

Communication lifting: fixed point computation for parallelism

Willem G. Vree and Pieter H. Hartel

Department of Computer Systems, University of Amsterdam
Kruislaan 403, 1098 SJ Amsterdam, The Netherlands
E-mail: {wimv, pieter}@fwi.uva.nl

Abstract

Communication lifting is a program transformation that can be applied to a synchronous process network to restructure the network. This restructuring in theory improves sequential and parallel performance. The transformation has been formally specified and proved correct and it has been implemented as an automatic program transformation tool. This tool has been applied to a small set of programs consisting of synchronous process networks. Measurements indicate performance gains in practice both with sequential and parallel evaluation. Communication lifting is a worthwhile optimisation to be included in a compiler for a lazy functional language.

1 Introduction

A process network is a system of communicating processes, which are connected by streams. The communicating processes are functions and the streams are potentially infinite lists of values upon which the functions operate. Programming with process networks has a long history, which dates back to the seminal work of Kahn (1974). Many special purpose languages such as Lucid (Ashcroft and Wadge, 1977), Esterel (Berry and Cosserat, 1984), Signal (Gautier *et al.*, 1987) and Lustre (Caspi *et al.*, 1987) have been developed to support programming with streams and process networks. The disadvantage of developing a new language is that it also requires a new implementation to be built. The approach that we will take is to build process networks using a subset of a standard lazy functional language, while taking special measures to guarantee good performance, both sequentially and in parallel. The evaluation mechanism of lazy functional languages naturally supports programming with potentially infinite lists of values (Peyton Jones, 1987). The process networks are thus embedded in a general purpose programming language, obviating the need for special compilers and language support systems. This approach has been advocated amongst others by Kelly (1989).

In an implementation of a lazy functional language, a stream is represented by a finite list, which is terminated by a *suspension*. This is a calculation that is suspended until further notice. A suspension can be revived and executed at any time to compute further elements of the stream, together with a new trailing suspension.

The latter can be executed in turn to build further list elements etc. In an implementation of a lazy functional language this mechanism of executing and building suspensions is completely automatic. Lazy evaluation of programs implementing process networks may incur considerable cost because each element in each stream requires executing a suspension and constructing a new one. When a large number of streams is involved, the cost may be prohibitively high. In practice large networks will indeed arise, for instance in simulations of digital circuits. Here each flip-flop is represented by two coupled *nand* functions that are mapped over streams of clock and data values. Even a small circuit will contain a large number of flip-flops, so that the simulation of such a circuit will require managing a large number of streams (Vree, 1989).

The cost of managing a large number of streams can be considerably reduced when the network is synchronous. In a synchronous network, executing one suspension will cause all suspensions on connected streams to be executed as well. All such closely related suspensions are said to belong to the same *generation*. It should thus be possible to revive and execute all suspensions in the same generation at the same time. The computations can be organised such that the management cost is shared between all streams. All streams are advanced by one generation at the same time.

The joint management of all suspensions in a synchronous network can be performed as follows. The *zip* of all streams in a network is a single stream, such that the original stream elements of one generation are gathered in one *state tuple*. The network as a whole will compute generation after generation, while managing only a single stream of tuples. Within a generation the original stream elements are accessible to the functions that used to operate on stream elements as the elements of one large tuple.

Practical networks are often synchronous; for instance a digital logic simulation is synchronous because all circuits are essentially driven by a common clock. Most of the special purpose languages that have been developed for programming with networks are also synchronous. It is thus important to develop efficient implementations of synchronous network programs. We claim that it is an advantage to be able to build such an efficient implementation without having to resort to developing a new language and a compiler for that language.

The purpose of this paper is to present a program transformation called *communication lifting* that takes a synchronous process network consisting of n streams into a network with a single stream of n -tuples. The transformation is rooted in the theory of recursive programs, based on the explicit calculation of fixed points of sets of recursive equations. This is the subject of section 2. Section 3 discusses the efficient implementation of programs that consist of synchronous process networks. Section 4 formally defines synchronous process networks. Communication lifting on simple process networks is described in Section 5. Section 6 describes a set of transformations that bring a more general synchronous process network in the form required for communication lifting proper. Performance measurements are reported in Section 7. A comparison with related work is given in Section 8 and the con-

clusions follow in Section 9. The correctness proofs of the program transformation may be found in the appendix.

2 Theoretical considerations: fixed points

Explicit calculation of the fixed point of a recursive program is both unusual (Allison, 1986) and inefficient (Manna *et al.*, 1973). Efficient computation rules such as the *normal order rule*, that can be shown to be safe (Vuillemin, 1973), are generally preferred. Direct fixed point iteration is inefficient because it calculates a sequence of approximations to the fixed point (if one exists). Each subsequent approximation is either the same as the previous, or better. The basic idea behind communication lifting is that as successive approximations are often the same, or almost the same, it may be more efficient to calculate the *changes* in the approximations only.

Consider as an example the system of Equations fib.a and fib.b over streams (infinite lists) to calculate the Fibonacci numbers $\bar{a} = 0, 1, 1, 2, 3, 5, \dots$ (The required auxiliary functions such as the family of functions map_n are defined in Figure 4.) The variables denoting streams are marked with a small wavy line (\bar{a}) to render these variables typographically distinct from other identifiers.

$$\begin{aligned} \bar{a} &= 0 : \bar{b} && \text{(fib.a)} \\ \bar{b} &= 1 : map_2 (+) \bar{a} \bar{b} && \text{(fib.b)} \end{aligned}$$

Using the function fix as the fixed point operator and the standard domain construction, the solution of these equations is given by:

$$\begin{aligned} \langle \bar{a}, \bar{b} \rangle &= fix\ fib && \text{(fib.fix)} \\ fib\ \langle \bar{a}, \bar{b} \rangle &= \langle 0 : \bar{b}, 1 : map_2 (+) \bar{a} \bar{b} \rangle && \text{(fib.def)} \end{aligned}$$

The communication lifting transformation causes the changes in the approximations

to be computed as follows. Define a new stream \bar{z} as the *zip* of the fixed point:

$$\begin{aligned} \bar{z} &= zip_2 (fix\ fib) && \\ \equiv &&& \text{(by fib.fix)} \\ \bar{z} &= zip_2 \langle \bar{a}, \bar{b} \rangle && \\ \equiv &&& \text{(by fib.a and fib.b)} \\ \bar{z} &= zip_2 \langle 0 : \bar{b}, 1 : map_2 (+) \bar{a} \bar{b} \rangle && \\ \equiv &&& \text{(unfold zip_2)} \\ \bar{z} &= \langle 0, 1 \rangle : zip_2 \langle \bar{b}, map_2 (+) \bar{a} \bar{b} \rangle && \\ \equiv &&& \text{(define plus } x = fst\ x + snd\ x \text{ and use the definition of } \bar{z}) \\ \bar{z} &= \langle 0, 1 \rangle : zip_2 \langle map_1\ snd\ \bar{z}, map_1\ plus\ \bar{z} \rangle && \\ \equiv &&& \text{(property of zip_n, see law 3.35 in Jeurig (1992)} \\ &&& \text{and define nextstate } x = \langle snd\ x, plus\ x \rangle) \\ \bar{z} &= \langle 0, 1 \rangle : map_1\ nextstate\ \bar{z} && \\ \equiv &&& \text{(property of iterate_0)} \\ \bar{z} &= iterate_0\ nextstate\ \langle 0, 1 \rangle && \end{aligned}$$

Hence (the *zip* of) the fixed point of a system of equations over streams can be calculated as follows: start with an initial state $\langle 0, 1 \rangle$ and $iterate_0$ over the state transformation function $nextstate$. Both the state transformation function and the initial state are systematically derived from the system of equations.

To complete the communication lifting transformation for the fib example, the required output stream \bar{a} must be recovered from the stream of pairs \bar{z} . This, and gathering the newly introduced definitions yields a complete program:

$$\begin{aligned} \bar{a} &= map_1\ fst\ (iterate_0\ nextstate\ \langle 0, 1 \rangle) \\ nextstate\ s &= \langle snd\ s, plus\ s \rangle \\ plus\ s &= fst\ s + snd\ s \end{aligned}$$

As a finishing touch, the definition of $nextstate$ can be simplified by unfolding the definitions of $plus$, fst and snd and by using pattern matching thus:

$$\begin{aligned} \bar{a} &= map_1\ fst\ (iterate_0\ nextstate\ \langle 0, 1 \rangle) \\ nextstate\ \langle a, b \rangle &= \langle b, c \rangle \\ &\quad where \\ c &= (+) a\ b \end{aligned}$$

The communication lifting transformation is discussed in more detail in the following sections of the paper. For now we note that for the application of map_2 in the network one definition is generated under the *where* of the $nextstate$ function, which captures basically the same calculation as the map_2 . Also for each application of $(:)$ in the network, there is one element in the state tuple processed by the $nextstate$ function.

The $iterate_0$ function controls the succession of generations of state tuples and

thereby captures all the recursion originally scattered throughout the network. Given the tuple produced by the previous generation, the *nextstate* function (which is not recursive!) calculates the tuple of the current generation. The expression $(\text{iterate}_0 \text{ nextstate } \langle 0, 1 \rangle)$ computes the sequence of changes in the fixed point approximation in the form of the sequence of state tuples:

$$\langle 0, 1 \rangle, \quad \langle 1, 1 \rangle, \quad \langle 1, 2 \rangle, \dots$$

This is a form of program synthesis as found in for instance Bird and Wadler (1988, Pages 131-132). Our method is a powerful generalisation of the procedure described there.

3 Practical considerations: parallelism

There are two practical aspects to communication lifting. The first is the reduction in the cost of managing a large number of streams, because after the transformation only a single stream remains to be managed. We will come back to this issue in section 7.

The second aspect is the possibility to evaluate the components of the state tuple in parallel. This has to be contrasted with the pipe-line parallelism of a process network. Before discussing this point further, the *fib* example must be extended slightly to introduce a possibility for parallel evaluation. The example as it stands does not allow parallel evaluation at all because only one addition is performed (on a stream of numbers). The extension consists of adding two more streams \bar{o} and \bar{p} to the network, to produce a running total of the Fibonacci sequence in the stream \bar{o} :

$$\begin{aligned} \bar{o} &= 0 : \bar{p} \\ \bar{p} &= \text{map_2 } (+) \bar{o} \bar{b} \\ \bar{a} &= 0 : \bar{b} \\ \bar{b} &= 1 : \bar{c} \\ \bar{c} &= \text{map_2 } (+) \bar{a} \bar{b} \end{aligned}$$

The original sub expression $\text{map_2 } (+) \bar{a} \bar{b}$ has been put in a separate equation \bar{c} . This makes it easier to draw a diagram for the network. The Fibonacci-sum program thus obtained will serve as the running example of the paper. The program generates a stream of Fibonacci numbers $\bar{a} = 0, 1, 1, 2, 3, 5, 8, \dots$ and then adds these numbers to produce the stream $\bar{o} = 0, 1, 2, 4, 7, 12, 20, \dots$

To compare different ways of parallel evaluation it is illustrative to look at diagrams of the untransformed process network and the communication lifted version. The diagram for a synchronous process network shows the streams as connections between the functions applied to the streams. The diagram of the Fibonacci-sum network is shown in Figure 1. The name of a stream in the network is used as the label on the corresponding edge. The label in a box is (the curried version of) the appropriate stream processing function.

Figure 1 shows that the two processes $\text{map_2 } (+)$ may perform the additions with pipe-line parallelism. With this example no speedup can be achieved in practice as




Fig. 1. The original Fibonacci-sum network with pipe-line parallelism.

the additions represent too little work. In practical networks sufficient coarse grain work should be available for the pipe-line to deliver speedups.

Communication lifting of the extended example produces the following program:

$$\begin{aligned} \bar{o} &= \text{map_1 } (\text{sel}_3 \ 1) (\text{iterate}_0 \ \text{nextstate } \langle 0, 0, 1 \rangle) \\ \text{nextstate } \langle o, a, b \rangle &= \langle p, b, c \rangle \\ &\text{where} \\ p &= (+) \ o \ b \\ c &= (+) \ a \ b \end{aligned}$$

This transformation is analogous to that of the Fibonacci program. The formal derivation will be given in Section 5.

