

**LABORATORY FOR
COMPUTER SCIENCE**



**MASSACHUSETTS
INSTITUTE OF
TECHNOLOGY**

Semantics of pH: A parallel dialect of Haskell

Computation Structures Group Memo 377
June 4, 1995

Shail Aditya Arvind Jan-Willem Maessen
MIT Laboratory for Computer Science
{shail,arvind,earwig}@lcs.mit.edu

Lennart Augustsson
Chalmers University
augustss@cs.chalmers.se

Rishiyur S. Nikhil
DEC Cambridge Research Laboratory
nikhil@crl.dec.com

The research described in this paper was funded in part by the Advanced Research Projects Agency of the Department of Defense under Office of Naval Research contract N00014-92-J-1310.

545 TECHNOLOGY SQUARE, CAMBRIDGE, MASSACHUSETTS 02139

Semantics of pH: A parallel dialect of Haskell*

Shail Aditya^a
MIT

Arvind^a
MIT

Lennart Augustsson^b
Chalmers University

Jan-Willem Maessen^a
MIT

Rishiyur S. Nikhil^c
DEC, CRL

Abstract

The semantics of kernel pH are defined in the form of a parallel, normalizing interpreter. A description of I-structure and M-structure operators is also given within the same framework. Semantics of barriers in pH are presented by translation into the kernel language without barriers. The framework presented is also suitable for multithreaded compilation of pH.

1 What is pH?

pH [11] is a parallel variant of the Haskell programming language [8] with extensions for loops, synchronized side-effect operations, and explicit sequentialization. This paper discusses the operational semantics of these variations and extensions. The concrete syntax of pH is presented in the preliminary pH manual [11] and will track future Haskell versions.

There is more than one approach to parallel implementations of functional languages. For example, it is possible to exploit the implicit parallelism of a Haskell program by concurrent evaluation of the arguments of strict operators. The absence of side effects ensures that this concurrent evaluation cannot change the result. Strictness analysis techniques widen the scope of this idea, by allowing parallel evaluation of any strict function argument. Further, it is possible to provide programmer annotations to indicate sub-expressions which should be evaluated in parallel, even when the compiler cannot prove that the value of all the sub-expressions will be required. Unlike similar annotations in imperative languages, the annotations can only affect termination, and not the results (if any). All these approaches take *demand-driven* evaluation as the starting point: in the absence of annotations, an expression is evaluated only if its value is required.

An alternative approach is to use a *parallel* evaluation order which reduces all redexes in parallel, not only those whose value is required. This strategy also implements non-strict semantics, provided that no redex has its evaluation delayed indefinitely. A parallel evaluation strategy is followed by the Id programming language [10], in which any redex not in the body of a conditional or lambda abstraction is reduced. pH follows the same strategy. A brief description of pH may be given as follows.

*The research described in this paper was funded in part by the Advanced Research Projects Agency of the Department of Defense under Office of Naval Research contract N00014-92-J-1310.

^aMIT Laboratory for Computer Science, Cambridge, MA 02139. Email: {shail,arvind,earwig}@lcs.mit.edu.

^bDepartment of Computer Science, Chalmers University, Göteborg, Sweden. Email: augustss@cs.chalmers.se.

^cCambridge Research Laboratory, Digital Equipment Corporation, One Kendall Square, Building 700, Cambridge, MA 02139. Email: nikhil@crl.dec.com.

pH = Haskell syntax and type system + Id evaluation strategy + Id side-effect operators

In practice, implementations of pH may not guarantee that every redex is reduced eventually. This will imply that some pH programs may fail to give a result and terminate when they would do so in “traditional” Haskell. However, if pH gives a proper (non-error, non-bottom) result for a particular program then so does “traditional” Haskell, and these results are the same.

1.1 Why pH?

pH is an attempt to bring together the lazy functional community (as represented by Haskell) and the dataflow community (as represented by Id and Sisal [7]). It is hoped that by sharing a common base language, it will be easier to share ideas and implementations and freely exchange programs. By doing so, programmers need only master a single syntax and type system. There are important differences between pH and Haskell, however. By choosing Haskell as a framework, it becomes easier to isolate and therefore understand the differences between the two.

1.2 Structure of pH

pH is a layered language. pH(F) is the purely-functional subset of pH. Its main addition to Haskell is for- and while-loops. For example, one could compute the sum of the integers between 1 and n as follows.

Example 1:

```
let sum = 0
in for i <- [1..n] do
    next sum = sum + i
finally sum
```

Loops are purely functional, and can easily be translated into tail-recursive functions. The eager semantics of pH permits tail-recursion to be implemented more efficiently than lazy semantics.

pH(I) adds I-structures [4] to pH(F). I-structures are *single-assignment* data structures that allow fine-grain producer-consumer synchronization among parallel tasks. pH(I) preserves determinacy, *i.e.*, *confluence*, although the language is no longer referentially transparent.

pH(M) adds M-structures [6] to pH(I). M-structures are *multiple-assignment* data structures that allow mutual exclusion synchronization. M-structures introduce side-effects and non-determinacy. pH(M) also provides the ability to group statements for sequential execution using *barriers* in order to avoid unwanted race conditions among side-effect operations.

Full pH is the same as pH(M).

