(\text{generate}(S, \lambda i.\text{reduce}(r, r0, T, \lambda j.e)), i') \\
& = \text{by G2} \\
& \lambda i.\text{reduce}(r, r0, T, \lambda j.e) (i') \\
& = \text{since } r, r0 \text{ and } T \text{ are independent of } i, \text{ move binding into reduction (lemma 3)} \\
& \text{reduce}(r, r0, T, \lambda i.(\lambda j.e) (i')) \\
& = \text{move binding of } i \text{ into abstraction of } j \text{ (lemma 2)} \\
& \text{reduce}(r, r0, T, \lambda j.(\lambda i.e) (i'))
\end{aligned}$$

Right side:

$$\begin{aligned}
& \text{element}(\text{reduce}(\epsilon(r), \text{generate}(\mathbf{S}, \lambda i \cdot r0), \mathbf{T}, \lambda j \cdot \text{generate}(\mathbf{S}, \lambda i \cdot e)), i') \\
&= \text{move element into reduction (lemma 7)} \\
& \text{reduce}(r, \text{element}(\text{generate}(\mathbf{S}, \lambda i \cdot r0), i'), \mathbf{T}, \\
& \quad \lambda j \cdot \text{element}(\text{generate}(\mathbf{S}, \lambda i \cdot e), i')) \\
&= \text{by G2} \\
& \text{reduce}(r, \lambda i \cdot r0 (i'), \mathbf{T}, \lambda j \cdot (\lambda i \cdot e (i'))) \\
&= \text{since } r0 \text{ is independent of } i \\
& \text{reduce}(r, r0, \mathbf{T}, \lambda j \cdot (\lambda i \cdot e (i')))
\end{aligned}$$

Hence, the left and right sides have the same shapes and the same elements; thus they are equivalent.

□

Proof 11: reduce-reduce combination

$$\begin{aligned}
& \text{reduce}(r, r0, \mathbf{S}, \lambda i \cdot \text{reduce}(r, r0, \mathbf{T}, \lambda j \cdot e)) \equiv \text{reduce}(r, r0, \mathbf{S} \times \mathbf{T}, \lambda ij \cdot e) \\
& \text{where } r0 \text{ is an identity element of } r \text{ and } r, r0 \text{ and } \mathbf{T} \text{ are independent of } i
\end{aligned}$$

Proof is by induction over \mathbf{S} .

Base Step: \emptyset

Left side:

$$\begin{aligned}
& \text{reduce}(r, r0, \emptyset, \lambda i \cdot \text{reduce}(r, r0, \mathbf{T}, \lambda j \cdot e)) \\
&= \text{by R1} \\
& r0
\end{aligned}$$

Right side:

$$\begin{aligned}
& \text{reduce}(r, r0, \emptyset \times \mathbf{T}, \lambda ij \cdot e) \\
&= \\
& \text{reduce}(r, r0, \emptyset, \lambda ij \cdot e) \\
&= \text{by R1} \\
& r0
\end{aligned}$$

Inductive Step: $\mathbf{S}+i'$

Assume identity holds for shape \mathbf{S} . Now consider shape $\mathbf{S}+i'$, where $i' \notin \mathbf{S}$.

Left side:

$$\begin{aligned}
& \text{reduce}(r, r0, \mathbf{S}+i', \lambda i \cdot \text{reduce}(r, r0, \mathbf{T}, \lambda j \cdot e)) \\
&= \text{by R2} \\
& r(\lambda i \cdot \text{reduce}(r, r0, \mathbf{T}, \lambda j \cdot e) (i'), \text{reduce}(r, r0, \mathbf{S}, \lambda i \cdot \text{reduce}(r, r0, \mathbf{T}, \lambda j \cdot e))) \\
&= \text{since } r, r0 \text{ and } \mathbf{T} \text{ are independent of } i, \\
& \quad \text{move binding for } i \text{ into reduce (lemma 2)} \\
& r(\text{reduce}(r, r0, \mathbf{T}, \lambda i \cdot (\lambda j \cdot e) (i')), \text{reduce}(r, r0, \mathbf{S}, \lambda i \cdot \text{reduce}(r, r0, \mathbf{T}, \lambda j \cdot e))) \\
&= \text{induction hypothesis; move binding for } i \text{ into abstraction of } j \text{ (lemma 3)} \\
& r(\text{reduce}(r, r0, \mathbf{T}, \lambda j \cdot (\lambda i \cdot e (i'))), \text{reduce}(r, r0, \mathbf{S} \times \mathbf{T}, \lambda ij \cdot e))
\end{aligned}$$