The communication lifted program can be executed with master-slave style parallelism as follows. The master process is responsible for generating the successive states, so it executes the calls to *iterate_0* and *nextstate*. The first state is the constant tuple $\langle 0, 0, 1 \rangle$, to which the master thus applies *nextstate*. This calls on two slaves to perform the two additions $p = o + a$ and $c = a + b$. The slaves have both access to the current tuple where the values of o , a and b are stored. Once the slave processors are finished, the next tuple is ready. The master process now selects the appropriate component of the new tuple for output via the application of $\text{map_1 } (\text{sel}_3 \ 1)$. The cycle is then complete and the whole sequence may begin again. This is schematically shown in Figure 2. Here the dashed lines represent the flow of data used by the master-slave communication. The solid lines are streams as before.

The master-slave parallelism shown in Figure 2 corresponds exactly to the pipe-line that is present in the original program of Figure 1, where two processes $\text{map_2 } (+)$ and $\text{map_2 } (+)$ are connected by streams. A pipe-line has thus been transformed into a master slave relation by communication lifting. The calculations

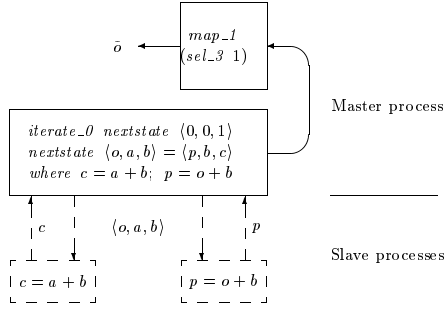


Fig. 2. Fibonacci-sum network after communication lifting with master-slave parallelism shown using dashed lines.

p	$::= d_1 \dots d_n$	
d_i	$::= \bar{v}_i = e_i$	(equations of the form $\bar{v} = \bar{w}$ are not permitted)
e	$::= s : e$	
	$map_n f e_1 \dots e_n$	
	$iterate_n f s e_1 \dots e_n$	
	$tl e$	
	\bar{v}	
s	$\in V$	(arbitrary values)
\bar{f}, \bar{g}	$\in \bar{F}$	(arbitrary functions on streams)
f, g	$\in F$	(arbitrary functions on stream elements)
$\bar{o}, \bar{v}, \bar{w}, \bar{x}, \bar{y}, \bar{z}$	$\in \bar{D}$	(streams)
s, v, w, x, y, z	$\in D$	(arbitrary stream elements)
i, j, k, l, m, n	$\in \mathbb{N}$	(natural numbers)

Fig. 3. Abstract syntax of the process network language.

in each step as performed by the slaves, are only dependent on the output produced by the previous step and are otherwise completely independent of each other.

It is possible to build a parallel implementation with a special primitive function that implements the required *master-slave* style parallelism. The tuples must then be constructed such, that at least two components require enough computation to outweigh the parallel overhead. The independence of the calculations on the tuple components should allow for a relatively cheap and simple mechanism to implement the parallel evaluation. Parallelism of a more general nature, such as pipe-line parallelism, is more difficult to harness efficiently. A more rigid paradigm (master-slave) allows the implementation more scope for optimisations than a more lenient paradigm (pipe-line). A more lenient paradigm offers the programmer better possibilities for clarity and conciseness. These claims are substantiated in Section 7.

$hd (x : \bar{x})$	$= x$
$tl (x : \bar{x})$	$= \bar{x}$
$\bar{x}!0$	$= hd \bar{x}$
$\bar{x}!(i+1)$	$= (tl \bar{x})!i$
$take 0 \bar{x}$	$= nil$
$take (i+1) \bar{x}$	$= hd \bar{x} : take i (tl \bar{x})$
$sel_n i (v_1, \dots, v_n)$	$= v_i \quad (\forall n \geq 1 \wedge 1 \leq i \leq n \text{ select } i\text{-th component})$
$map_n f \bar{v}_1 \dots \bar{v}_n$	$= w : map_n f (tl \bar{v}_1) \dots (tl \bar{v}_n)$ where $w = f (hd \bar{v}_1) \dots (hd \bar{v}_n)$
$iterate_n f s \bar{v}_1 \dots \bar{v}_n$	$= s : iterate_n f s' (tl \bar{v}_1) \dots (tl \bar{v}_n)$ where $s' = f s (hd \bar{v}_1) \dots (hd \bar{v}_n)$
$zip_n \langle \bar{v}_1, \dots, \bar{v}_n \rangle$	$= w : zip_n \langle tl \bar{v}_1, \dots, tl \bar{v}_n \rangle$ where $w = \langle hd \bar{v}_1, \dots, hd \bar{v}_n \rangle$

Fig. 4. Definition of stream processing and auxiliary functions and operators.

4 The definition of a synchronous process network

In this section the notations involved in the communication lifting program transformation are formally introduced. A network of synchronous processes is a graph, with synchronous processes as vertices and streams as edges. A process is synchronous if it is one of $(:)$, map_n , $iterate_n$ or tl . A synchronous process network should be represented by a number of equations over streams, according to the syntax in Figure 3. There must be one equation in the network for the stream \bar{o} , which by convention, generates the output of the network. No two equations in a network may have the same left hand side. The network graph must be connected, with the equation for \bar{o} as the root. Equations of the form $\bar{v} = \bar{w}$ are not permitted, they can always be eliminated by substitution. All these restrictions are necessary to allow communication lifting to be implemented as an automatic program transformation, for instance as part of a compiler.

The free stream variables of a synchronous process network are taken to be provided as input to the network from outside. A network may have any number of input streams, but need not have any.

The definitions of a number of useful stream processing and auxiliary functions are shown in Figure 4. According to the syntax, only tl , $(:)$, map_n and $iterate_n$ can be used as stream processing functions. The other functions and operators are also shown here because of their use in the transformation process. The choice of functions that may be applied to streams seems rather limited. However, considering that the functions that may be applied to *stream elements* are not constrained, the abstract syntax is actually quite general. Only three functions are required to build a synchronous process network: one to extend a stream up front $(:)$, one to trim the first element off a stream (tl) and a third function to perform an arbitrary computation on a stream element (map_n). A fourth function $iterate_n$ has been

included because it captures the concept of a function that carries its own local state.

To be completely general, functions that add and remove arbitrary stream elements should have been supported as well, for instance *filter*. We have chosen not to include such functions as it makes it more difficult to guarantee that the networks constructed are synchronous. Work is in progress on an extension of the method to also support some forms of asynchronous networks, in which functions such as *filter* play a role.

Functions with a suffix *_n* represent a whole family of functions, because *n* is a natural number. In an enumeration such as $e_1 \dots e_n$, *n* may also be equal to 0, which means that there is not even a single expression e_i present. The function *sel_0* is ill defined, but *map_0* and *iterate_0* are valid functions. Note that the definition of *map_n* does not correspond with the usual definition, because there is no test whether any of the input streams $\bar{v}_1 \dots \bar{v}_n$ are empty. This is consistent with the view that streams are infinite (lists).

5 Communication lifting of a synchronous process network

The communication lifting transformation of a synchronous process network, as defined according to the syntax of Figure 3, will now be presented in two steps. As the first step, we present the communication lifting of a simplified process network. The second step (Section 6) brings a more general network that conforms to the abstract syntax of Figure 3 into the simplified form.

The communication lifting transformation proper is given by Rule *T0* of Figure 5. The notation employed is more or less standard (see for example Ferguson and Wadler (1988)). The transformation rules take a syntactic argument enclosed in emphatic brackets \llbracket and \rrbracket . Pattern matching is used to choose between alternative clauses. The matching order is top down, thus Clause 0b is a *catch-all* clause, fitting any network that does not match Clause 0a. The matching of some clauses (e.g. Clause 0a) is further constrained by a guard, written as a conditional (*if...*) \implies connecting the left and right hand side of the clause. A clause protected with a guard matches only if both the pattern and the guard are satisfied. If either fails, the next clause will be tried.

A simplified synchronous process network must match the left hand side of Clause 0a. This means that the first *k* equations must contain an application of $(:)$ and that the next *m* equations must contain an application of *map_n*. The remaining *l* equations must use *tl*. At this stage applications of *iterate_n*, equations of the form $\bar{v} = \bar{w}$, or nested expressions are not permitted. These restrictions are necessary to make the presentation of the communication lifting transformation proper reasonably succinct. The lifting of the restrictions is the subject of the next section.

Some of the stream variables \bar{x}_i and \bar{y}_{ij} in the left hand side of Clause 0a will be the same as some of the variables \bar{o} , \bar{v}_i , \bar{w}_i or \bar{t}_i . The \bar{x}_i and \bar{y}_{ij} that are not defined within the network act as the input streams to the network. These input streams are identified by the set $\{\bar{u}_1 \dots \bar{u}_h\}$.

The guard in Clause 0a ensures that the streams $\bar{z}_1 \dots \bar{z}_l$, and the states $s_1 \dots s_k$,

$$\begin{aligned}
T0 \llbracket & \begin{array}{l} \bar{o} = s_1 : \bar{x}_1 \\ \bar{v}_2 = s_2 : \bar{x}_2 \\ \vdots \\ \bar{v}_k = s_k : \bar{x}_k \\ \bar{w}_1 = \text{map_}n_1 f_1 \bar{y}_{11} \dots \bar{y}_{1n_1} \\ \vdots \\ \bar{w}_m = \text{map_}n_m f_m \bar{y}_{m1} \dots \bar{y}_{mn_m} \\ \bar{t}_1 = \text{tl } \bar{z}_1 \\ \vdots \\ \bar{t}_l = \text{tl } \bar{z}_l \end{array} \rrbracket \quad (\text{if } \begin{array}{l} \{\bar{z}_1 \dots \bar{z}_l\} \cap \mathcal{N} = \emptyset \wedge \\ \text{freevars}\{s_1 \dots s_k\} \cap \mathcal{N} = \emptyset \wedge \\ \text{freevars}\{f_1 \dots f_m\} \cap \mathcal{N} = \emptyset \end{array} \implies \\ & \begin{array}{l} \bar{o} = \text{map_}1 (\text{sel_}k \ 1) (\text{iterate_}h \ \text{nextstate } (s_1, \dots, s_k) \ \bar{u}_1 \dots \bar{u}_h) \\ \bar{t}_1 = \text{tl } \bar{z}_1 \\ \vdots \\ \bar{t}_l = \text{tl } \bar{z}_l \\ \text{nextstate } s \ u_1 \dots u_h = \langle x_1, \dots, x_k \rangle \\ \text{where} \\ \langle o, v_2, \dots, v_k \rangle = s \\ w_1 = f_1 \ y_{11} \dots y_{1n_1} \\ \vdots \\ w_m = f_m \ y_{m1} \dots y_{mn_m} \end{array} \\ & \text{where } \bar{u}_1 \dots \bar{u}_h \text{ are defined as} \\ & \begin{array}{l} \{\bar{u}_1 \dots \bar{u}_h\} = (\{\bar{x}_1 \dots \bar{x}_k\} \cup \{\bar{y}_{ij} : i \in \{1 \dots m\} \wedge j \in \{1 \dots n_i\}\}) \setminus (\mathcal{N} \cup \{\bar{t}_1 \dots \bar{t}_l\}) \\ \mathcal{N} = \{\bar{w}_1 \dots \bar{w}_m\} \cup \{\bar{o}, \bar{v}_2 \dots \bar{v}_k\} \end{array} \\ & T0 \llbracket p \rrbracket \quad (\text{otherwise}) \implies p \end{aligned} \tag{0a}$$

Fig. 5. Communication lifting transformation.

and the functions $f_1 \dots f_m$ are not dependent on any of the streams $\bar{o}, \bar{v}_2 \dots \bar{v}_k$ and $\bar{w}_1 \dots \bar{w}_m$. This is necessary because the transformation removes the definitions of the streams \bar{v}_i and \bar{w}_j .