1.3 Outline

The rest of this paper concentrates on semantic issues in which pH differs from Haskell. Section 2 describes the parallel evaluation order of pH in the context of a kernel language and its multi-threaded interpreter. In Section 3 we extend the kernel language with I-structure and M-structure operations. Section 4 defines the semantics of barriers by translating them into the kernel language using explicit termination signals from side-effect operations. Finally, Section 5 concludes with some notes about current implementations.

c	\in	Constant
$f, t, x, y, z \dots$	\in	Identifier
SE, X, Y, Z, \dots	\in	Simple Expression
E	\in	Expression
PF^n	\in	Primitive Fn. with n arguments
CF^n	\in	Data Constructor with n arguments
S	\in	Statement
Constant	$::=$	$Integer \mid Float \mid Boolean \mid Nil$
SE	$::=$	Identifier \mid Constant
PF^1	$::=$	$hd \mid tl \mid select_k$
PF^2	$::=$	$+ \mid - \mid \dots \mid < \mid > \mid \dots$
CF^2	$::=$	Cons
CF^n	$::=$	make_n_tuple
E	$::=$	$SE \mid PF^n(x_1, \dots, x_n) \mid CF^n(x_1, \dots, x_n)$ $\mid \lambda x. E \mid case(x, E_1, \dots, E_n)$ $\mid ap(f, z) \mid sap(f, z) \mid Block$
Block	$::=$	$\{ S \text{ in } x \}$
S	$::=$	$\epsilon \mid x = E \mid S_1; \dots; S_n$
Program	$::=$	Block

Figure 1: The kernel pH language.

2 Parallel Evaluation Order

pH follows an *eager* evaluation strategy: all tasks execute in parallel, restricted only by the data dependencies among them. This strategy automatically exposes large amounts of parallelism both within and across procedures. This is in contrast with a *lazy* evaluation strategy followed by Haskell: only those tasks are evaluated which are required to produce the result. This strategy imposes a sequential constraint on the overall computation, although the exact ordering of tasks is decided dynamically.

In this section we describe the eager evaluation model of pH by giving a parallel interpreter for the kernel language of pH. Our interpreter is based on a parallel abstract machine, which is in the spirit of the G-machine and its later variants [9, 5, 12]. We first describe the kernel language (Section 2.1) and then our parallel abstract machine (Section 2.2) and its instruction set (Section 2.3). It is followed by a description of the interpreter (Section 2.4).

2.1 The Kernel pH Language

The abstract syntax of the kernel pH language is shown in Figure 1. The kernel language ensures that every intermediate result of a complex expression is explicitly named using an identifier. This is convenient for expressing and preserving the sharing of subexpressions within a computation [3]. In order to simplify the description of the evaluation rules in this paper, we do not allow constants to appear as arguments to primitive functions.

Aside from the usual arithmetic primitives, the kernel language provides functions for constructing and selecting elements of lists and n -ary tuples. The language also provides n -ary `case`

ι	\in	Instruction	
ℓ	\in	Location	
ν	\in	Value	::= Constant Location $\langle \lambda x. E, \rho \rangle$
ρ	\in	Environment Frame	= Identifier \rightarrow Location
σ	\in	Store	= Location \rightarrow $\{ \langle \text{full } \nu \rangle \mid \langle \text{defer } \delta s \rangle \}$
ω	\in	Thread	= Code \times Environment $\#$
δ	\in	Suspension	= Thread
δs	\in	Suspensions	= List(Thread)
ωs	\in	Work Queue	= List(Thread)
		Accumulator	= Value
ιs	\in	Code	= List(Instruction)
ρs	\in	Environment	= List(Environment Frame)
		Processor	= Accumulator \times Code \times Environment
		Machine	= List(Processor) \times Work Queue \times Store

State	Processor			Global Memory	
	Accumulator	Code	Environment	Work Queue	Store
Initial (starting proc.)	-	ιs_0	$[\]$	$[\]$	$\{ \}$
Initial (other proc.)	-	$[\]$	$[\]$	$[\]$	$\{ \}$
Final (all proc.)	-	$[\]$	$[\]$	$[\]$	σ
Error (any proc.)	storerr	ιs	ρs	ωs	σ

Where $\iota s_0 = [\text{eval}(\text{Program}), \text{print}, \text{schedule}]$

Figure 2: Dynamic entities used by the abstract machine.

expressions which select one of the branches based on the value of the dispatch identifier, nested λ -expressions which may contain free identifiers, a general function application operator **ap**, and a parallel block construct that controls lexical scoping and enables precise sharing of subexpression values. The order of bindings in a block is not significant. The identifier following the **in** keyword in a block expression denotes the result of the block.

2.2 A Parallel Abstract Machine

Our parallel abstract machine consists of a number of sequential processors connected to a global shared memory. The important features of our machine are described in Figure 2. Each memory location can hold a value or a list of suspended threads, and is tagged with its status. A full location contains a value. A newly allocated location is tagged with **defer** and it is initialized to be empty. A *value* is either a constant, an allocated store location, or a function closure. The global memory holds three types of entities: heap storage, activation frames and a *work queue* which is a list of *threads*, that is, (*code*, *environment*) pairs.

The processor state consists of code, environment, and an *accumulator* which can hold a value. Any idle processor can get work by dequeuing a thread from the work queue and loading its code and environment space from the thread. Initially, it is assumed that the whole machine is empty. A program E is started by executing the following code on some processor in the empty environment.

[eval(E), print, schedule]

The `eval` instruction interprets an expression by structural decomposition in a given environment. Environments map program identifiers to frame locations, and are structured in a stack like manner. `eval` of a block allocates a new *activation frame* in the global memory to store the values of the bound identifiers of the block. It also creates a new *environment frame* for these identifiers and pushes it on the processor environment. The bound identifiers of the block point to locations in the newly allocated activation frame. New environment frames (but not activation frames) are also created in case of λ -expressions and function applications. The sharing of computations is achieved entirely through the activation frame locations: all references to an identifier within the scope of its definition lead to the same frame location. The separation of environment frames and activation frames permits environments to be copied freely and therefore allows us to treat function closures in the same way as simple values.