Right side:

$$\begin{aligned}
& \text{reduce}(r, r0, (S+i') \times T, \lambda ij \cdot e) \\
& = \text{by definition of set insertion} \\
& \text{reduce}(r, r0, (S \cup -i'') \times T, \lambda ij \cdot e) \\
& = \\
& \text{reduce}(r, r0, (S \times T) \cup (-i'' \times T), \lambda ij \cdot e) \\
& = \text{split reduction (lemma 8)} \\
& r(\text{reduce}(r, r0, S \times T, \lambda ij \cdot e), \text{reduce}(r, r0, -i'' \times T, \lambda ij \cdot e)) \\
& = \text{collapse singleton dimensions in reduction (lemma 9); } \tau \text{ commutes} \\
& r(\text{reduce}(r, r0, T, \lambda j \cdot (\lambda i \cdot e (i'))), \text{reduce}(r, r0, S \times T, \lambda ij \cdot e))
\end{aligned}$$

Hence, by induction, the identity holds for all shapes.

□

6 Conclusions

The transformations required for converting basic array expressions into whole-array form have been shown to preserve the meaning of expressions. The majority of the proofs of correctness are simple and many are trivial; even the more complex proofs are straightforward to carry out, and use only well-known techniques (primarily induction over sets).

The simplicity of the proofs is due in large measure to the decomposition of a derivation into independent stages and to the decomposition of each stage into a sequence of simple transformations — the effect of each transformational step is small and is consequently readily amenable to formal analysis. In addition, the postponement of consideration of imperative details until very late in a derivation allows most of the transformational steps to be made within a purely functional framework; indeed, in this paper, it was not necessary to consider imperative details at all, even though the motive for applying the transformations is to tailor a program to the peculiarities of a particular type of imperative system.

Transformation sequences provide a means of characterising the special features of particular parallel architectures and/or problems domains (e.g. array processors or sparse matrix problems). The transformations relating to array processors presented in this paper may be reused in the derivation of efficient implementations of algorithms for the solution of a range of problems. Further, they are applied automatically by a tool. For these reasons, it is particularly important that the transformation sequence be *proved* to be meaning preserving.

A Proofs of Preliminary Lemmas

Some basic properties of `generate` and `reduce` are established below.

A.1 Properties of Elementwise Applications

Proof of Lemma 4: Shape of an elementwise application

$$\text{shape}(\epsilon(f)(A, B)) \equiv \text{shape}(B)$$

Proof.

$$\begin{aligned}
& \text{shape}(\epsilon(f)(A, B)) \\
& = \text{definition of } \epsilon \text{ (definition 9)} \\
& \text{shape}(\lambda X, Y \cdot \text{generate}(\text{shape}(Y), \lambda i \cdot f(X@i, Y@i)) (A, B)) \\
& = \beta\text{-reduce} \\
& \text{shape}(\text{generate}(\text{shape}(B), \lambda i \cdot f(A@i, B@i))) \\
& = \text{by G1} \\
& \text{shape}(B)
\end{aligned}$$

□

Proof of Lemma 5: Element of an elementwise application

$$\begin{aligned} \text{element}(\epsilon(f)(A, B), i) &\equiv f(\text{element}(A, i), \text{element}(B, i)) \equiv f(A@i, B@i) \\ &\text{for } i \in \text{shape}(\epsilon(f)(A, B)) \end{aligned}$$

Proof.