The best way to understand Rule *T0* is to try it out on a simple example. The program produced by applying Rule *T0* to the Fibonacci network of Section 3 yields:

$$\begin{array}{l|l} \bar{o} = s_1 : \bar{x}_1 & \bar{o} = 0 : \bar{p} \\ \bar{v}_2 = s_2 : \bar{x}_2 & \bar{a} = 0 : \bar{b} \\ \bar{v}_3 = s_3 : \bar{x}_3 & \bar{b} = 1 : \bar{c} \\ \bar{w}_2 = \text{map_}2 f_1 \bar{y}_{11} \bar{y}_{12} & \bar{p} = \text{map_}2 (+) \bar{o} \bar{b} \\ \bar{w}_2 = \text{map_}2 f_2 \bar{y}_{21} \bar{y}_{22} & \bar{c} = \text{map_}2 (+) \bar{a} \bar{b} \end{array} \xrightarrow{T0}$$

Here we have shown the correspondence between formal and actual identifiers of the Rule *T0*. The guard of Clause 0a is satisfied: because there are no equations of type $\bar{t} = \text{tl } \bar{z}$, and because $s_1 = 0, s_2 = 0, s_3 = 1, f_1 = (+)$, and $f_2 = (+)$ do not depend on $\mathcal{N} = \{\bar{o}, \bar{a}, \bar{b}, \bar{p}, \bar{c}\}$. For the Fibonacci-sum program the set $\{\bar{u}_1 \dots \bar{u}_h\}$ is

empty, which means that the network does not use external input streams:

$$\begin{aligned} \bar{o} &= \text{map_1 } (\text{sel_3 } 1) (\text{iterate_0 } \text{nextstate } \langle 0, 0, 1 \rangle) \\ \text{nextstate } s &= \langle p, b, c \rangle \\ \xrightarrow{T0} \quad &\text{where} \\ &\langle o, a, b \rangle = s \\ &p = (+) o b \\ &c = (+) a b \end{aligned}$$

$$T1 \llbracket d_1 \dots d_n \rrbracket \Rightarrow T2 \llbracket d_1 \rrbracket \dots T2 \llbracket d_n \rrbracket \quad (1)$$

$$\begin{aligned} T2 \llbracket \bar{v} = f \dots (g \dots) \dots \rrbracket &\Rightarrow T2 \llbracket \bar{v} = f \dots \bar{w} \dots \rrbracket \\ &T2 \llbracket \bar{w} = (g \dots) \rrbracket \end{aligned} \quad (2a)$$

where \bar{w} is a new variable

$$T2 \llbracket \bar{v} = e \rrbracket \Rightarrow \bar{v} = e \quad (2b)$$

$$T3 \llbracket d_1 \dots d_n \rrbracket \Rightarrow T4 \llbracket d_1 \rrbracket \dots T4 \llbracket d_n \rrbracket \quad (3)$$

$$\begin{aligned} T4 \llbracket \bar{v} = \text{iterate_n } f s \bar{x}_1 \dots \bar{x}_n \rrbracket &\Rightarrow \bar{v} = s : \bar{w} \\ &\bar{w} = \text{map_m } f \bar{v} \bar{x}_1 \dots \bar{x}_n \end{aligned} \quad (4a)$$

where \bar{v} and \bar{w} are new variables and $m = n + 1$

$$T4 \llbracket d \rrbracket \Rightarrow d \quad (4b)$$

$$T5 \llbracket \bar{o} =_k \text{map_n } f \bar{x}_1 \dots \bar{x}_n \quad d_2 \dots d_m \rrbracket \rho \Rightarrow T5 \llbracket \bar{o} =_k f (hd \bar{x}_1) \dots (hd \bar{x}_n) : \bar{v} \quad \bar{v} =_k tl \bar{w} \quad \bar{w} =_k \text{map_n } f \bar{x}_1 \dots \bar{x}_n \quad d_2 \dots d_m \rrbracket \rho \quad (5a)$$

where \bar{v} and \bar{w} are new variables

$$T5 \llbracket \bar{o} =_k tl \bar{w} \quad \bar{w} =_j s : \bar{x} \quad d_3 \dots d_m \rrbracket \rho \Rightarrow T5 \llbracket (\bar{w} =_j s : \bar{o} \quad d_3 \dots d_m) [\bar{o}/\bar{x}] \rrbracket \rho \quad (5b)$$

$$T5 \llbracket \bar{v} =_k tl \bar{w} \quad \bar{w} =_j s : \bar{x} \quad d_3 \dots d_m \rrbracket \rho \Rightarrow T5 \llbracket (\bar{w} =_j s : \bar{x} \quad d_3 \dots d_m) [\bar{x}/\bar{v}] \rrbracket \rho \quad (5c)$$

$$T5 \llbracket \bar{v} =_k tl \bar{w} \quad \bar{w} =_j \text{map_n } f \bar{x}_1 \dots \bar{x}_n \quad d_3 \dots d_m \rrbracket \rho \Rightarrow T5 \llbracket \bar{w} =_j f (hd \bar{x}_1) \dots (hd \bar{x}_n) : \bar{v} \quad \bar{v} =_j \text{map_n } f \bar{y}_1 \dots \bar{y}_n \quad \bar{y}_1 =_k tl \bar{x}_1 \quad \vdots \quad \bar{y}_n =_k tl \bar{x}_n \quad d_3 \dots d_m \rrbracket (\rho \cup \{(j, k)\}) \quad (5d)$$

$$(if (j, k) \notin \rho) \Rightarrow T5 \llbracket \bar{w} =_j f (hd \bar{x}_1) \dots (hd \bar{x}_n) : \bar{v} \quad \bar{v} =_j \text{map_n } f \bar{y}_1 \dots \bar{y}_n \quad \bar{y}_1 =_k tl \bar{x}_1 \quad \vdots \quad \bar{y}_n =_k tl \bar{x}_n \quad d_3 \dots d_m \rrbracket (\rho \cup \{(j, k)\})$$

where $\bar{y}_1 \dots \bar{y}_n$ are new variables

$$T5 \llbracket p \rrbracket \rho \quad (\text{otherwise}) \Rightarrow p \quad (5e)$$

Fig. 6. Simplifying transformations: $p' = T5 (T3 (T1 p)) \emptyset$.

6 Simplifying transformations

A synchronous process network specified according to the syntax of Figure 3 has to be transformed into a simplified form that is acceptable to Rule $T0$. Figure 6 shows the simplifications that are performed by the successive application of the Rules $T1$, $T3$ and $T5$. We will explain the purpose of each of these rules in turn.

First we note, that the Fibonacci-sum program as discussed in Section 3 is actually the outcome of applying the simplifying transformations to the following program:

$$\begin{aligned} \bar{o} &= \text{iterate_1 } (+) 0 (tl \bar{a}) \\ \bar{a} &= 0 : 1 : \text{map_2 } (+) \bar{a} (tl \bar{a}) \end{aligned} \quad \xrightarrow{T1}$$

This program cannot be transformed directly by Rule $T0$ because: it uses the iterate_1 function; it uses nested function applications and the output stream \bar{o} is not an application of $(:)$. The purpose of the simplifying transformations is to eliminate these constructs.

Rule $T1$: removing nested function applications

The first problem to solve is to remove nested function applications. This is the purpose of Rule $T1$, which defines a new equation for each nested expression.

Rule $T1$ applies Rule $T2$ to all equations of the network. Clause 2a introduces extra stream equations for all nested expressions that occur in the synchronous process network. The pattern $\bar{v} = f \dots (g \dots) \dots$ is matched by an equation that starts with an application of some function f and that contains a nested application $(g \dots)$. If there is more than one nested application that can be matched by $(g \dots)$, the left most is chosen. This choice is arbitrary as $T2$ is reapplied to the results $\bar{v} = f \dots \bar{w} \dots$ and $\bar{w} = (g \dots)$ so that remaining nested applications will be dealt with eventually. According to the syntax, the functions f and g must be one of tl , $(:)$, map_n or iterate_n . On the right hand side of Clause 2a the newly introduced equation bears a name \bar{w} that must not appear anywhere else in the network. Rule $T2$ is applied recursively until all arguments of each application of f are simple stream identifiers. The default Clause 2b terminates the recursion.

Applying Rule $T1$ to the Fibonacci-sum program as given above introduces four new equations for \bar{b} , \bar{c} , \bar{d} and \bar{q} to remove the nested function applications. The result of this transformation is:

$$\begin{aligned} \bar{o} &= \text{iterate_1 } (+) 0 \bar{q} \\ \bar{q} &= tl \bar{a} \\ \xrightarrow{T1} \quad \bar{a} &= 0 : \bar{b} \\ \bar{b} &= 1 : \bar{c} \\ \bar{c} &= \text{map_2 } (+) \bar{a} \bar{d} \\ \bar{d} &= tl \bar{a} \end{aligned} \quad \xrightarrow{T3}$$

Rule T3: removing calls to the function *iterate_n*

The purpose of Rule *T3* is to simplify the network by replacing all occurrences of the function *iterate_n* by applications of the simpler *map_m* function.

Rule *T3* applies Rule *T4* to all the equations of the network, which after application of Rule *T1* contain a single function application each. Each equation with *iterate_n* is transformed into an equation with $(:)$ and an equation with *map_m*. Clause 4b retains all other equations as they are.

Application of Rule *T3* to the Fibonacci-sum network as delivered by Rule *T1* replaces the equation for \bar{o} by two new equations. The equations for $\bar{a} \dots \bar{d}$ and \bar{q} remain unchanged:

$$\begin{array}{lcl} \bar{o} & = & 0 : \bar{p} \\ \bar{p} & = & \text{map_2 } (+) \bar{o} \bar{q} \\ \bar{q} & = & tl \bar{a} \\ \xrightarrow{T3} \bar{a} & = & 0 : \bar{b} \qquad \xrightarrow{T3} \\ \bar{b} & = & 1 : \bar{c} \\ \bar{c} & = & \text{map_2 } (+) \bar{a} \bar{d} \\ \bar{d} & = & tl \bar{a} \end{array}$$

Rule T5: removing redundant calls to *tl*

The goal of the final simplifying Rule *T5* is to remove as many applications of *tl* as possible. Rule *T5* also ensures that the output stream is defined as an application of $(:)$. These apparently different purposes have to be served by one transformation, as both involve calls to the function *tl*.

Before discussing Rule *T5* proper, a further explanation is appropriate about the pattern matching of clauses such as Clause 5c, which process several equations simultaneously. The variable \bar{w} on the left hand side of the Clause 5c occurs twice in the pattern, which means that both occurrences must match the same stream. This links two equations of the process network. Thus far no ordering on the equations in the process network has been assumed, because rules *T2* and *T4* can be applied in any order. For Rule *T5* it is convenient to regard the equations that define the process network under consideration as a proper set. Any subset of equations satisfying the constraints specified by the patterns may be chosen. Clause 5c will thus select any equation $\bar{v} = tl \bar{w}$ and the corresponding equation $\bar{w} = s : \bar{x}$. The remaining equations are named $d_3 \dots d_m$ and retained so that they can be processed by the recursive call to Rule *T5*.

Figure 6 shows that all equations of Rule *T5* are labelled with a subscript ($=_i$). This is necessary to ensure that Rule *T5* will terminate on all inputs. The labels are used in the guard (*if* $(j, k) \notin \rho$) \implies of Clause 5d to avoid this clause from looping on a pair of definitions such as $\bar{v} = tl \bar{w}; \bar{w} = \text{map_1 } f \bar{v}$. The labelling can be added to a system of equations, by numbering each equation and using the equation number as the label number. The only assumption about the labels is that they are all different when they are first assigned.

Clause 5a unfolds the definition of *map_n* once and introduces two new equations so that $\bar{o} = \dots$ as required.

Clause 5c resolves a combination of $\bar{v} = tl \bar{w}$ and $\bar{w} = s : \bar{x}$ by replacing every free occurrence of \bar{v} in the process network by \bar{x} . This replacement is expressed as $(\dots) [\bar{x}/\bar{v}]$. The equation for \bar{v} is then removed. The *tl* is thus cancelled against the $(:)$. A combination of $\bar{v} = tl \bar{w}$ and $\bar{w} = \text{map_n } \dots$ (Clause 5d) can be resolved in a similar way, after unfolding the definition of *map_n* once. Then the *tl* can be cancelled against the $(:)$.

Clause 5b is a special case of Clause 5c. Both clauses cancel an application of *tl* against an application of $(:)$, but the output stream of the network \bar{o} must be treated specially. If Clause 5b were omitted, Clause 5c applied to $\bar{v} = tl \bar{o}$ would remove the equation for \bar{o} from the network. This would make it impossible for Rule *T0* later to retrieve the output stream from the network.