At each step, the processor executes the instruction at the head of the current code sequence, modifying the processor state in the process. The instruction set of the processor has instructions to load and store the accumulator, perform arithmetic and logic operations on store locations, and suspend the current thread in case of missing values. This is in contrast with lazy functional languages where unevaluated locations contain *thunks* that compute its value on-demand. Typically an instruction is popped from the code sequence after execution, but some instructions may also add new instructions to the code sequence. This is the primary way to alter the control flow. There are also instructions to push and pop environment frames, and enqueue and dequeue threads from the work queue.

The work queue is maintained as a FIFO queue in order to guarantee fair scheduling. This queue must be manipulated atomically. It is easy to modify the machine so that each processor has its own work queue which is maintained in a FIFO manner. In case a processor goes idle, it can pick up a thread from the back of the work queue of any other processor. As long as we do not arbitrarily suspend the currently running thread on any processor, this strategy still implements a fair schedule while considerably reducing the contention due to the atomic manipulation of the work queue.

The overall state of our parallel machine is described by the state of all the processors and the global memory and the work queue. Figure 2 also shows the initial, the final and the error states of the processors of our machine. The machine starts by scheduling an initial sequence of instructions on some processor that evaluates the program, prints the result present in the accumulator¹, and then continues to schedule the remaining work present in the work queue. This emphasizes the non-strict nature of this evaluation model: the program may print a result before final termination. The machine halts when both the code sequence and the work queue are empty. The machine is said to be *deadlocked* when there are suspended computations in the store but there is no work to be done. If there are no suspended computations, then the machine is said to have *terminated normally*. An error may be generated during execution due to exceptional conditions such as a type mismatch, an arithmetic overflow, or out of bounds access to data structures. However, if an attempt is made to write into a memory location that is already full then the whole machine is said to have reached an error state as shown in Figure 2.

2.3 Instruction Set

Figure 3 shows the instructions used by our interpreter. We will describe the semantics of these instructions in terms of a state transition involving a 5-tuple $(\nu, \iota s, \rho s, \omega s, \sigma)$. The first three

¹The main thread is the only place where a print instruction may appear.

<code>eval(E)</code>	:: Evaluate E in the current environment and leave the result in the accumulator
-----------------------------------	--

<code>touch(ℓ)</code>	:: Check if the location ℓ is full, if not suspend the current thread
<code>rtouch(ℓ)</code>	:: Check if the location ℓ is full, if not suspend the current thread including the <code>rtouch</code> instruction. (Thus, <code>rtouch</code> can be retried when the thread is activated again).
<code>load(ℓ)</code>	:: Load the accumulator with the value present in the location ℓ (the location must be full)
<code>take(ℓ)</code>	:: Load the accumulator with the value present in the location ℓ and mark the location empty (the location must be full)
<code>loadi_k(ℓ)</code>	:: Load the accumulator with the value in the location at an offset k from the location pointed to by ℓ , or suspend the thread if the location is empty
<code>store(ℓ)</code>	:: Store the value in the accumulator in the location ℓ and reactivate any suspensions
<code>storerr(ℓ)</code>	:: Generate an error if the location ℓ is full
<code>PFn(ℓ_1, \dots, ℓ_n)</code>	:: Apply strict primitive operator PF^n (such as $+$) to the values at locations ℓ_1, \dots, ℓ_n (the locations must be full)
<code>switch($\ell_x, \iota s_1, \dots, \iota s_n$)</code>	:: Switch to the branch indexed by the value at location ℓ_x
<code>pushenv(ρ)</code>	:: Push the environment frame ρ onto the environment stack
<code>popenv</code>	:: Pop the top-level environment frame from the environment stack
<code>schedule</code>	:: Schedule the next thread from the work queue
<code>print</code>	:: Print the accumulator contents

Figure 3: The instruction set of the parallel abstract machine.

elements of the tuple refer to the accumulator carrying a value ν , the code sequence ιs , and the environment ρs on the processor where the instruction is executed. The last two elements of the tuple are the work queue ωs , and the dynamic store σ which are shared amongst all processors.

All the instructions that modify the global store must be executed atomically. That is, while a processor reads and modifies a given location within the store, no other processor should be allowed to read or write that location. Transactions on different locations may proceed in parallel. In case where an instruction sequence ιs has to be executed atomically, we will indicate that by writing *atomic*(ιs).

Touch Instructions

The `touch` instruction tests the status of a location. If the status is full, it does nothing. Otherwise, (status is *defer*) the code sequence following the current instruction is suspended at that location along with the current environment. The `rtouch` instruction is similar except that the `touch` operation is retried.

touch:

$$\frac{\langle _ \text{touch}(\ell) : \iota s \ \rho s \ \omega s \ \sigma[\ell \mapsto \langle \text{full } \nu \rangle] \rangle}{\langle _ \ \iota s \ \rho s \ \omega s \ \sigma \rangle}$$

$$\frac{\langle _ \text{touch}(\ell) : \iota s \ \rho s \ \omega s \ \sigma[\ell \mapsto \langle \text{defer } \delta s \rangle] \rangle}{\langle _ \text{[schedule]} \ \rho s \ \omega s \ \sigma[\ell \mapsto \langle \text{defer } (\iota s, \rho s) : \delta s \rangle] \rangle}$$

rtouch:

$$\frac{\langle _ \text{rtouch}(\ell) : \iota s \ \rho s \ \omega s \ \sigma[\ell \mapsto \langle \text{full } \nu \rangle] \rangle}{\langle _ \ \iota s \ \rho s \ \omega s \ \sigma \rangle}$$

$$\frac{\langle _ \text{rtouch}(\ell) : \iota s \ \rho s \ \omega s \ \sigma[\ell \mapsto \langle \text{defer } \delta s \rangle] \rangle}{\langle _ \text{[schedule]} \ \rho s \ \omega s \ \sigma[\ell \mapsto \langle \text{defer } (\text{rtouch}(\ell) : \iota s, \rho s) : \delta s \rangle] \rangle}$$

Note that these instructions do not start the evaluation of the expression that will fill the location being touched. In lazy evaluation this expression is always known and a thunk for it can be stored in this location. In that case, the above instructions can be easily generalized to enqueue the thunk into the work queue.

