



IST-510255

EmBounded

Automatic Prediction of Resource Bounds for Embedded Systems

Specific Targeted Research Project (STReP)
FET Open

D04 (WP2a): Cost Model (Revised)

Due date of deliverable: 1st January 2006

Actual submission date: 3rd March 2006

Start date of project: 1st March 2005

Duration: 36 months

Lead contractor: Ludwig-Maximilians-Universität

Revision: 1.10

Purpose: Build a low-level *formal cost model*, reflecting the costs of executing Hume Abstract Machine (HAM) expressions on one particular architecture.

Results: We present cost models for heap space, stack space and time consumption of HAM code.

Conclusion: This document presents both a formal specification of the HAM, as well as cost models for the aforementioned resources. We give a small-step, operational semantics for HAM instructions, and we define resource algebras for heap, stack and time that can be added to this operational semantics in a modular way and describe both expression level and system level resource consumption.

Project co-funded by the European Commission within the 6 th Framework Programme (2002-06)		
Dissemination Level		
PU	Public	*
PP	Restricted to other programme participants (including the Commission Services)	
RE	Restricted to a group specified by the consortium (including the Commission Services)	
CO	Confidential only for members of the consortium (including the Commission Services)	

Specification and Cost Model for the Hume Abstract Machine (Version 1.1)

Hans-Wolfgang Loidl <hwloidl@tcs.ifi.lmu.de>,
Institut für Informatik, Theoretische Informatik, Ludwig-Maximilians Universität, München
Kevin Hammond <kh@dcs.st-and.ac.uk>
School of Computing Science, Univ of St Andrews, St Andrews, Scotland

Abstract

This document is both a formal specification of the Hume Abstract Machine (HAM) semantics as well as a cost model for its time, stack and heap space consumption. The specification reflects changes made to the HAM since project start. It contains instructions to support higher-order functions and exceptions, but no explicit support for timeouts (these are still in the design phase and subject to changes). The cost models are described as resource algebras that can be freely added to the HAM specification. The latter is encoded as a 2-level, small step operational semantics.

Version 1.1: This update reflects and accounts for the Renesas M32C being used as target processor (mainly affecting the time cost model).

1 Introduction

For the execution of Hume [HM] programs we define an abstract machine, the Hume Abstract Machine (HAM). This definition is an extension of the initial design, described in [Ham03], by constructs for higher-order functions and exceptions. Additionally, we formally specify the components of the machine and its behaviour in the form of a 2-level, small-step, operational semantics and give a reference implementation of the instructions of the HAM. The operational semantics uses the approach of resource algebras, which we have developed in a previous project [ABM06], to collect information on the resource consumption during execution. The resource algebras are designed in a modular way and can be instantiated without modifying the rules of the operational semantics. We define the cost model for the HAM by giving resource algebras for stack space, heap space and time consumption. The values for stack and heap space are independent of the underlying processor. For the time information we have performed measurements of the HAM interpreter running on a PowerPC architecture.

As the specification of the abstract machine used to execute Hume programs, this document serves as interface between groups working on the compilation infrastructure for Hume. Together with the document on the Hume formal semantics (D12), and the formal description of the compilation of Hume to HAM (D3), it gives a description about the costs for executing Hume and HAM programs. This is in turn a prerequisite for developing static analyses of resource consumption. Finally, this formal specification is the basis for the work in WP7 on the certification of the resource consumption for Hume code.

The structure of this document is as follows. Section 2 describes the HAM design in general, the data structures used and the behaviour of the machine by presenting a reference implementation of the HAM instructions in pseudo-C. Section 3 describes the concept of a resource algebra and instantiates it for stack space, heap space and time consumption. Section 4 provides a formal specification of the HAM as a 2-level, small-step operational semantics. Finally, Section 5 summarises.

Comments on Version 1.1 of the cost model: This version of the HAM-level cost model updates the original version in the following aspects. Firstly, as target platform the Renesas M32C/85U embedded processor is used, and the time cost model as been replaced to reflect costs for HAM operations on this processor. Secondly, the space cost models reflect changes in the Hume compiler, in particular the costs as incurred by the new native code generator. The cost model in Version 1.1 has been validated against measured costs on the target platform in Deliverable D28.

2 Hume Abstract Machine Design and Reference Implementation

name	interpretation	name	interpretation
<i>S</i>	stack	<i>pc</i>	program counter
<i>H</i>	heap	<i>pcr</i>	restart program counter
<i>sp</i>	stack pointer	<i>blocked</i>	box blocked
<i>hp</i>	heap pointer	<i>blockedon</i>	output on which blocked
<i>fp</i>	frame pointer	<i>INITPC</i>	initial program counter
<i>slp</i>	function frame pointer	<i>ins</i>	input buffers
<i>mp</i>	match pointer	<i>outs</i>	output buffers
<i>inp</i>	input pointer	<i>nIns</i>	number of inputs
<i>rs</i>	current ruleset	<i>nOuts</i>	number of outputs
<i>base</i>	base ruleset		

Figure 1: Box-specific registers, constants and memory areas — the *box state record*

name	interpretation
<i>rules</i>	array of rule entry points
<i>nRules</i>	number of rules
<i>rp</i>	current rule pointer

Figure 2: Ruleset-specific registers and constants

The goal of the Hume Abstract Machine (HAM) design is to provide a credible basis for research into bounded time and space computation, allowing formal cost models to be verified against a realistic implementation. Absolute space- and time-performance (while an important long-term objective for Hume) is thus less important in this initial design than predictability, simplicity and ease of implementation.