Rule *T5* has no case for combinations such as $\bar{v} = tl \bar{w}$ with $\bar{w} = tl \bar{x}$, because a combination of two or more *tl* applications is resolved by cancellation of the last *tl* against either *map_n* or $(:)$, followed by cancellation of the penultimate *tl* etc. A combination of equations involving applications of *tl* can not be removed by cancellation if the last *tl* is applied to an input stream. This case is adequately handled by Rule *T0* and will not concern Rule *T5*.

When applied to the Fibonacci-sum example program, Rule *T5* cancels the applications of *tl* by using Clause 5c twice. This yields the Fibonacci-sum network in a compact form, ready for the final Rule *T0*. It is the same program as the one we started with in Section 3 and also in Section 5.

$$\begin{array}{lcl} \bar{o} & = & 0 : \bar{p} \\ \bar{p} & = & \text{map_2 } (+) \bar{o} \bar{b} \\ \xrightarrow{T5} \bar{a} & = & 0 : \bar{b} \\ \bar{b} & = & 1 : \bar{c} \\ \bar{c} & = & \text{map_2 } (+) \bar{a} \bar{b} \end{array}$$

The simplifying transformations as defined in Figure 6 are necessary and sufficient to bring a system of equations over streams as specified according to the syntax of Figure 3 into the form required by the communication lifting transformation proper as given by Rule *T0*.

7 The performance of a number of small networks

The communication lifting transformation consists of a number of fold/unfold steps (Burstall and Darlington, 1977) and uses algebraic properties of stream functions such as *map_n* and *zip_n*. The communication lifting transformation incorporates a strategy that decides which steps to take, to guarantee delivery of an equivalent but completely restructured program. Communication lifting can thus be viewed as a transformation skeleton (Darlington *et al.*, 1991).

Communication lifting as a programming tool is only useful if the transformed programs will run faster, and/or use less space, than the original programs. Three important performance issues can be distinguished. The first is the gain or loss in

Table 1. Synchronous process networks with an indication of their size and purpose. Execution times are reported in milliseconds.

system program	source lines	sequential		master-slave trans. (ms)	pipe orig. (ms)	tuple size	stream length
		orig. (ms)	trans. (ms)				
<code>fibsum</code> ^a	6	627	493	477	536	3	3000
<code>flipflop</code> ^b	28	2434	760	581	974	13	1600
<code>fft</code> ^c	190	3045	2956	1364	1504	8	5
<code>wave4</code> ^d	264	2041	2131	1111	1151	3	200

^a Sums the first 3000 Fibonacci numbers.

^b Simulation of a D-flipflop over 1600 state transitions (Vree, 1989).

^c A 1024-point fast Fourier transform using arrays (Hartel and Vree, 1992).

^d Predicts the water heights and velocities in a square area of 4×4 grid points of the North Sea over 200 time steps (Vree, 1989).

sequential performance due to transformation, the second is the difference between sequential and master-slave parallel performance, the third issue is the difference between master-slave parallel performance of the transformed programs and pipe-line parallel performance of the untransformed programs. These issues will be discussed in the three following sections. To make measurements possible, communication lifting has been implemented and applied to a small set of synchronous process networks. The transformation has been implemented in Miranda† (Turner, 1985) and the input and output of the transformations are also Miranda programs. The abstract syntax of Figure 3 can be embedded in that of Miranda.

7.1 Sequential performance

To assess the impact of communication lifting on sequential performance, the set of synchronous process networks has been compiled and executed both before and after the complete set of transformations. The programs are compiled by the FAST compiler (Hartel *et al.*, 1991; Langendoen and Hartel, 1992), which amongst others, provides efficient arrays for the benefit of the `fft` and the `wave4` applications. The programs are executed on a stand alone Motorola 88000 processor board with 64Mb of memory. This allows execution time to be measured with an accuracy of 1 millisecond. For each program, Table 1 gives an indication of the size, the execution times of four versions of the program, the size of the state tuple, the number of elements of the streams that are evaluated and an explanation of the purpose of the program. The programs are sorted by the number of lines of source text. This count is exclusive of standard library functions as provided by Miranda, comments and blank lines.

† Miranda is a trademark of Research software Ltd.

The columns *sequential/orig.* and *sequential/trans.* show the sequential execution times (in milli-seconds) of a program before and after transformation. The sequential performance of most programs is improved, which shows that communication lifting is a viable optimisation technique for sequential programs. The significant performance gain in the `flipflop` program is due to the fact that instead of managing 13 lists, as is the case before the transformation, the transformed program only needs to manage 2 lists, which can be done more efficiently. The effect is strongest for the `flipflop` program, because it uses more streams than the others. For the `fibsum` program, there are 3 streams before the transformation and still 2 streams after the transformation. So only a small reduction of the stream management effort is the result.

The large networks that occur in real programs will lead to huge tuples. However, the implementation creates the state tuples in a single heap allocation, whereas a network of streams gives rise to more heap claims of smaller cells, which is thus more costly. The selection of an element of a tuple is performed in unit time, which is the same as in a stream network. When a small amount of work is involved in computing the stream elements, communication lifting will allow sequential performance gains on practical programs. For programs that involve large amounts of work on the stream elements, such as `wave4` and `fft`, sequential performance will not be affected much.

7.2 Parallel performance of the master-slave system

The sequential performance improvement for most of the programs indicates that communication lifting is a valid point of departure for parallel evaluation of independent tuple elements. Parallel evaluation always introduces overhead. This overhead should be kept small in comparison to the amount of real work involved in the evaluation of the state tuples. For example, two of the three tuple elements in the transformed `wave4` program represent a large amount of computation. To assess the parallel performance of the programs after communication lifting, an implementation has been built that supports the required master-slave parallelism. The Motorola 88000 system has 4 CPUs, 4 instruction caches and 4 coherent data caches and 64Mb shared memory. Cache coherency is handled by the hardware. The four processors are numbered 0, 1, 2 and 3. Processor 0 is the master processor, which begins execution. The other three processors are initially idle. The runtime system of the FAST compiler supports master-slave parallelism through a special primitive function *pforce*, which when applied to a tuple allocates the evaluation of each component of the tuple to a separate processor. The scheduling strategy is as follows: processor *p* first evaluates component *p* to full normal form (not just to head normal form). Then, as soon as this terminates, processor *p* evaluates component *p* + 4 to full normal form, then *p* + 8 etc. Parallel evaluation continues until all components of the state tuple have been evaluated to full normal form. The function *pforce* thus has a completely strict semantics.

During parallel evaluation the master processor behaves as an ordinary slave processor. Processor 0 becomes master again as soon as all tuple components have

been evaluated, at which point a new tuple is formed by the master processor. All other processors remain idle until the master encounters the next application of *pforce*.

In the communication lifted version of the Fibonacci-sum program the forming of the tuple is expressed as follows:

$$\begin{aligned} \bar{o} &= \text{map_1 } (\text{sel_3 } 1) (\text{iterate_0 } \text{nextstate } \langle 0, 0, 1 \rangle) \\ \text{nextstate } s &= \text{pforce } \langle p, b, c \rangle \\ &\text{where} \\ \langle o, a, b \rangle &= s \\ p &= (+) o b \\ c &= (+) a b \end{aligned}$$

The two additions $(+) a b$ and $(+) o b$ are thus evaluated in parallel. The first tuple ever created by the program contains three numbers, all of which are normal forms: $\langle 0, 0, 1 \rangle$. Because of the strict semantics of *pforce* as required by the master-slave style parallel evaluation, each subsequent tuple will be fully normalised when created. The sharing between computations is maintained as usual. In the example, both additions use the variable b , which is obtained from the current tuple and made available to all processors that need the value through the shared heap.

In all four programs that we have used, the only sharing that occurs is between the elements of the state-tuple that is passed as an argument to *nextstate*. Because of the completely strict semantics of *pforce*, all state tuple elements are fully normalised before *nextstate* is entered. Therefore, no locking/blocking mechanism is needed to prevent concurrent reduction of shared expressions. The implementation of master/slave parallelism for communication lifted programs is thus simple and fast. No locking overhead is incurred and scheduling is equally simple and fast. As we will see in the next section it is more complicated to implement pipe-line parallelism.

Not all programs that one might wish to transform by communication lifting will have the property that only normal forms are shared. In general, CAF's may be shared between processes so that a locking mechanism is needed. In future work we will identify extra conditions on synchronous process networks to guarantee that only normal forms are shared after communication lifting.

The column *master-slave/trans.* in Table 1 shows the parallel performance of the four test programs after communication lifting on the 4-processor system. The performance of three programs is improved by parallel execution, that of the *fibsum* program is not affected.

The granularity of the *fibsum* program is only one addition, just about sufficient to compensate the overhead of process creation by *pforce*. But the *flipflop* program, still quite fine grained, is already faster than the sequential version. This shows that the overhead of *pforce* is indeed low.

The *fft* program has four coarse grain and four fine grain processes (the tuple size is 8). However, there is a substantial amount of sequential processing before and after these processes are created. About half of the time is spent in the sequential parts of the program. The speedup is about 2.2

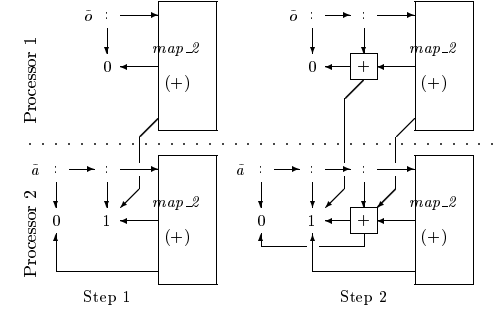


Fig. 7. Two processors performing pipe-line parallel graph reduction in a shared memory system. Processor 1 computes the Fibonacci-sum stream \bar{o} and Processor 2 computes the Fibonacci stream \bar{a} . The boxes represent suspended computations, all other nodes represent data.

The *wave4* program has three processes (the tuple size is 3), two of which contain a significant amount of work. This explains the speedup of about 1.9.

7.3 Parallel performance of a pipe-line implementation

The performance of the master-slave parallel system will now be compared with a pipe-line parallel system. The best way to do this is by using two implementations that are similar in as many ways as possible, so that the differences can be attributed to the differences between master-slave and pipe-line parallelism.

The master-slave parallel system provides most of the mechanism required to implement pipe-lines on the shared memory system. The process networks that must be executed in a pipe-line parallel fashion may be annotated as shown below:

$$\begin{aligned} \text{pforce } \langle \text{take } x \bar{o}, \text{take } x \bar{a} \rangle \text{ where } & \bar{o} = 0 : \text{map_2 } (+) \bar{o} \bar{b} \\ & \bar{a} = 0 : \bar{b} \\ & \bar{b} = 1 : \text{map_2 } (+) \bar{a} \bar{b} \end{aligned}$$

The *pforce* primitive takes the same steps as before, which in this case means that both the expressions *take* $x \bar{o}$ and *take* $x \bar{a}$ will be fully evaluated in parallel. The variable x gives the number of elements of each stream that we wish to be evaluated, which should be 3000 in the case of the Fibonacci-sum example. See the column *stream length* in Table 1 for the values required by the other network programs.

The *pforce* expression above behaves as a pipe-line that computes the streams \bar{a} and \bar{o} in parallel. Figure 7 shows the slightly simplified configurations of the graphs that arise during the first two graph reduction steps taken by the two processors. The processors are notionally separated by the dotted line, there is no physical separation as the processors use a shared heap. The boxes represent suspended

computations, all other nodes represent data. Applications of *hd* and *tl* (see the definitions of *map* and *take* in Figure 4) have been omitted to avoid clutter.

Initially, both Processor 1 and Processor 2 are in a state whereby the next action will be to evaluate the suspended computations *map_2 (+) ...*. Processor 1 requests input from Processor 2. This is indicated by the pointer that crosses the dotted line, and which points at the shared object 1. Both processors will be able to use this shared object without synchronisation because it is data. After some time, both processors will have progressed to the state shown as Step 2 in Figure 7. At this stage, both processors will start to evaluate the suspended applications of *+*. It is now apparent that for the pipe-line parallel system to work properly, a locking/blocking mechanism is required.

To support pipe-lines, a low overhead locking/blocking mechanism has been built into the runtime system to avoid two or more processors from reducing the same sub-graph. The locking mechanism uses the *XMEM* machine instruction to read a memory location and to replace its contents immediately with a known lock value. The hardware implements this as an atomic transaction. Should the lock thus accessed be unavailable, then the processor requesting the lock will block, which is implemented as a busy-wait until the locked object becomes available.