Load and Store Instructions

The load instruction loads the contents of a given location into the accumulator. The location must already be full. The take instruction is similar to load except that it leaves the location empty.

load:

$$\frac{\langle _ \text{load}(\ell) : \iota s \ \rho s \ \omega s \ \sigma[\ell \mapsto \langle \text{full } \nu \rangle] \rangle}{\langle \nu \ \iota s \ \rho s \ \omega s \ \sigma \rangle}$$

take:

$$\frac{\langle _ \text{take}(\ell) : \iota s \ \rho s \ \omega s \ \sigma[\ell \mapsto \langle \text{full } \nu \rangle] \rangle}{\langle \nu \ \iota s \ \rho s \ \omega s \ \sigma[\ell \mapsto \langle \text{defer } [] \rangle] \rangle}$$

loadi_k is an indexed, indirect load instruction. It desugars into touching an indexed location and then loading its value into the accumulator.

loadi_k:

$$\frac{\langle _ \text{loadi}_k(\ell_x) : \iota s \ \rho s \ \omega s \ \sigma[\ell_x \mapsto \langle \text{full } \ell \rangle] \rangle}{\langle _ \text{touch}(\ell + k) : \text{load}(\ell + k) : \iota s \ \rho s \ \omega s \ \sigma \rangle}$$

The store instruction stores the contents of the accumulator into the given location. It also reactivates any suspensions waiting in the location by enqueueing them in the work queue.

store:

$$\frac{\langle \nu \ \text{store}(\ell) : \iota s \ \rho s \ \omega s \ \sigma[\ell \mapsto \langle \text{defer } \delta s \rangle] \rangle}{\langle \nu \ \iota s \ \rho s \ \omega s ++ \delta s \ \sigma[\ell \mapsto \langle \text{full } \nu \rangle] \rangle}$$

Storing in a location that is already full is regarded as an error. The storerr instruction tests the status of a location and produces an error if it is already full. This instruction can precede a store instruction to avoid multiple stores to the same location.

storerr:

$$\frac{\langle - \text{storerr}(\ell) : \iota s \ \rho s \ \omega s \ \sigma [\ell \mapsto \langle \text{full } \nu \rangle] \rangle}{\langle \text{storerr} \ \iota s \ \rho s \ \omega s \ \sigma \rangle}$$

$$\frac{\langle - \text{storerr}(\ell) : \iota s \ \rho s \ \omega s \ \sigma [\ell \mapsto \langle \text{defer } \delta s \rangle] \rangle}{\langle - \ \iota s \ \rho s \ \omega s \ \sigma \rangle}$$

Arithmetic and Logic Instructions

The standard arithmetic and logic primitives are straightforward. All of them leave the result in the accumulator.

+:

$$\frac{\langle - \ +(\ell_1, \ell_2) : \iota s \ \rho s \ \omega s \ \sigma [\ell_1 \mapsto \langle \text{full } \underline{m} \rangle, \ell_2 \mapsto \langle \text{full } \underline{n} \rangle] \rangle}{\langle \underline{m+n} \ \iota s \ \rho s \ \omega s \ \sigma \rangle}$$

Switch Instruction

The switch instruction selects one of its branches based on the value present in the location being dispatched upon.

switch ($1 \leq m \leq n$):

$$\frac{\langle - \ \text{switch}(\ell_x, \iota s_1, \dots, \iota s_n) : \iota s \ \rho s \ \omega s \ \sigma [\ell_x \mapsto \langle \text{full } \underline{m} \rangle] \rangle}{\langle - \ \iota s_m \ ++ \ \iota s \ \rho s \ \omega s \ \sigma \rangle}$$

Environment Manipulation Instructions

The pushenv instructions pushes the given environment frame onto the local environment stack of a processor, while the popenv instruction pops one frame from the top. The contents of the accumulator are not disturbed while popping.

pushenv:

$$\frac{\langle - \ \text{pushenv}(\rho) : \iota s \ \rho s \ \omega s \ \sigma \rangle}{\langle - \ \iota s \ \rho : \rho s \ \omega s \ \sigma \rangle}$$

popenv:

$$\frac{\langle \nu \ \text{popenv} : \iota s \ \rho : \rho s \ \omega s \ \sigma \rangle}{\langle \nu \ \iota s \ \rho s \ \omega s \ \sigma \rangle}$$

Thread Scheduling Instruction

The execution of a schedule instruction on a processor signifies either the termination or the suspension of a thread. The processor schedules a new thread from the head of the work queue or starts *idling* if the work queue is empty. The machine halts when all processors are idling.

schedule:

$$\frac{\langle - \ [\text{schedule}] \ \rho s \ (\iota s', \rho s') : \omega s \ \sigma \rangle}{\langle - \ \iota s' \ \rho s' \ \omega s \ \sigma \rangle}$$

$$\frac{\langle - \ [\text{schedule}] \ \rho s \ [] \ \sigma \rangle}{\langle - \ [] \ \rho s \ [] \ \sigma \rangle}$$

2.4 Expression Evaluation

In this section we describe the `eval` instruction which is used to destructure and evaluate expressions. It leaves the value of the expression in the accumulator.