$$\begin{aligned} &\text{element}(\epsilon(f)(A, B), i') \\ &= \text{definition of } \epsilon \text{ (definition 9)} \\ &\text{element}(\lambda X, Y \cdot \text{generate}(\text{shape}(Y), \lambda i \cdot f(X@i, Y@i)) (A, B), i') \\ &= \beta\text{-reduce} \\ &\text{element}(\text{generate}(\text{shape}(B), \lambda i \cdot f(A@i, B@i)), i') \\ &= \text{by } G2' \\ &\text{if } (i' \in \text{shape}(\epsilon(f)(A, B))) \text{ then } \lambda i \cdot f(A@i, B@i) (i') \text{ else } \perp \\ &= i' \in \text{shape}(\epsilon(f)(A, B)) \text{ is true by assumption} \\ &\lambda i \cdot f(A@i, B@i) (i') \\ &= \beta\text{-reduce} \\ &f(A@i', B@i') \end{aligned}$$

□

A.2 Properties of Reductions

Proof of Lemma 6: Shape of an ϵ -reduction

$$\text{shape}(\text{reduce}(\epsilon(r), R0, S, \lambda i \cdot g)) \equiv \text{shape}(R0)$$

Proof.

Proof is by induction on S.

Base Step: \emptyset

$$\begin{aligned} &\text{shape}(\text{reduce}(\epsilon(r), R0, \emptyset, \lambda i \cdot g)) \\ &= \text{by } R1 \\ &\text{shape}(R0) \end{aligned}$$

Inductive Step: $S+i'$

Assume lemma holds for shape S. Now consider shape $S+i'$, where $i' \notin S$.

$$\begin{aligned} &\text{shape}(\text{reduce}(\epsilon(r), R0, S+i', \lambda i \cdot g)) \\ &= \text{by } R2 \\ &\text{shape}(\epsilon(r)(\lambda i \cdot g (i'), \text{reduce}(\epsilon(r), R0, S, \lambda i \cdot g))) \\ &= \text{by lemma 4} \\ &\text{shape}(\text{reduce}(\epsilon(r), R0, S, \lambda i \cdot g)) \\ &= \text{by induction hypothesis} \\ &\text{shape}(R0) \end{aligned}$$

Hence, by induction, the lemma holds for all shapes.

□

Proof of Lemma 7: Element of an ϵ -reduction

$$\begin{aligned} &\equiv \text{element}(\text{reduce}(\epsilon(r), R0, S, \lambda i \cdot g), j) \\ &\text{reduce}(r, \text{element}(R0, j), S, \lambda i \cdot \text{element}(g, j)) \\ &\text{where } S \text{ is independent of } j \end{aligned}$$

Proof. Proof is by induction on the shape over which the reduction is performed.

Base Step: \emptyset

Left side:

$$\begin{aligned} & \text{element}(\text{reduce}(\epsilon(r), R0, \emptyset, \lambda i \cdot g), j) \\ &= \text{by } R1 \\ & \text{element}(R0, j) \end{aligned}$$

Right side:

$$\begin{aligned} & \text{reduce}(r, \text{element}(R0, j), \emptyset, \lambda i \cdot \text{element}(g, j)) \\ &= \text{by } R1 \\ & \text{element}(R0, j) \end{aligned}$$

Inductive Step: $S+i'$

Assume the lemma holds for shape S . Now consider shape $S+i'$, where $i' \notin S$.

Left side:

$$\begin{aligned} & \text{element}(\text{reduce}(\epsilon(r), R0, S+i', \lambda i \cdot g), j) \\ &= \text{by } R2 \\ & \text{element}(\epsilon(r)(\lambda i \cdot g(i')), \text{reduce}(\epsilon(r), R0, S, \lambda i \cdot g), j) \\ &= \text{by lemma 5} \\ & r(\text{element}(\lambda i \cdot g(i'), j), \text{element}(\text{reduce}(\epsilon(r), R0, S, \lambda i \cdot g), j)) \\ &= \text{by induction hypothesis} \\ & r(\text{element}(\lambda i \cdot g(i'), j), \text{reduce}(r, \text{element}(R0, j), S, \lambda i \cdot \text{element}(g, j))) \end{aligned}$$

Right side:

$$\begin{aligned} & \text{reduce}(r, \text{element}(R0, j), S+i', \lambda i \cdot \text{element}(g, j)) \\ &= \text{by } R2 \\ & r(\lambda i \cdot \text{element}(g, j)(i'), \text{reduce}(r, \text{element}(R0, j), S, \lambda i \cdot \text{element}(g, j))) \\ &= \text{since } j \text{ does not contain } i, \text{ propagate } \lambda\text{-binding into element (lemma 3)} \\ & r(\text{element}(\lambda i \cdot g(i'), j), \text{reduce}(r, \text{element}(R0, j), S, \lambda i \cdot \text{element}(g, j))) \end{aligned}$$

Hence, by induction, the lemma holds for all shapes.