Mapping the two parallel pipe-line processes onto a system with two processors is straight forward, but as the number of processes exceeds the number of processors, a suitable static mapping of pipe-line processes onto processors is hard to find. We have tried several static mappings, using both fine grained and course grained parallel execution. The best results are shown in the column marked *pipe/orig.* of Table 1. We do not expect a dynamically scheduled pipe-line to be much better, because measurements have shown that the statically scheduled pipe-line spends at most 8% of its time busy waiting. A dynamically scheduled pipe-line requires some execution time of its own, so it may be at most 8% better than a statically scheduled pipe-line.

Because a low overhead locking mechanism has been combined with the best possible static process schedule, the figures in Table 1 represent the best speed up for a statically scheduled pipe-line implementation. The master-slave implementation is consistently faster than the pipe-line. The parts of the programs that run sequentially on the master slave system offer some additional parallelism for the pipe line system. This effect is particularly strong on the *fft* program. The overhead of the required locking mechanism is apparently not compensated by the extra parallelism available to the pipe-line.

It is possible to conceive a pipe-line mechanism that does not require locks (Kelly, 1989). Such a mechanism will not be able to exploit more parallelism than the master-slave implementation combined with communication-lifting. The advantage of a real pipe-line is lost. Moreover, we expect that such an implementation will result in more overhead than our implementation, because a complex synchronisation mechanism is required, that exchanges normalised stream elements between the pipe processes. Communication lifting essentially extracts the synchronisation mechanism at compile time, resulting in the iteration of the function *nextstate*.

This function contains the knowledge of the communication pattern, while *iterate* provides the synchronisation in the most efficient way.

8 Related work

Many languages have been developed to support programming with process networks, which gives an indication of the importance of work in this area. *Lucid* (Ashcroft and Wadge, 1977) is one of the first languages based on the notion that variables represent not a single value, but a potentially infinite history of values. This notion is also central to the work on *Esterel* (Berry and Cosserat, 1984), *Signal* (Gautier *et al.*, 1987), *Lustre* (Caspi *et al.*, 1987) and others. Languages such as these offer special operators to manipulate the histories, whereby the aim has often been to make programs look like more conventional programs by hiding the history character of the variables involved.

The language that bears most resemblance to our work is *Lustre* (Caspi *et al.*, 1987). The differences lie in the realisation: the *Lustre* implementation is conventional in the sense that it uses a special purpose compiler. Our synchronous process network language is a true subset of a standard lazy functional language, and thus requires no special purpose compiler. However, to achieve a good sequential performance for the synchronous networks embedded in a lazy functional program we use program transformation techniques. Program annotations are used to achieve speedup through parallel evaluation of components of the process networks.

The core of *Lustre* (its data and sequence operators) is equivalent to the synchronous process network language, as defined in Figure 3. To illustrate this point consider the implementation of the *Lustre* data operators *+* and its four sequence operators *pre*, *→*, *when* and *current* using only the four stream functions (*:*), *tl*, *map_n* and *iterate_n* as defined in Figure 4.

Each *Lustre* sequence represents a stream of values, with which a clock is associated. A clock can be thought of as a stream of boolean values. The *zip* of the stream of values and the associated clock is thus a sensible representation of a *Lustre* sequence:

$$\begin{aligned} \text{clock } \bar{x} &= \text{map-1 } c \bar{x} \\ &\text{where} \\ c \ x &= \langle \text{True}, x \rangle \end{aligned}$$

The *Lustre* sequence operators *pre* and *→* correspond to (*:*) and *tl* respectively:

$$\begin{aligned} \text{pre } \bar{x} &= \langle \text{True}, \perp \rangle : \bar{x} \\ \bar{x} \rightarrow \bar{y} &= \text{hd } \bar{x} : \text{tl } \bar{y} \end{aligned}$$

A *Lustre* data operator applies some function to the elements of sequences. The *+* operator for example performs pairwise addition of the elements of two input sequences. The *Lustre* semantics specify that the two sequences must be *on the same clock*, so that additions will only take place when the clocks of both input streams are *True* (see the first clause of the function *p* below). When both clocks are *False*, an undefined value $\langle \text{False}, \perp \rangle$ appears in the output stream (second clause of *p*). Should the sequences be on different clocks, a semantic error is produced (last

clause of p).

$$\begin{aligned} \bar{x} + \bar{y} &= \text{map_2 } p \ \bar{x} \ \bar{y} \\ &\text{where} \\ p \ \langle \text{True}, x \rangle \ \langle \text{True}, y \rangle &= \langle \text{True}, x + y \rangle \\ p \ \langle \text{False}, x \rangle \ \langle \text{False}, y \rangle &= \langle \text{False}, \perp \rangle \\ p \ x \ y &= \perp \end{aligned}$$

All other Lustre data operators can be expressed in a similar way using $\text{map_}n$ and an appropriate auxiliary function. The Lustre sampling operator when can also be implemented using $\text{map_}2$:

$$\begin{aligned} \bar{e} \ \text{when} \ \bar{b} &= \text{map_2 } w \ \bar{e} \ \bar{b} \\ &\text{where} \\ w \ \langle \text{True}, e \rangle \ \langle \text{True}, \text{True} \rangle &= \langle \text{True}, e \rangle \\ w \ e \ b &= \langle \text{False}, \perp \rangle \end{aligned}$$

The implementation of the Lustre projection operator current requires the use of tl and iterate_1 . The state maintained by iterate_1 is used to remember the last stream element, so that this can be inserted in place of stream elements with a False clock value. The tl is necessary to remove the first, irrelevant stream element that is produced by iterate_1 :

$$\begin{aligned} \text{current } \bar{y} &= \text{tl } (\text{iterate_1 } c \ \langle \text{True}, \perp \rangle \ \bar{y}) \\ &\text{where} \\ c \ s \ \langle \text{True}, e \rangle &= \langle \text{True}, e \rangle \\ c \ s \ \langle \text{False}, e \rangle &= s \end{aligned}$$

The Lustre operators can thus all be implemented using the four stream functions (\cdot) , tl , $\text{map_}n$ and $\text{iterate_}n$ as defined in Figure 4. The functional forms of Lustre can all be implemented in Miranda without difficulty.

When viewed as a high level optimisation technique, communication lifting is related to deforestation (Wadler, 1988; Gill *et al.*, 1993). Restricted to lists, deforestation removes the need for intermediate list structure. An expression such as $\text{map_1 } f \ (\text{map_1 } g \ \bar{x})$ is transformed into $\text{map_1 } (f.g) \ \bar{x}$. For deforestation it is essential, that the producer (here $\text{map_1 } g$) and the consumer (here $\text{map_1 } f$) of an intermediate list can be identified. Communication lifting operates on any set of lists and does not require such lists to be in a producer-consumer relationship. Instead communication lifting requires the lists to be manipulated in a synchronous fashion. Communication lifting is thus supplementary to deforestation.

9 Conclusions

Many languages have been developed to support programming with synchronous streams (Esterel, Signal, Lustre etc.). This indicates how important programming with synchronous networks is as a technique for developing practical applications.

The disadvantage of developing a new language is that it also requires a new implementation to be built. The approach we take is to build process networks

using a subset of a standard lazy functional language. As an example we show that this subset is equivalent to the special purpose stream-language Lustre.

When a large number of streams is involved, which will be often the case in practical applications, the cost of lazy evaluation is high. Therefore we developed a program transformation called *communication lifting*, which takes a synchronous process network consisting of n streams into a network with only a single stream carrying n -tuples.

Measurements show that for three out of four test programs, managing the single stream of n -tuples is cheaper than managing the original n streams. The performance of the fourth program stays within 5% of the untransformed version. This result indicates that the use of communication lifting in sequential applications is worthwhile.

We have also investigated the possibility to evaluate the components of the n -tuples, resulting from communication lifting, in parallel. This gives rise to a simple master/slave kind of parallelism, which we have implemented on a four processor shared memory machine.

Comparing this master/slave implementation to the conventional way of parallelising process networks, in a pipe-line fashion, shows that communication lifting outperforms a pipe-line implementation which uses an optimal static schedule.

The communication lifting transformation has been specified formally. This makes it possible to prove the correctness of the transformation (see appendix) and to implement communication lifting as an automatic tool. Annotation by the programmer is necessary to indicate which set of streams must be transformed.

Acknowledgements

We thank Marcel Beemster, Andy Gravell, Koen Langendoen, Henk Muller and the two referees for their comments on a draft version of the paper. Andy Gravell suggested the derivation of the communication lifting transformation in section 2.

References

- L. Allison. *A practical introduction to denotational semantics*. Cambridge Univ. Press, Cambridge, England, 1986.
- E. A. Ashcroft and W. W. Wadge. Lucid, a non procedural language with iteration. *Communications ACM*, 20(7):519–526, Jul 1977.
- G. Berry and L. Cosserat. The ESTEREL synchronous programming language and its mathematical semantics. In S. D. Brookes, A. W. Roscoe, and G. Winksel, editors, *Seminar on concurrency, LNCS 197*, pages 389–448, Pittsburgh, Pennsylvania, Jul 1984. Springer-Verlag, Berlin.
- R. S. Bird and P. L. Wadler. *Introduction to functional programming*. Prentice Hall, New York, 1988.
- R. M. Burstall and J. Darlington. A transformation system for developing recursive programs. *Journal ACM*, 24(1):44–67, Jan 1977.
- P. Caspi, D. Pilaud, N. Halbwachs, and J. A. Plaice. LUSTRE: A declarative language for programming synchronous systems. In *14th Principles of programming languages*, pages 178–188, Munich, West Germany, Jan 1987. ACM.

- J. Darlington, A. J. Field, P. G. Harrison, D. Harper, G. K. Jouret, P. H. J. Kelly, K. M. Sephton, and D. W. Sharp. Structured parallel functional programming. In H. W. Glaser and P. H. Hartel, editors, *3rd Implementation of functional languages on parallel architectures*, pages 31–51, Southampton, England, Jun 1991. CSTR 91-07, Dept. of Electr. and Comp. Sci, Univ. of Southampton, England.
- A. B. Ferguson and P. L. Wadler. When will deforestation stop? In C. Hall, R. J. M. Hughes, and J. T. O'Donnell, editors, *Functional programming*, pages 39–56, Rothesay, Isle of Bute, Scotland, Aug 1988. Research report 89/R4, Dept. of Comp. Sci, Univ. of Glasgow, Scotland.
- T. Gautier, P. le Gueric, and L. Besnard. SIGNAL: A declarative language for synchronous programming of real-time systems. In G. Kahn, editor, *3rd Functional programming languages and computer architecture, LNCS 274*, pages 257–277, Portland, Oregon, Sep 1987. Springer-Verlag, Berlin.
- A. Gill, J. Launchbury, and S. L. Peyton Jones. A short cut to deforestation. In *6th Functional programming languages and computer architecture*, pages 223–232, Copenhagen, Denmark, Jun 1993. ACM.
- P. H. Hartel, H. W. Glaser, and J. M. Wild. Compilation of functional languages using flow graph analysis. *Software—practice and experience*, 24(2):127–173, Feb 1994.
- P. H. Hartel and W. G. Vree. Arrays in a lazy functional language – a case study: the fast Fourier transform. In G. Hains and L. M. R. Mullin, editors, *2nd Arrays, functional languages, and parallel systems (ATABLE)*, pages 52–66. Publication 841, Dept. d'informatique et de recherche opérationelle, Univ. de Montréal, Canada, Jun 1992.
- J. Jeuring. *Theories for algorithm calculation*. PhD thesis, Dept. of Comp. Sci, Univ. of Utrecht, The Netherlands, 1992.
- G. Kahn. The semantics of a simple language for parallel programming. In J. L. Rosenfeld, editor, *Information processing*, pages 471–475, Stockholm, Sweden, Aug 1974. North Holland.
- P. H. J. Kelly. *Functional programming for loosely-coupled multiprocessors*. Pitman publishing, London, England, 1989.
- K. G. Langendoen and P. H. Hartel. FCG: a code generator for lazy functional languages. In U. Kastens and P. Pfahler, editors, *Compiler construction (CC), LNCS 641*, pages 278–296, Paderborn, Germany, Oct 1992. Springer-Verlag, Berlin.
- Z. Manna, S. Ness, and J. E. Vuillemin. Inductive methods for proving properties of programs. *Communications ACM*, 16(8):491–502, Aug 1973.
- S. L. Peyton Jones. *The implementation of functional programming languages*. Prentice Hall, Englewood Cliffs, New Jersey, 1987.
- D. A. Turner. Miranda: A non-strict functional language with polymorphic types. In J.-P. Jouannaud, editor, *2nd Functional programming languages and computer architecture, LNCS 201*, pages 1–16, Nancy, France, Sep 1985. Springer-Verlag, Berlin.
- W. G. Vree. *Design considerations for a parallel reduction machine*. PhD thesis, Dept. of Comp. Sys, Univ. of Amsterdam, Dec 1989.
- J. E. Vuillemin. *Proof techniques for recursive programs*. PhD thesis, Comp. Sci Dept, Stanford Univ., Oct 1973. Technical report STAN-CS-73-393.
- P. L. Wadler. Deforestation: Transforming programs to eliminate trees. In H. Ganzinger, editor, *European symp. on programming (ESOP 88), LNCS 300*, pages 344–358, Nancy, France, Mar 1988. Springer-Verlag, Berlin.