Evaluation of a constant simply loads that constant into the accumulator.

eval constant:

$$\frac{\langle _ \text{eval}(c) : \iota s \ \rho s \ \omega s \ \sigma \rangle}{\langle c \ \iota s \ \rho s \ \omega s \ \sigma \rangle} \quad \checkmark$$

Evaluation of an identifier first touches it and then loads its value in the accumulator.

eval identifier:

$$\frac{\langle _ \text{eval}(x) : \iota s \ \rho s \ \omega s \ \sigma \rangle}{\langle _ \text{touch}(\rho s(x)) : \text{load}(\rho s(x)) : \iota s \ \rho s \ \omega s \ \sigma \rangle} \quad \checkmark$$

In case of a strict primitive function, we touch all its arguments before executing the primitive application.

eval strict PF:

$$\frac{\langle _ \text{eval}(PF^n(x_1, \dots, x_n)) : \iota s \ \rho s \ \omega s \ \sigma \rangle}{\langle _ \text{touch}(\rho s(x_1)) : \dots : \text{touch}(\rho s(x_n)) : PF^n(\rho s(x_1), \dots, \rho s(x_n)) : \iota s \ \rho s \ \omega s \ \sigma \rangle} \quad \checkmark$$

All data selector functions such as `hd`, `tl`, and `select_k` are directly implemented using the indexed, indirect load instruction `loadi_k`.

The non-strict constructor `Cons` allocates a pair of empty locations in the store and immediately returns the pointer to the first location as its value. It also pushes work into the work queue to evaluate the arguments of `Cons` and to fill these locations eagerly. The n -ary tuple constructor `make_n_tuple` operates in the same fashion. Note that under lazy evaluation we would have stored thunks into these locations which would be evaluated on demand.

eval `Cons`:

$$\frac{\langle _ \text{eval}(\text{Cons}(x_1, x_2)) : \iota s \ \rho s \ \omega s \ \sigma \rangle}{\langle \ell \ \iota s \ \rho s \ \omega s \ ++ [\omega_1, \omega_2] \ \sigma' \rangle} \quad \checkmark$$

$$\begin{aligned} \text{where } \sigma' &= \sigma + \{ \ell \mapsto \langle \text{defer } [] \rangle, (\ell + 1) \mapsto \langle \text{defer } [] \rangle \} \\ \omega_1 &= ([\text{eval}(x_1), \text{store}(\ell), \text{schedule}], \rho s) \\ \omega_2 &= ([\text{eval}(x_2), \text{store}(\ell + 1), \text{schedule}], \rho s) \end{aligned}$$

λ -expressions are evaluated to a closure value. The closure records the locations for the free identifiers of the λ -body in a new environment frame.

eval λ -expression:

$$\frac{\langle _ \text{eval}(\lambda x. E) : \iota s \ \rho s \ \omega s \ \sigma \rangle}{\langle \langle \lambda x. E, \rho \rangle \ \iota s \ \rho s \ \omega s \ \sigma \rangle} \quad \checkmark$$

$$\begin{aligned} \text{where } y_1, \dots, y_m &= FV(\lambda x. E) \\ \rho &= \{ y_1 \mapsto \rho s(y_1), \dots, y_m \mapsto \rho s(y_m) \} \end{aligned}$$

Evaluation of a function application proceeds by first touching the function and then applying the resulting closure to the argument. The function argument is touched prior to application only if the function application is strict (`sap`).

eval ap:

$$\frac{\langle _ \text{eval}(\text{ap}(f, z)) : \iota s \ \rho s \ \omega s \ \sigma \rangle}{\langle _ \text{touch}(\rho s(f)) : \text{ap}(\rho s(f), \rho s(z)) : \iota s \ \rho s \ \omega s \ \sigma \rangle}$$

eval sap:

$$\frac{\langle _ \text{eval}(\text{sap}(f, z)) : \iota s \ \rho s \ \omega s \ \sigma \rangle}{\langle _ \text{touch}(\rho s(f)) : \text{touch}(\rho s(z)) : \text{ap}(\rho s(f), \rho s(z)) : \iota s \ \rho s \ \omega s \ \sigma \rangle}$$

Given a closure, the ap instruction extends the closure environment frame with the given argument location and starts the evaluation of the function body after pushing the new environment frame onto the current environment. The old environment is restored after the evaluation of the body produces a result in the accumulator².

ap:

$$\frac{\langle _ \text{ap}(\ell_f, \ell_z) : \iota s \ \rho s \ \omega s \ \sigma[\ell_f \mapsto \langle \text{full}(\lambda x. E, \rho) \rangle] \rangle}{\langle _ (\text{pushenv}(\rho') : \text{eval}(E) : \text{popenv} : \iota s \ \rho s \ \omega s \ \sigma) \rangle}$$

where $\rho' = \rho + \{ x \mapsto \ell_z \}$ [Schedule]

The evaluation of the case expression proceeds by first touching the dispatch expression. This must yield an integer within the range of the dispatch which is used to select the appropriate branch.

eval case:

$$\frac{\langle _ \text{eval}(\text{case}(x, E_1, \dots, E_n)) : \iota s \ \rho s \ \omega s \ \sigma \rangle}{\langle _ \text{touch}(\rho s(x)) : \text{switch}(\rho s(x), [\text{eval}(E_1)], \dots, [\text{eval}(E_n)]) : \iota s \ \rho s \ \omega s \ \sigma \rangle}$$

The evaluation of a block expression allocates a new activation frame in the store with a fresh location for each bound identifier of the block. It also allocates new environment frame which points to the locations in the new activation frame. This environment frame is pushed onto the current environment. Following eager evaluation, auxiliary threads are added to the work queue to evaluate each of the right-hand side expressions in the extended environment. The result identifier is also evaluated under this extended environment.

eval block:

$$\frac{\langle _ \text{eval}(\{ x_1 = E_1; \dots; x_n = E_n \text{ in } x \}) : \iota s \ \rho s \ \omega s \ \sigma \rangle}{\langle _ \text{pushenv}(\rho) : \text{eval}(x) : \text{popenv} : \iota s \ \rho s \ \omega s \ \omega s' \ \sigma' \rangle}$$

where $\rho = \{ x_1 \mapsto \ell_1, \dots, x_n \mapsto \ell_n \}$
 $\sigma' = \sigma + \{ \ell_1 \mapsto \langle \text{defer} [] \rangle, \dots, \ell_n \mapsto \langle \text{defer} [] \rangle \}$
 $\omega s' = [([\text{eval}(E_1), \text{store}(\ell_1), \text{schedule}], \rho : \rho s),$
 $\dots,$
 $([\text{eval}(E_n), \text{store}(\ell_n), \text{schedule}], \rho : \rho s)]$

It is also possible to initialize the newly allocated block locations with the thunks for their respective right-hand side expressions, and not enqueue threads into the work queue. These thunks will be enqueued when the locations are touched.