□

Proof of Lemma 8: Reduction over a union

$$\begin{aligned} & \text{reduce}(r, r0, S \cup T, \lambda i \cdot e) \equiv r(\text{reduce}(r, r0, S, \lambda i \cdot e), \text{reduce}(r, r0, T, \lambda i \cdot e)) \\ & \text{where } r0 \text{ is an identity element of } r, \text{ and } S \text{ and } T \text{ are disjoint} \end{aligned}$$

Proof. Proof is by induction over S .

Base Step: \emptyset

Left side:

$$\begin{aligned} & \text{reduce}(r, r0, \emptyset \cup T, \lambda i \cdot e) \\ &= \\ & \text{reduce}(r, r0, T, \lambda i \cdot e) \end{aligned}$$

Right side:

$$\begin{aligned} & r(\text{reduce}(r, r0, \emptyset, \lambda i \cdot e), \text{reduce}(r, r0, T, \lambda i \cdot e)) \\ &= \text{by } R1 \\ & r(r0, \text{reduce}(r, r0, T, \lambda i \cdot e)) \\ &= \text{since } r0 \text{ is an identity of } r \\ & \text{reduce}(r, r0, T, \lambda i \cdot e) \end{aligned}$$

Inductive Step: S+i'

Assume lemma holds for shape S. Now consider shape S+i', where i' ∉ S ∪ T.

Left side:

$$\begin{aligned} & \text{reduce}(r, r0, (S+i') \cup T, \lambda i \cdot e) \\ &= \text{interchanging set insertion and union} \\ & \text{reduce}(r, r0, (S \cup T)+i', \lambda i \cdot e) \\ &= \text{by R2} \\ & r(\lambda i \cdot e (i'), \text{reduce}(r, r0, S \cup T, \lambda i \cdot e)) \\ &= \text{by induction hypothesis} \\ & r(\lambda i \cdot e (i'), r(\text{reduce}(r, r0, S, \lambda i \cdot e), \text{reduce}(r, r0, T, \lambda i \cdot e))) \end{aligned}$$

Right side:

$$\begin{aligned} & r(\text{reduce}(r, r0, S+i', \lambda i \cdot e), \text{reduce}(r, r0, T, \lambda i \cdot e)) \\ &= \text{by R2} \\ & r(r(\lambda i \cdot e (i'), \text{reduce}(r, r0, S, \lambda i \cdot e)), \text{reduce}(r, r0, T, \lambda i \cdot e)) \\ &= \text{since r is associative} \\ & r(\lambda i \cdot e (i'), r(\text{reduce}(r, r0, S, \lambda i \cdot e), \text{reduce}(r, r0, T, \lambda i \cdot e))) \end{aligned}$$

Hence, by induction, the lemma holds for all shapes.

□

Proof of Lemma 9: Collapsing singleton dimensions

$$\begin{aligned} & \text{reduce}(r, r0, -i'' \times S, \lambda ij \cdot e) \equiv \text{reduce}(r, r0, S, \lambda j \cdot (\lambda i \cdot e (i''))) \\ & \text{where } r0 \text{ is an identity element of } r \\ & \text{and where } i' \text{ and } i \text{ have the same dimensionality} \end{aligned}$$

Proof. Proof is by induction over S.

Base Step: ∅

Left side:

$$\begin{aligned} & \text{reduce}(r, r0, -i'' \times \emptyset, \lambda ij \cdot e) \\ &= \\ & \text{reduce}(r, r0, \emptyset, \lambda ij \cdot e) \\ &= \text{by R1} \\ & r0 \end{aligned}$$

Right side:

$$\begin{aligned} & \text{reduce}(r, r0, \emptyset, \lambda j \cdot (\lambda i \cdot e (i''))) \\ &= \text{by R1} \\ & r0 \end{aligned}$$

Inductive Step: S+j'

Assume lemma holds for shape S. Now consider shape S+j', where j' ∉ S.