Appendix - Correctness of the transformation

To establish the correctness of the Rules $T0$ to $T5$, a few auxiliary lemmas are used, which are shown below. The zip-, and map-lemmas can easily be proved by complete

induction on i . The proofs of the remaining lemmas are given. It is essential for all these proofs that streams are infinite lists; the stream elements may assume any value, including \perp .

The zip-lemma

Given $\forall n \in \mathbb{N} \wedge \bar{x}_1, \dots, \bar{x}_n \in \bar{D}$, then:

$$\forall i \in \mathbb{N} : (\text{zip_}n \langle \bar{x}_1, \dots, \bar{x}_n \rangle) ! i = \langle (\bar{x}_1 ! i), \dots, (\bar{x}_n ! i) \rangle$$

This can be proved by induction on i .

The map-lemma

Given $\forall \bar{x}_1 \dots \bar{x}_n \in \bar{D} \wedge f \in F$, then:

$$\forall i \in \mathbb{N} : (\text{map_}n f \bar{x}_1 \dots \bar{x}_n) ! i = f (\bar{x}_1 ! i) \dots (\bar{x}_n ! i)$$

This can be proved by induction on i .

The iterate-lemma

Given $\forall \bar{x}_1 \dots \bar{x}_n \in \bar{D} \wedge f \in F \wedge s \in D$, then:

$$\begin{aligned} \forall i \in \mathbb{N} : (\text{iterate_}n f s \bar{x}_1 \dots \bar{x}_n) ! i &= \bar{y} ! i \\ &\text{where} \\ \bar{y} &= s : \text{map_}m f \bar{y} \bar{x}_1 \dots \bar{x}_n \\ m &= n + 1 \end{aligned}$$

To prove this lemma, consider the special case that $n = 1$. Then given $\bar{x} \in \bar{D}$ and the binary function $f \in F$, the iterate-lemma is:

$$\begin{aligned} \forall i \in \mathbb{N} : (\text{iterate_}1 f s \bar{x}) ! i &= \bar{y} ! i \\ &\text{where} \\ \bar{y} &= s : \text{map_}2 f \bar{y} \bar{x} \end{aligned}$$

The proof is by induction on i . Base case:

$$\begin{aligned} &(\text{iterate_}1 f s \bar{x}) ! 0 \\ &= \text{hd} (\text{iterate_}1 f s \bar{x}) && \text{(unfold !)} \\ &= \text{hd} (s : \text{iterate_}1 f (f s (\text{hd} \bar{x})) (tl \bar{x})) && \text{(unfold iterate_}1) \\ &= s && \text{(unfold hd)} \\ &= \bar{y} ! 0 && \text{(new definition, fold !)} \\ &\text{where} \\ &\bar{y} = s : \text{map_}2 f \bar{y} \bar{x} \quad \square \end{aligned}$$

Induction step:

$$\begin{aligned}
& (\text{iterate_1 } f \ s \ \bar{x})!(i+1) \\
= & \ (tl \ (\text{iterate_1 } f \ s \ \bar{x}))!i && \text{(unfold !)} \\
= & \ (tl \ (s : \text{iterate_1 } f \ (f \ s \ (hd \ \bar{x})) \ (tl \ \bar{x})))!i && \text{(unfold iterate_1)} \\
= & \ (\text{iterate_1 } f \ (f \ s \ (hd \ \bar{x})) \ (tl \ \bar{x}))!i && \text{(unfold tl)} \\
= & \ \bar{y}!i && \text{(hypothesis)} \\
& \text{where} \\
& \ \bar{y} = f \ s \ (hd \ \bar{x}) : \text{map_2 } f \ \bar{y} \ (tl \ \bar{x}) \\
= & \ \bar{y}!i && \text{(fold map_2)} \\
& \text{where} \\
& \ \bar{y} = \text{map_2 } f \ (s : \bar{y}) \ (hd \ \bar{x} : tl \ \bar{x}) \\
= & \ (s : \bar{y})!(i+1) && \text{(fold !, hd and tl)} \\
& \text{where} \\
& \ (s : \bar{y}) = s : \text{map_2 } f \ (s : \bar{y}) \ \bar{x} \\
= & \ \bar{z}!(i+1) && \text{(define } \bar{z} = s : \bar{y}\text{)} \\
& \text{where} \\
& \ \bar{z} = s : \text{map_2 } f \ \bar{z} \ \bar{x} \quad \square
\end{aligned}$$

A similar proof can be given of the general case for $n \geq 2$.

Corollary of the iterate-lemma

Given $\forall \bar{x}_1 \dots \bar{x}_n \in \bar{D} \wedge f \in F \wedge s \in D \wedge m = n + 1$, then:

$$\begin{aligned}
& \bar{y} = s : \text{map_m } f \ \bar{y} \ \bar{x}_1 \dots \bar{x}_n \\
\equiv & && \text{(iterate-lemma)} \\
\forall i \in \mathbb{N} : & \ \bar{y}!i = (\text{iterate_n } f \ s \ \bar{x}_1 \dots \bar{x}_n)!i \\
\equiv & && \text{(element-lemma)} \\
& \bar{y} = \text{iterate_n } f \ s \ \bar{x}_1 \dots \bar{x}_n
\end{aligned}$$

The element-lemma

Given $\forall \bar{x}, \bar{y} \in \bar{D}$ then:

$$\bar{x} = \bar{y} \equiv \forall i \in \mathbb{N} : \bar{x}!i = \bar{y}!i$$

To prove the element-lemma an auxiliary result is needed to give the correspondence between list comprehensions and the function *take*, so that the take-lemma (Bird and Wadler, 1988, Pages 182-183) can be used. Given $\bar{x} \in \bar{D}$, then:

$$\forall i \in \mathbb{N} : \text{take } i \ \bar{x} = [\bar{x}!k \mid k \leftarrow [0..i-1]]$$

The proof of the auxiliary lemma is by induction on i . Base case:

$$\begin{aligned}
& \text{take } 0 \ \bar{x} \\
= & \ [] && \text{(unfold take)} \\
= & \ [\bar{x}!k \mid k \leftarrow []] && \text{(list comprehension)} \\
= & \ [\bar{x}!k \mid k \leftarrow [0..-1]] \quad \square && \text{(arithmetic)}
\end{aligned}$$

Induction step:

$$\begin{aligned}
& \text{take } (i+1) \ \bar{x} \\
= & \ hd \ \bar{x} : \text{take } i \ (tl \ \bar{x}) && \text{(unfold take)} \\
= & \ hd \ \bar{x} : [\bar{x}!k \mid k \leftarrow [0..i-1]] && \text{(hypothesis)} \\
= & \ hd \ \bar{x} : [\bar{x}!k \mid k \leftarrow [0..i-1]] && \text{(} k = h - 1\text{)} \\
= & \ \bar{x}!0 : [\bar{x}!h \mid h \leftarrow [1..i]] && \text{(fold !)} \\
= & \ [\bar{x}!h \mid h \leftarrow [0..i]] \quad \square && \text{(list comprehension)}
\end{aligned}$$

The unfold step is permitted because streams are infinite lists and not partial lists.

The element-lemma can now be proved as follows:

$$\begin{aligned}
& \bar{x} = \bar{y} \\
\equiv & && \text{(take-lemma)} \\
\forall i \in \mathbb{N} : & \ \text{take } i \ \bar{x} = \text{take } i \ \bar{y} \\
\equiv & && \text{(auxiliary lemma)} \\
\forall i \in \mathbb{N} : & \ [\bar{x}!k \mid k \leftarrow [0..i-1]] = [\bar{y}!k \mid k \leftarrow [0..i-1]] \\
\equiv & && \text{(list comprehension)} \\
\forall i \in \mathbb{N} : & \ \bar{x}!i = \bar{y}!i \quad \square
\end{aligned}$$

The stream-lemma

When given $\bar{z}_1 \dots \bar{z}_h \in \bar{D} \wedge z_1 \dots z_h \in D \wedge A \in F \wedge \bar{A} \in \bar{F}$ and the following 4 conditions are met:

- (i) $\bar{A} \bar{z}_1 \dots \bar{z}_h = \bar{e}_0[\bar{z}_1 \dots \bar{z}_h, \bar{a}_1 \dots \bar{a}_m]$
where
 $\bar{a}_1 = \bar{e}_1[\bar{z}_1 \dots \bar{z}_h, \bar{a}_1 \dots \bar{a}_m]$
 \vdots
 $\bar{a}_m = \bar{e}_m[\bar{z}_1 \dots \bar{z}_h, \bar{a}_1 \dots \bar{a}_m]$
- (ii) $A z_1 \dots z_h = e_0[z_1 \dots z_h, a_1 \dots a_m]$
where
 $a_1 = e_1[z_1 \dots z_h, a_1 \dots a_m]$
 \vdots
 $a_m = e_m[z_1 \dots z_h, a_1 \dots a_m]$
- (iii) there are no free occurrences of either \bar{A} or A in $\bar{e}_0 \dots \bar{e}_m$ or $e_0 \dots e_m$
- (iv) $\forall i \in \mathbb{N} \wedge 0 \leq j \leq m$ we have:
 $(\bar{e}_j[\bar{z}_1 \dots \bar{z}_h, \bar{a}_1 \dots \bar{a}_m])^i = e_j[z_1^i/z_1 \dots \bar{z}_h^i/z_h, \bar{a}_1^i/a_1 \dots \bar{a}_m^i/a_m]$

Then the stream-lemma asserts that:

- (v) $\bar{A} = \text{map_h } A$

In the stream-lemma and its proof some special notation will be used. For an expression e , in which the variables $a_1 \dots a_m$ occur free we write $e[a_1 \dots a_m]$ and the notation $e[b_1/a_1 \dots b_m/a_m]$ is an expression e in which the free variables $a_1 \dots a_m$ are simultaneously replaced by respectively $b_1 \dots b_m$. For brevity we use a superscript in the proof to denote projection on tuples rather than the function sel_m :

$$\forall 1 \leq k \leq m : \langle \bar{a}_1, \dots, \bar{a}_m \rangle^k = \bar{a}_k$$

We begin the proof by associating a function \bar{f}_j with each of the expressions $\bar{e}_j[\bar{z}_1 \dots \bar{z}_h, \bar{a}_1 \dots \bar{a}_m]$ in (i), such that:

$$\forall 0 \leq j \leq m : \bar{f}_j = \lambda u t. \bar{e}_j[u^1/\bar{z}_1 \dots u^h/\bar{z}_h, t^1/\bar{a}_1 \dots t^m/\bar{a}_m]$$

Thus $\bar{f}_j \langle \bar{z}_1, \dots, \bar{z}_h \rangle \langle \bar{a}_1, \dots, \bar{a}_m \rangle = \bar{e}_j$. In the same way associate functions $f_0 \dots f_m$ with the expressions $e_0 \dots e_m$.

As the second step rewrite (iv) in terms of the functions $\bar{f}_0 \dots \bar{f}_m$ and $f_0 \dots f_m$.