3 Side-Effect Operations

Functional programming languages purposely do not permit assignment or mutable storage. Instead, such features are often offered through various implementation tricks such as by using

²Note that restoring the environment does not imply that the function environment can be deallocated because some threads of the function body may not have terminated yet.

Constant	$::=$	\dots	$ $	\bullet
PF^1	$::=$	\dots	$ $	$\text{iAlloc} \mid \text{mAlloc} \mid \text{iFetch} \mid \text{mFetch} \mid \text{W}$
PF^2	$::=$	\dots	$ $	$\text{iStore} \mid \text{mStore} \mid \&$
S	$::=$	\dots	$ $	$(S_1 \text{ --- } S_2)$

Figure 4: Kernel language extensions for side-effect operations.

monadic programming techniques. Alternatively, side-effects may be introduced *via* the traditional assignment operator after specifying a precise sequential ordering on all operations (*e.g.*, Scheme and ML). `Id`, motivated by concerns for parallelism, offers yet another alternative for incorporating side-effects. In `Id`, it is possible to specify only a partial order on side-effect operations and still retain an overall consistent picture of the computation.

`pH` supports imperative operations on I-structure [4] and M-structure [6] objects. I-structures allow the creation of a data structure to be separated from the definition of its components: attempts to use the value of a component are automatically delayed until that component is defined; attempts to redefine a component lead to an error state. M-structures, on the other hand, are fully mutable data structures whose components can be redefined repeatedly: an `mFetch` operation reads and empties a full component; an `mStore` operation (re)defines it; two successive `mStore` operations on the same component lead to an error state unless separated by an `mFetch` operation.

For some programs in `pH` we need a way to sequentialize M-structure operations in order to avoid race conditions implied by its parallel evaluation order. This is accomplished by the use of *control regions* and *barriers* [2]. A control region is informally defined as a set of concurrent threads that are under the same control dependence and therefore always execute together. Threads within the same control region may execute in any order or in an interleaved manner as long as the data dependencies among them are respected. Barriers provide a mechanism to detect the termination of a set of parallel activities enclosed within a control region. A barrier (`---`) creates two sub-regions within a given control region — one above the barrier called the *pre-region* and the other below the barrier called the *post-region*. Intuitively, no computation within the post-region is allowed to proceed until all the side-effect computations within the pre-region have terminated. This semantics is different from those proposed in [2] where a barrier waits for the termination of *all* computations rather than just those which cause side-effects.

In the rest of this section we describe the `pH` extensions to deal with I-structure and M-structure operations. The semantics of barriers is described in Section 4.

3.1 Kernel Language Extensions

We extend the kernel language to incorporate side-effect operations and barriers as shown in Figure 4. We add primitives to allocate, read, and write I-structures and M-structures. The fetch and store operations directly address store locations: indexed addressing is handled separately in the kernel language.

We also extend the syntax for block bindings to allow barriers—two sets of parallel block bindings may now be sequentialized by using a barrier (`---`) between them. Additional constants (\bullet), primitive operators (`W` and `&`) and strict function application (`sap`) are used to translate barriers into an ordinary set of bindings.

3.2 Evaluation Rules for Side-Effect Operations

In this section we describe additional evaluation rules for internal primitive operators involving I-structures, M-structures, and barriers. All these operators are strict on all their inputs.

I-structure Instructions

The `iAlloc` instruction allocates an empty I-structure array. The `iFetch` and `iStore` instructions address an I-structure location directly; all address arithmetic must be done separately. The `iFetch` instruction is similar to `loadi.k`. The `iStore` instruction updates an empty location with a value and reactivates any suspended continuations. Storing to an already full location is considered to be an error.

`iAlloc`:

$$\frac{\langle _ \text{ iAlloc}(\ell_x) : \iota s \ \rho s \ \omega s \ \sigma[\ell_x \mapsto \langle \text{full } \underline{n} \rangle] \rangle}{\langle \ell \ \iota s \ \rho s \ \omega s \ \sigma' \rangle}$$

where $\sigma' = \sigma + \{ \ell \mapsto \langle \text{defer } [] \rangle, \dots, (\ell + n - 1) \mapsto \langle \text{defer } [] \rangle \}$

`iFetch`:

$$\frac{\langle _ \text{ iFetch}(\ell_x) : \iota s \ \rho s \ \omega s \ \sigma[\ell_x \mapsto \langle \text{full } \ell \rangle] \rangle}{\langle _ \text{ touch}(\ell) : \text{load}(\ell) : \iota s \ \rho s \ \omega s \ \sigma \rangle}$$

`iStore`:

$$\frac{\langle _ \text{ iStore}(\ell_x, \ell_z) : \iota s \ \rho s \ \omega s \ \sigma[\ell_x \mapsto \langle \text{full } \ell \rangle] \rangle}{\langle _ \text{ atomic}(\{\text{storerr}(\ell), \text{load}(\ell_z), \text{store}(\ell)\}) : \iota s \ \rho s \ \omega s \ \sigma \rangle}$$

M-structure Instructions

The `mAlloc` and `mStore` instructions are identical to the corresponding instructions for I-structures. The `mFetch` instruction is similar to the `iFetch` except that it empties the location being accessed and has to be retried if the location is already empty.