Left side:

$$\begin{aligned}
& \text{reduce}(r, r0, -i'' \times (S+j'), \lambda ij \cdot e) \\
& = \text{by definition of set insertion} \\
& \text{reduce}(r, r0, -i'' \times (S \cup -j''), \lambda ij \cdot e) \\
& = \\
& \text{reduce}(r, r0, (-i'' \times S) \cup (-i'' \times -j''), \lambda ij \cdot e) \\
& = \text{split index set (lemma 8)} \\
& r(\text{reduce}(r, r0, -i'' \times S, \lambda ij \cdot e), \text{reduce}(r, r0, -i'' \times -j'', \lambda ij \cdot e)) \\
& = \text{induction hypothesis; cartesian product of singleton sets is a singleton set} \\
& r(\text{reduce}(r, r0, S, \lambda j \cdot (\lambda i \cdot e (i'))), \text{reduce}(r, r0, -i'' \times -j'', \lambda ij \cdot e)) \\
& = \text{by R2} \\
& r(\text{reduce}(r, r0, S, \lambda j \cdot (\lambda i \cdot e (i'))), r(\lambda ij \cdot e (i'j'), \text{reduce}(r, r0, \emptyset, \lambda ij \cdot e))) \\
& = \text{by R1 and by definition of index concatenation} \\
& r(\text{reduce}(r, r0, S, \lambda j \cdot (\lambda i \cdot e (i'))), r(\lambda i, j \cdot e (i', j'), r0)) \\
& = r0 \text{ is an identity of } r; \text{Currying} \\
& r(\text{reduce}(r, r0, S, \lambda j \cdot (\lambda i \cdot e (i'))), \lambda j \cdot (\lambda i \cdot e (i')) (j'))
\end{aligned}$$

Right side:

$$\begin{aligned}
& \text{reduce}(r, r0, S+j', \lambda j \cdot (\lambda i \cdot e (i'))) \\
& = \text{by R2} \\
& r(\lambda j \cdot (\lambda i \cdot e (i')) (j'), \text{reduce}(r, r0, S, \lambda j \cdot (\lambda i \cdot e (i')))) \\
& = \text{since } r \text{ is commutative} \\
& r(\text{reduce}(r, r0, S, \lambda j \cdot (\lambda i \cdot e (i'))), \lambda j \cdot (\lambda i \cdot e (i')) (j'))
\end{aligned}$$

Hence, by induction, the lemma holds for all shapes.

□

References

- [1] M Clint, Stephen Fitzpatrick, T J Harmer, P L Kilpatrick, and J M Boyle. A family of data-parallel derivations. In Wolfgang Gentzsch and Uwe Harms, editors, *Proceedings of High Performance Computing and Networking, Volume II*, volume 797 of *Lecture Notes in Computer Science*, pages 457–462. Springer-Verlag, April 1994.
- [2] M. Clint, Stephen Fitzpatrick, T.J. Harmer, P. Kilpatrick, and J.M. Boyle. A family of intermediate forms. Technical Report 1993/Nov-MC.SF.TJH.PLK.JMB, Department of Computer Science, The Queen's University of Belfast, Belfast BT7 1NN, Northern Ireland, UK, November 1993. [url:http://www.cs.qub.ac.uk/pub/TechReports/1993/Nov-MC.SF.TJH.PLK.JMB/](http://www.cs.qub.ac.uk/pub/TechReports/1993/Nov-MC.SF.TJH.PLK.JMB/).
- [3] Stephen Fitzpatrick, Terence J. Harmer, Alan Stewart, Maurice Clint, and James M. Boyle. The automated transformation of abstract specifications of numerical algorithms into efficient array processor implementations. *The Science Of Computer Programming*, 28(1), January 1997.
- [4] A. Wikström. *Functional Programming using Standard ML*. Prentice Hall, 1987.
- [5] Dennis Parkinson and John Litt, editors. *Massively Parallel Computing with the DAP*. Research Monographs in Parallel and Distributed Computing. The MIT Press, 1990.