$\forall i \in \mathbb{N} \wedge 0 \leq j \leq m :$

$$(\bar{e}_j[\bar{z}_1 \dots \bar{z}_h, \bar{a}_1 \dots \bar{a}_m])^i = e_j[z_1^i/z_1 \dots \bar{z}_h^i/z_h, \bar{a}_1^i/a_1 \dots \bar{a}_m^i/a_m]$$

\equiv (use the definitions of $f_0 \dots f_m$ and $\bar{f}_0 \dots \bar{f}_m$)

$$(\bar{f}_j \langle \bar{z}_1, \dots, \bar{z}_h \rangle \langle \bar{a}_1, \dots, \bar{a}_m \rangle)^i = f_j \langle \bar{z}_1^i, \dots, \bar{z}_h^i \rangle \langle \bar{a}_1^i, \dots, \bar{a}_m^i \rangle$$

\equiv (substitute $z = \langle \bar{z}_1, \dots, \bar{z}_h \rangle$ and $a = \langle \bar{a}_1, \dots, \bar{a}_m \rangle$ and use the zip-lemma)

$$(vi) \quad (\bar{f}_j z a)^i = f_j ((zip_h z)^i) ((zip_m a)^i)$$

The third step is to reformulate the definition of \bar{A} from given (i):

$$\begin{aligned} (vii) \quad \bar{A} \bar{z}_1 \dots \bar{z}_h &= \bar{f}_0 \langle \bar{z}_1, \dots, \bar{z}_h \rangle \langle \bar{a}_1, \dots, \bar{a}_m \rangle && \text{(given (i))} \\ &\text{where} \\ \bar{a}_1 &= \bar{f}_1 \langle \bar{z}_1, \dots, \bar{z}_h \rangle \langle \bar{a}_1, \dots, \bar{a}_m \rangle \\ &\vdots \\ \bar{a}_m &= \bar{f}_m \langle \bar{z}_1, \dots, \bar{z}_h \rangle \langle \bar{a}_1, \dots, \bar{a}_m \rangle \\ &= \bar{f}_0 \langle \bar{z}_1, \dots, \bar{z}_h \rangle \langle \bar{a}_1, \dots, \bar{a}_m \rangle && \text{(tupling)} \\ &\text{where} \\ \langle \bar{a}_1, \dots, \bar{a}_m \rangle &= \langle \bar{f}_1 \langle \bar{z}_1, \dots, \bar{z}_h \rangle \langle \bar{a}_1, \dots, \bar{a}_m \rangle, \dots, \bar{f}_m \langle \bar{z}_1, \dots, \bar{z}_h \rangle \langle \bar{a}_1, \dots, \bar{a}_m \rangle \rangle \\ &= \text{(abstraction and fixed point)} \\ &\bar{f}_0 \langle \bar{z}_1, \dots, \bar{z}_h \rangle (fix(\lambda a. \langle \bar{f}_1 \langle \bar{z}_1, \dots, \bar{z}_h \rangle a, \dots, \bar{f}_m \langle \bar{z}_1, \dots, \bar{z}_h \rangle a \rangle)) \\ &= \bar{B} \langle \bar{z}_1, \dots, \bar{z}_h \rangle && \text{(introduce } \bar{B}, \bar{E} \text{)} \\ &\text{where} \\ \bar{B} &= \lambda z. \bar{f}_0 z \bar{E}[z] \\ \bar{E}[z] &= fix(\lambda a. \langle \bar{f}_1 z a, \dots, \bar{f}_m z a \rangle) \end{aligned}$$

Rewrite definition (ii) of A in a similar way to yield:

$$\begin{aligned} (viii) \quad A z_1 \dots z_h &= B \langle z_1, \dots, z_h \rangle \\ &\text{where} \\ B &= \lambda z. f_0 z E[z] \\ E[z] &= fix(\lambda a. \langle f_1 z a, \dots, f_m z a \rangle) \end{aligned}$$

Omitting the *where* expression of B (viii) looks like this:

$$\begin{aligned} A z_1 \dots z_h &= B \langle z_1, \dots, z_h \rangle \\ \equiv & \text{(Let } \bar{z}_1 \dots \bar{z}_h \in \bar{D} \text{ and use the element-lemma)} \end{aligned}$$

$\forall i \in \mathbb{N} :$

$$A (\bar{z}_1^i) \dots (\bar{z}_h^i) = B \langle \bar{z}_1^i, \dots, \bar{z}_h^i \rangle$$

\equiv

(zip-lemma)

$\forall i \in \mathbb{N} :$

$$A (\bar{z}_1^i) \dots (\bar{z}_h^i) = B ((zip_h \langle \bar{z}_1, \dots, \bar{z}_h \rangle)^i)$$

\equiv

(map- and element-lemma)

$$\text{map_h } A \bar{z}_1 \dots \bar{z}_h = \text{map_l } B (\text{zip_h} \langle \bar{z}_1, \dots, \bar{z}_h \rangle)$$

Compare this to (vii), which relates \bar{A} and \bar{B} :

$$\bar{A} \bar{z}_1 \dots \bar{z}_h = \bar{B} \langle \bar{z}_1, \dots, \bar{z}_h \rangle$$

To prove the stream-lemma (v) it thus remains to show that:

$$\begin{aligned} \bar{B} \langle \bar{z}_1, \dots, \bar{z}_h \rangle &= \text{map_1 } B \text{ (zip_h } \langle \bar{z}_1, \dots, \bar{z}_h \rangle) \\ \equiv & \hspace{15em} \text{(element-lemma)} \end{aligned}$$

$$\begin{aligned} \forall i \in \mathbb{N} : \\ (\bar{B} \langle \bar{z}_1, \dots, \bar{z}_h \rangle) \bar{l}_i &= (\text{map_1 } B \text{ (zip_h } \langle \bar{z}_1, \dots, \bar{z}_h \rangle)) \bar{l}_i \\ \equiv & \hspace{15em} \text{(map-lemma)} \end{aligned}$$

$$\begin{aligned} \forall i \in \mathbb{N} : \\ (\bar{B} \langle \bar{z}_1, \dots, \bar{z}_h \rangle) \bar{l}_i &= B \text{ ((zip_h } \langle \bar{z}_1, \dots, \bar{z}_h \rangle) \bar{l}_i) \end{aligned}$$

Using the definitions of \bar{B} and B we can derive the following fixed point equation for $(\text{zip_m } \bar{E}[z]) \bar{l}_i$:

$$\begin{aligned} \forall i \in \mathbb{N} : \\ (\text{zip_m } \bar{E}[z]) \bar{l}_i &= (\text{zip_m } (\text{fix}(\lambda a. \langle \bar{f}_1 z a, \dots, \bar{f}_m z a \rangle))) \bar{l}_i && \text{(unfold } \bar{E}) \\ &= (\text{zip_m } \langle \bar{f}_1 z \bar{E}[z], \dots, \bar{f}_m z \bar{E}[z] \rangle) \bar{l}_i && \text{(unfold fix)} \\ &= \langle (\bar{f}_1 z \bar{E}[z]) \bar{l}_i, \dots, (\bar{f}_m z \bar{E}[z]) \bar{l}_i \rangle && \text{(zip-lemma)} \\ &= && \text{(by (vi))} \\ &\langle \bar{f}_1 \text{ ((zip_h } z) \bar{l}_i) \text{ ((zip_m } \bar{E}[z]) \bar{l}_i), \dots, \bar{f}_m \text{ ((zip_h } z) \bar{l}_i) \text{ ((zip_m } \bar{E}[z]) \bar{l}_i) \rangle && \\ &= && \text{(abstraction)} \\ &(\lambda a. \langle \bar{f}_1 \text{ ((zip_h } z) \bar{l}_i) a, \dots, \bar{f}_m \text{ ((zip_h } z) \bar{l}_i) a \rangle) \text{ ((zip_m } \bar{E}[z]) \bar{l}_i) \\ &= \text{fix}(\lambda a. \langle \bar{f}_1 \text{ ((zip_h } z) \bar{l}_i) a, \dots, \bar{f}_m \text{ ((zip_h } z) \bar{l}_i) a \rangle) && \text{(fixed point)} \\ &= E[(\text{zip_h } z) \bar{l}_i / z] && \text{(fold } E) \end{aligned}$$

So for all h -tuples z we also have:

$$(ix) \forall i \in \mathbb{N} : (\text{zip_m } \bar{E}[z]) \bar{l}_i = E[(\text{zip_h } z) \bar{l}_i / z]$$

The proof of the stream-lemma can now be completed as follows:

$$\begin{aligned} (\bar{B} \langle \bar{z}_1, \dots, \bar{z}_h \rangle) \bar{l}_i &= ((\lambda z. \bar{f}_0 z \bar{E}[z]) \langle \bar{z}_1, \dots, \bar{z}_h \rangle) \bar{l}_i && \text{(unfold } \bar{B}) \\ &= (\bar{f}_0 \langle \bar{z}_1, \dots, \bar{z}_h \rangle \bar{E}[\langle \bar{z}_1, \dots, \bar{z}_h \rangle]) \bar{l}_i && \text{(reduction)} \\ &= \bar{f}_0 \text{ ((zip_h } \langle \bar{z}_1, \dots, \bar{z}_h \rangle) \bar{l}_i) \text{ ((zip_m } \bar{E}[\langle \bar{z}_1, \dots, \bar{z}_h \rangle]) \bar{l}_i) && \text{(by (vi))} \\ &= \bar{f}_0 \text{ ((zip_h } \langle \bar{z}_1, \dots, \bar{z}_h \rangle) \bar{l}_i) E[(\text{zip_h } \langle \bar{z}_1, \dots, \bar{z}_h \rangle) \bar{l}_i / z] && \text{(by (ix))} \\ &= (\lambda z. \bar{f}_0 z \bar{E}[z]) \text{ ((zip_h } \langle \bar{z}_1, \dots, \bar{z}_h \rangle) \bar{l}_i) && \text{(abstraction)} \\ &= B \text{ ((zip_h } \langle \bar{z}_1, \dots, \bar{z}_h \rangle) \bar{l}_i) \quad \square && \text{(fold } B) \end{aligned}$$

Correctness of the simplifying transformations

The correctness of each of the Rules $T1$, $T3$ and $T5$ will now be proved. The correctness of the composition of these transformations then follows, because each

represents a total function ($p \rightarrow p$). The partial correctness of the clauses in Figure 6 will be proved first. Then the termination of the transformations will be established.

Partial correctness of the simplifying transformations

Using the auxiliary lemmas, the correctness of Clause 4a can be proved as follows. Define $\forall n \in \mathbb{N} \wedge f \in F \wedge s \in D \wedge x_1 \dots x_n \in D$:

$$\begin{aligned} \bar{y} &= s : \text{map_m } f \bar{y} \bar{x}_1 \dots \bar{x}_n \\ m &= n + 1 \end{aligned}$$

Then:

$$\bar{v} = \text{iterate_n } f s \bar{x}_1 \dots \bar{x}_n$$

$$\equiv \hspace{15em} \text{(element-lemma)}$$

$$\forall i \in \mathbb{N} : \bar{v} \bar{l}_i = (\text{iterate_n } f s \bar{x}_1 \dots \bar{x}_n) \bar{l}_i$$

$$\equiv \hspace{15em} \text{(iterate-lemma)}$$

$$\forall i \in \mathbb{N} : \bar{v} \bar{l}_i = f (\bar{y} \bar{l}_i) (\bar{x}_1 \bar{l}_i) \dots (\bar{x}_n \bar{l}_i)$$

$$\equiv \hspace{15em} \text{(map-lemma)}$$

$$\forall i \in \mathbb{N} : \bar{v} \bar{l}_i = (\text{map_m } f \bar{y} \bar{x}_1 \dots \bar{x}_n) \bar{l}_i$$

$$\equiv \hspace{15em} \text{(element-lemma)}$$

$$\bar{v} = \text{map_m } f \bar{y} \bar{x}_1 \dots \bar{x}_n$$

The fact that for all $n, f, s, \bar{x}_1 \dots \bar{x}_n$ the stream \bar{v} prior to transformation is equal to that after transformation establishes the correctness of Clause 4a.

The proofs for the remaining clauses present no difficulties. The Clauses 1, 2b, 3, 4b and 5e by themselves make no changes. Clause 2a introduces a new equation, which is equivalence preserving. An unfold in recursive equations such as those under consideration here is always equivalence preserving (Manna *et al.*, 1973). Rule $T5$ uses unfolds and introduces new definitions: Clause 5a unfolds the definition of map_n as given in Figure 4 and introduces new equations; Clauses 5b and 5c unfold the definitions of \bar{w} and tl . For Clause 5b this gives the equation $\bar{o} = \bar{x}$ and for (5c) we have $\bar{v} = \bar{x}$. These equations can be eliminated by renaming \bar{x} to \bar{o} and \bar{v} to \bar{x} respectively. Clause 5d unfolds the definitions of map_n , \bar{w} and tl and introduces new equations. Rule $T5$ is thus equivalence preserving, which concludes the proofs of the partial correctness of the simplifying transformations.