`mAlloc`:

$$\frac{\langle _ \text{ mAlloc}(\ell_x) : \iota s \ \rho s \ \omega s \ \sigma[\ell_x \mapsto \langle \text{full } \underline{n} \rangle] \rangle}{\langle \ell \ \iota s \ \rho s \ \omega s \ \sigma' \rangle}$$

where $\sigma' = \sigma + \{ \ell \mapsto \langle \text{defer } [] \rangle, \dots, (\ell + n - 1) \mapsto \langle \text{defer } [] \rangle \}$

`mFetch`:

$$\frac{\langle _ \text{ mFetch}(\ell_x) : \iota s \ \rho s \ \omega s \ \sigma[\ell_x \mapsto \langle \text{full } \ell \rangle] \rangle}{\langle _ \text{ atomic}(\{\text{rtouch}(\ell), \text{take}(\ell)\}) : \iota s \ \rho s \ \omega s \ \sigma \rangle}$$

`mStore`:

$$\frac{\langle _ \text{ mStore}(\ell_x, \ell_z) : \iota s \ \rho s \ \omega s \ \sigma[\ell_x \mapsto \langle \text{full } \ell \rangle] \rangle}{\langle _ \text{ atomic}(\{\text{storerr}(\ell), \text{load}(\ell_z), \text{store}(\ell)\}) : \iota s \ \rho s \ \omega s \ \sigma \rangle}$$

4 Semantics of M-Barriers

In this section we describe the semantics of barriers by translating them into ordinary block bindings. The idea is to systematically construct a composite termination signal from all the side-effect

	EXPRESSIONS
TE	:: Expression \rightarrow Expression
TE [<i>c</i>]	= <i>c</i> , •
TE [<i>x</i>]	= <i>x</i> , •
TE [mFetch (<i>x</i>)]	= { <i>y</i> = mFetch (<i>x</i>); in <i>y</i> , W (<i>y</i>) }
TE [mStore (<i>x</i> , <i>z</i>)]	= { <i>y</i> = mStore (<i>x</i> , <i>z</i>); in <i>y</i> , W (<i>y</i>) }
TE [$PF^n(x_1, \dots, x_n)$]	= $PF^n(x_1, \dots, x_n)$, •
TE [$\lambda x. E$]	= $\lambda x. \mathbf{TE}[E]$, •
TE [case (<i>x</i> , E_1, \dots, E_n)]	= case (<i>x</i> , TE [E_1], ..., TE [E_n])
TE [ap (<i>f</i> , <i>x</i>)]	= ap (<i>f</i> , <i>x</i>)
TE [{ <i>S</i> in <i>x</i> }]	= { S' in <i>x</i> , <i>s</i> }
where $S', s = \mathbf{TS}[S]$	
STATEMENTS	
TS [[]]	:: Statement $\rightarrow List(\text{Statement}) \times \text{Identifier}$
TS [ϵ]	= (<i>s</i> = •), <i>s</i>
TS [<i>x</i> = E]	= (<i>x</i> , <i>s</i> = TE [E]), <i>s</i>
TS [$S_1; \dots; S_n$]	= ($S'_1; \dots; S'_n; s = s_1 \& \dots \& s_n$), <i>s</i>
where $S'_i, s_i = \mathbf{TS}[S_i]$ $1 \leq i \leq n$	
TS [$S_1 \text{ --- } S_2$]	= ($S'_1;$ $f = \lambda s. \{ S'_2 \text{ in } y_1, \dots, y_m, s_2 \};$ $y_1, \dots, y_m, s'_2 = \mathbf{sap}(f, s_1)$), s'_2
where $S'_i, s_i = \mathbf{TS}[S]$ $1 \leq i \leq 2$	
$y_1, \dots, y_m = BV(S_2)$	

Figure 5: Semantics of M-barriers.

operations taking place within the pre-region of a barrier (including its child regions) which would then be used to trigger the operations present in the post-region of the barrier. The notion of signal generation and composition may be understood by looking at the translation of a simple parallel block as shown below.

$$\begin{array}{lcl}
 \mathbf{TE}[\{ x_1 = E_1; & = & \{ x_1, s_1 = \mathbf{TE}[E_1]; \\
 \dots & & \dots \\
 x_n = E_n; & & x_n, s_n = \mathbf{TE}[E_n]; \\
 \text{in } x \}] & & \text{in } x, s_1 \& \dots \& s_n \}
 \end{array}$$

An expression E_i is translated as **TE**[E_i] which dynamically returns a value (bound to x_i) along with an explicit termination signal s_i . The operator **W** is used to detect the termination of each side-effect operation (**mFetch** and **mStore**) within E_i . The value of s_i would become • as soon as all the side-effect operations in E_i have terminated. The operator **&** is then used to combine all such signals into a single signal for the whole block.

The primitive **W** operator produces a signal • when a given identifier becomes a value. The general evaluation rule for strict primitive functions would ensure that the identifier being tested

is touched before applying the W operator. Similarly, the primitive $\&$ operator is used to combine signals from two subexpressions into one signal after they have been touched.

W :

$$\frac{\langle _ \quad W(l_x) : \iota s \quad \rho s \quad \omega s \quad \sigma [l_x \mapsto \langle \text{full } \nu \rangle] \rangle}{\langle \bullet \quad \iota s \quad \rho s \quad \omega s \quad \sigma \rangle}$$

$\&$:

$$\frac{\langle _ \quad \&(l_1, l_2) : \iota s \quad \rho s \quad \omega s \quad \sigma [l_1 \mapsto \langle \text{full } \bullet \rangle, l_2 \mapsto \langle \text{full } \bullet \rangle] \rangle}{\langle \bullet \quad \iota s \quad \rho s \quad \omega s \quad \sigma \rangle}$$

The complete barrier translation appears in Figure 5. Note that tuple return values and non-refutable pattern matching is used only for clarity—these operations are directly desugared into primitive operators in the obvious way. The only non-trivial base cases for signal generation are those for the `mFetch` and `mStore` primitive operators. The `mFetch` operation is considered to have terminated when the value of the location being fetched is returned. Similarly, the `mStore` operation is considered to have terminated when it returns the value being stored.

In the translation of $S_1 \text{ --- } S_2$, note that S_2 gets protected by a λ -expression, so that it does not get evaluated until the λ -expression is applied to something. That something is the termination signal of S_1 , and we use `sap` to ensure that the application does not commence until the termination signal is available. Finally, since the bound variables of S_2 have now gone into an inner scope (in the λ -expression), we return them all and rebind them again in the outer scope.

The barrier semantics shown here are different from the semantics presented in [2] in that here we are concerned only with the termination of the side-effect operations present within the pre-region of a barrier, while the semantics presented in [2] waited for the termination of the entire computation within the pre-region.

5 Conclusion

In this paper we have presented the various semantic issues in the design of pH at which it differs from or extends the Haskell programming language. The major difference is that pH uses a parallel eager evaluation strategy as opposed to the lazy evaluation strategy of Haskell. First, we described the kernel pH language and a normalizing interpreter for it that implements this parallel evaluation order. Next, we extended the language and the interpreter with synchronizing side-effect operations, I-structures and M-structures. Finally, we described a notion of sequentialization of side-effects using barriers. We showed a systematic translation of a kernel pH program with barriers into one without barriers using primitive termination detection operators.

The readers may note that our interpreter used an explicit instruction stream rather than directly evaluating kernel language expressions. This organization allows us to define a multithreaded compilation scheme for the kernel language within the same framework. The process of compilation can be defined as simply generating all the threads statically instead of manipulating the instruction stream dynamically. It essentially gets rid of the `eval` instruction. Such a compilation scheme for the kernel language is described in [1].

Currently, there is a working implementation of the pH language using the Haskell/pH HBCC front-end (written in Haskell) from Chalmers University and the Id compiler Monsoon back-end (written in Lisp) from MIT. The HBCC front-end produces a kernel pH intermediate format that is converted into dataflow graphs and fed into the Monsoon back-end. Currently, we have exercised

this compiler with purely functional Haskell programs or pH programs derived by automatically transliterating Id programs. At MIT, we are currently working on a new pH back-end for commercial uniprocessor and multiprocessor architectures that is closer in spirit to the interpretation scheme described in this paper.

References

- [1] Shail Aditya. Normalizing Strategies for Multithreaded Interpretation and Compilation of Non-Strict Languages. CSG Memo 374, MIT Laboratory for Computer Science, Cambridge, MA 02139, May 1995.
- [2] Shail Aditya, Arvind, and Joseph E. Stoy. Semantics of Barriers in a Non-Strict, Implicitly-Parallel Language. In *Proc. Functional Programming Languages and Computer Architecture, La Jolla, CA*, Cambridge, MA 02139, June 1995. Also available as CSG Memo 367-1, MIT Lab. for Computer Sc., Cambridge, MA 02139.
- [3] Zena M. Ariola and Arvind. Properties of a First-order Functional Language with Sharing. CSG Memo 347-1, Laboratory for Computer Science, MIT, Cambridge, MA 02139, June 1994. To appear in *Theoretical Computer Science*, September 1995.
- [4] Arvind, Rishiyur S. Nikhil, and Keshav K. Pingali. I-Structures: Data Structures for Parallel Computing. *ACM Transactions on Programming Languages and Systems*, 11(4):598–632, 1989.
- [5] Lennart Augustsson and Thomas Johnsson. Parallel Graph Reduction with the $\langle \nu, G \rangle$ -machine. In *Proc. Fourth Intl. Conf. on Functional Programming Languages and Computer Architecture, London*, pages 202–213. ACM Press, September 1989.
- [6] Paul S. Barth, Rishiyur S. Nikhil, and Arvind. M-Structures: Extending a Parallel, Non-Strict, Functional Language with State. In *Proc. Functional Programming Languages and Computer Architecture*, pages 538–568. Springer-Verlag, 1991. LNCS 523.
- [7] A. P. W. Böhm, D. C. Cann, J. T. Feo, and R. R. Oldehoeft. SISAL 2.0 Reference Manual. Technical Report UCRL-MA-109098, Lawrence Livermore National Laboratory, December 1991.
- [8] Paul Hudak, Simon Peyton Jones, and Philip Wadler (editors). Report on the Programming Language Haskell: A Non-strict Purely Functional Language, Version 1.2. *ACM SIGPLAN Notices*, 27(5), May 1992.
- [9] Thomas Johnsson. Efficient Compilation of Lazy Evaluation. *Proc. ACM SIGPLAN '84 Symposium on Compiler Construction, SIGPLAN Notices*, 19(6):58–69, June 1984.
- [10] Rishiyur S. Nikhil. Id Language Reference Manual Version 90.1. Technical Report CSG Memo 284-2, Laboratory for Computer Science, MIT, Cambridge, MA 02139, July 15 1991.
- [11] Rishiyur S. Nikhil, Arvind, James Hicks, Shail Aditya, Lennart Augustsson, Jan-Willem Maessen, and Yuli Zhou. pH Language Reference Manual, Version 1.0—preliminary. CSG Memo 369, Laboratory for Computer Science, MIT, Cambridge, MA 02139, January 1995.
- [12] Simon L. Peyton Jones. Implementing lazy functional languages on stock hardware: the Spineless Tagless G-machine. *Journal of Functional Programming*, 2(2):127–202, April 1992.