Termination of the simplifying transformations

All simplifying transformations terminate because a bound can be given for the number of times each individual clause is applied.

Clauses 1, 3 and 5e are each applied once. The number of function applications in the original process network is an upper bound for number of times Clauses 2a, 2b, 4a or 4b is applied. Clause 5a will be applied at most once. The only clause of Rule $T5$ that matches an equation of the form $\bar{o} = \dots$ is (5c), which will never replace that equation by one of the form that can be matched again by (5a).

To derive an upper bound on the number of times Clauses 5b, 5c or 5d may be called, let the number of applications of tl and map_n be t respectively m . The way the labelling of the equations is created and maintained guarantees that Clause 5d will be applied at most $m \times t$ times. The upper bound on the number of times Clauses 5b or 5c are applied is given by the number of tl applications, which are either present originally, or introduced by Clause 5a or 5d. This number is bounded because there is an upper bound on the number of times (5d) is applied.

Summarising, we have now established the fact that the simplifying transformations preserve equivalence and terminate.

left and right hand side of Clause 0a	stream-lemma functions \bar{A} and A
$\begin{aligned} & \bar{A} \bar{s} \bar{u}_1 \dots \bar{u}_h \\ & = \text{zip_k} \langle \bar{x}_1, \dots, \bar{x}_k \rangle \\ & \quad \text{where} \\ & \quad \bar{v}_1 = \text{sel_k } 1 \text{ (unzip_k } \bar{s}) \\ & \quad \vdots \\ & \quad \bar{v}_k = \text{sel_k } k \text{ (unzip_k } \bar{s}) \\ & \quad \bar{w}_1 = \text{map_n}_1 f_1 \bar{y}_{11} \dots \bar{y}_{1n_1} \\ & \quad \vdots \\ & \quad \bar{w}_m = \text{map_n}_m f_m \bar{y}_{m1} \dots \bar{y}_{mn_m} \end{aligned}$	$\begin{aligned} & \bar{A} \bar{z}_1 \dots \bar{z}_h \\ & = \bar{e}_0 \\ & \quad \text{where} \\ & \quad \bar{a}_1 = \bar{e}_1 \\ & \quad \vdots \\ & \quad \bar{a}_k = \bar{e}_k \\ & \quad \bar{a}_{k+1} = \bar{e}_{k+1} \\ & \quad \vdots \\ & \quad \bar{a}_{k+m} = \bar{e}_{k+m} \end{aligned}$
$\begin{aligned} & A s u_1 \dots u_h \\ & = \langle x_1, \dots, x_k \rangle \\ & \quad \text{where} \\ & \quad v_1 = \text{sel_k } 1 s \\ & \quad \vdots \\ & \quad v_k = \text{sel_k } k s \\ & \quad w_1 = f_1 y_{11} \dots y_{1n_1} \\ & \quad \vdots \\ & \quad w_m = f_m y_{m1} \dots y_{mn_m} \end{aligned}$	$\begin{aligned} & A z_1 \dots z_h \\ & = e_0 \\ & \quad \text{where} \\ & \quad a_1 = e_1 \\ & \quad \vdots \\ & \quad a_k = e_k \\ & \quad a_{k+1} = e_{k+1} \\ & \quad \vdots \\ & \quad a_{k+m} = e_{k+m} \end{aligned}$

Fig. 8. Structural correspondence between the stream-lemma and $T0$.

Correctness of the communication lifting transformation

The termination proof of Rule $T0$ (Figure 5) is immediate. To prove the partial correctness of Rule $T0$, the $(:)$ and map_n equations on the left hand side of Rule $T0$ are tupled; the tl equations are left unchanged. The guard in Clause 0a ensures that the tl equations are independent of the remaining equations. Therefore it is correct to consider the network without the tl equations. From now on let $\bar{v}_1 = \bar{o}$.

Then:

$$\left. \begin{aligned} \bar{v}_1 &= s_1 : \bar{x}_1 \\ &\vdots \\ \bar{v}_k &= s_k : \bar{x}_k \\ \bar{w}_1 &= \text{map_n}_1 f_1 \bar{y}_{11} \dots \bar{y}_{1n_1} \\ &\vdots \\ \bar{w}_m &= \text{map_n}_m f_m \bar{y}_{m1} \dots \bar{y}_{mn_m} \end{aligned} \right\} \begin{aligned} & \langle \bar{v}_1, \dots, \bar{v}_k \rangle \\ &= \langle s_1 : \bar{x}_1, \dots, s_k : \bar{x}_k \rangle \\ & \quad \text{where} \\ & \bar{w}_1 = \text{map_n}_1 f_1 \bar{y}_{11} \dots \bar{y}_{1n_1} \\ & \quad \vdots \\ & \bar{w}_m = \text{map_n}_m f_m \bar{y}_{m1} \dots \bar{y}_{mn_m} \end{aligned}$$

The function composition $\text{unzip_k} \circ \text{zip_k}$ is the identity function on a tuple of streams, thus the left hand side of Clause 0a is equivalent to:

$$\begin{aligned} \langle \bar{v}_1, \dots, \bar{v}_k \rangle &= \text{unzip_k} (\text{zip_k} \langle s_1 : \bar{x}_1, \dots, s_k : \bar{x}_k \rangle)^\dagger \\ &\equiv \end{aligned} \tag{unfold zip_k}$$

$$\langle \bar{v}_1, \dots, \bar{v}_k \rangle = \text{unzip_k} (\langle s_1, \dots, s_k \rangle : (\text{zip_k} \langle \bar{x}_1, \dots, \bar{x}_k \rangle))^\dagger$$

Next a new equation \bar{s} for the stream of k -tuples is introduced, and the free variables of the network $\bar{u}_1 \dots \bar{u}_h$ are made explicit. This yields the following set of equations, again equivalent to the left hand side of Clause 0a:

$$\begin{aligned} \langle \bar{v}_1, \dots, \bar{v}_k \rangle &= \text{unzip_k } \bar{s} \\ \bar{s} &= \langle s_1, \dots, s_k \rangle : (\text{zip_k} \langle \bar{x}_1, \dots, \bar{x}_k \rangle)^\dagger \\ &\equiv \tag{make unzip_k equation local} \\ \bar{s} &= \langle s_1, \dots, s_k \rangle : (\text{zip_k} \langle \bar{x}_1, \dots, \bar{x}_k \rangle)^\dagger \\ & \quad \text{where} \\ & \quad \langle \bar{v}_1, \dots, \bar{v}_k \rangle = \text{unzip_k } \bar{s} \\ &\equiv \tag{introduce function } \bar{A} \text{ and make free variables explicit)} \\ \bar{s} &= \langle s_1, \dots, s_k \rangle : (\bar{A} \bar{s} \bar{u}_1 \dots \bar{u}_h) \\ \bar{A} \bar{s} \bar{u}_1 \dots \bar{u}_h &= \text{zip_k} \langle \bar{x}_1, \dots, \bar{x}_k \rangle^\dagger \\ & \quad \text{where} \\ & \quad \langle \bar{v}_1, \dots, \bar{v}_k \rangle = \text{unzip_k } \bar{s} \end{aligned}$$

In the last step of this derivation, we have used the fact that the states $s_1 \dots s_k$ are independent of the stream variables $\bar{v}_1 \dots \bar{v}_k$ and $\bar{w}_1 \dots \bar{w}_m$ (see Figure 5).

As shown in Figure 8, the final step brings the system of equations in a form that fits the structural requirements of the stream-lemma. The function \bar{A} as derived from the equations on the left hand side of Clause 0a is shown to the left, while the corresponding elements of the definitions for the stream-lemma are shown to the right.

The stream-lemma states that $\bar{A} = \text{map_h } A$, provided stream-lemma condition (iii) holds $\forall i \in \mathbb{N} \wedge 0 \leq j \leq k + m$. This is verified as follows:

For $j = 0$ we have that $\bar{e}_0 \equiv \text{zip_k} \langle \bar{x}_1, \dots, \bar{x}_k \rangle$ and also that $e_0 \equiv \langle x_1, \dots, x_k \rangle$. This

\dagger The *where* equations for $w_1 \dots w_m$ have been omitted.

can be proved as follows:

$$\begin{aligned}
& \forall i \in \mathbb{N} : \\
& (\text{zip_k } \langle \bar{x}_1, \dots, \bar{x}_k \rangle) \text{!i} \\
& = \langle \bar{x}_1 \text{!i}, \dots, \bar{x}_k \text{!i} \rangle \quad (\text{zip-lemma}) \\
& = \langle x_1, \dots, x_k \rangle [\bar{x}_1 \text{!i} / x_1, \dots, \bar{x}_k \text{!i} / x_k] \quad \square
\end{aligned}$$

For $1 \leq j \leq k$ we have that $\bar{e}_j \equiv \text{sel_k } j \text{ (unzip_k } \bar{s})$ and also $e_j \equiv \text{sel_k } j \text{ s}$ because:

$$\begin{aligned}
& \forall i \in \mathbb{N} : \\
& (\text{sel_k } j \text{ (unzip_k } \bar{s})) \text{!i} \\
& = (\text{sel_k } j \text{ (unzip_k (zip_k } \langle s_1 : \bar{x}_1, \dots, s_k : \bar{x}_k \rangle))) \text{!i} \quad (\text{unfold } \bar{s}) \\
& = (\text{sel_k } j \langle s_1 : \bar{x}_1, \dots, s_k : \bar{x}_k \rangle) \text{!i} \quad (\text{identity}) \\
& = (s_j : \bar{x}_j) \text{!i} \quad (\text{unfold sel_k}) \\
& = (\text{sel_k } j \langle (s_1 : \bar{x}_1) \text{!i}, \dots, (s_k : \bar{x}_k) \text{!i} \rangle) \quad (\text{fold sel_k}) \\
& = \text{sel_k } j \langle (\text{zip_k } \langle s_1 : \bar{x}_1, \dots, s_k : \bar{x}_k \rangle) \text{!i} \rangle \quad (\text{zip-lemma}) \\
& = \text{sel_k } j \text{ (}\bar{s}\text{!i)} \quad (\text{fold } \bar{s}) \\
& = (\text{sel_k } j \text{ s}) [\bar{s} \text{!i} / s] \quad \square
\end{aligned}$$

For $k+1 \leq k+j \leq k+m$ we have $\bar{e}_{k+j} \equiv \text{map_n}_j \text{ } f_j \text{ } \bar{y}_{j1} \dots \bar{y}_{jn_j}$ and $e_{k+j} \equiv f_j \text{ } y_{j1} \dots y_{jn_j}$ since:

$$\begin{aligned}
& \forall i \in \mathbb{N} : \\
& (\text{map_n}_j \text{ } f_j \text{ } \bar{y}_{j1} \dots \bar{y}_{jn_j}) \text{!i} \\
& = f_j (\bar{y}_{j1} \text{!i}) \dots (\bar{y}_{jn_j} \text{!i}) \quad (\text{map-lemma}) \\
& = (f_j \text{ } y_{j1} \dots y_{jn_j}) [\bar{y}_{j1} \text{!i} / y_{j1}, \dots, \bar{y}_{jn_j} \text{!i} / y_{jn_j}] \quad \square
\end{aligned}$$

Given the definition of *nextstate* as in Figure 5, (which is the same as *A* here) we have:

$$\begin{aligned}
\bar{s} & = \langle s_1, \dots, s_k \rangle : (\bar{A} \bar{s} \bar{u}_1 \dots \bar{u}_h) \\
& \equiv \quad (\text{stream-lemma}) \\
\bar{s} & = \langle s_1, \dots, s_k \rangle : (\text{map_h } A \bar{s} \bar{u}_1 \dots \bar{u}_h) \\
& \equiv \quad (A = \text{nextstate} \text{ and corollary of iterate-lemma}) \\
\bar{s} & = \text{iterate_h nextstate } \langle s_1, \dots, s_k \rangle \bar{u}_1 \dots \bar{u}_h \quad \square
\end{aligned}$$

This completes the correctness proof of transformation Rule *T0*.