The Hume Abstract Machine is loosely based on the design of the classical G-Machine [Aug87] or SECD-Machine [Lan64], with extensions to manage concurrency and asynchronicity. Each Hume box is implemented as a thread with its own dynamic stack (*S*) and heap (*H*) and associated stack and heap pointers (*sp* and *hp*). These and the other items that form part of the individual *state record* are shown in Figure 1. Each function and box has an associated *ruleset* (Figure 2). The ruleset is used for two purposes: it gives the address of the next rule to try if matching fails; and it is used to reorder rules if fair matching is specified. The box ruleset is specified as the *base* field of the state record. Function rulesets are set as part of a function call.

Separate stacks are needed to maintain independent state records. Separate heaps allow a simple model of garbage collection where the entire heap becomes garbage each time a box completes. Small

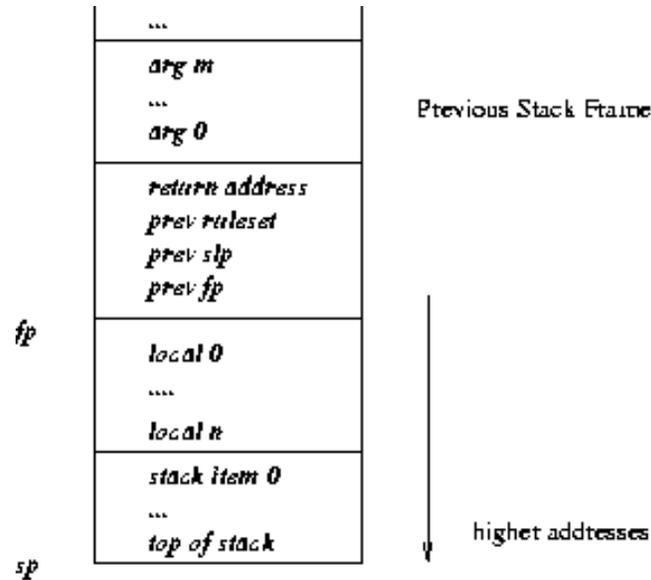


Figure 3: Stack frame layout in the Hume Abstract Machine

pointer ranges can be used in both cases (8-bit stack and heap pointers are possible for a number of applications). This helps conserve space. The corresponding disadvantage of this design is the need to communicate arguments and results between boxes rather than using a physically or virtually shared heap. This is achieved in the HAM reference implementation by copying such values between heaps at the beginning and end of the box execution. There is an analogy with the *working copies* of global variables that may be obtained by implementations of the JVM [LY99]. However, variable accesses in the JVM may occur at any point during thread execution, not only at the beginning/end as in the HAM. Moreover, unlike the HAM, which is stateless, the JVM maintains a virtually shared heap containing *master copies* of each variable. Our design is thus closer to that of Eden [BLOP96]: a reactive functional language based on Haskell.

The layout of a typical stack frame is shown in Figure 3. The abstract machine design uses a pure stack calling convention. Function arguments are followed by a four-item frame-header containing the return address, a pointer to the previous ruleset, the static link pointer and the previous frame pointer. In this description, the size of this subframe is given by the constant \mathcal{S}_{frame} . The local frame pointer *fp* points immediately after the frame-header, to the address of the first local variable. The function frame pointer *slp* points to the beginning of the nearest frame representing a function (note that let etc instructions in Hume also allocate frames). For consistency, the same layout is used at the outer box level. In this case, the box inputs are stored in the argument position, and the return address item is redundant. Note that all values on the stack other than the saved return address, ruleset and frame pointer are pointers to the local heap (i.e. they are *boxed* [PL92]): in the current design there is no separate *basic value stack* to handle scalar values as in some versions of the G-Machine [PL92], STG-Machine [Pey92] etc. nor are scalars and heap objects mixed on the stack as in the JVM [LY99] or recent versions of the STG-Machine.

2.1 Box Scheduling

The implementation maintains a vector of boxes, each with its own state record. Boxes are connected through *wires*, which are shared communication buffers each connecting an *in* of one box to an *out* of another (or the same) box. Each wire comprises a pair of a value (*value*) and a flag indicating that the value is valid (*available*), used to ensure correct locking.

```

for  $i = 1$  to  $nBoxes$  do
   $runnable := false$ ;
  for  $j = 1$  to  $box[i].base.nRules$  do
    if  $\neg runnable$  then
       $runnable := true$ ;
      for  $k = 1$  to  $box[i].nIns$  do
         $runnable \& = box[i].required[j,k] \Rightarrow box[i].ins[k].available$ 
      endfor
    endif
  endfor
  if  $runnable$  then  $schedule(box[i])$  endif
endfor

```

Figure 4: Scheduling Algorithm

Boxes are scheduled under the control of a built-in scheduler. The exact scheduling order is not fixed by the HAM semantics. Implementations are free to realise any order that satisfies the conditions stated in Section 4.4. A box is deemed to be *runnable* if all the required inputs are available for any of its rules to be executed (Figure 4). A compiler-specified matrix is used to determine whether an input is needed: for some box t , $box[t].required[r,i]$ is true if input i is required to run rule r of that box. A single execution cycle comprises the following (all realised via explicit HAM instructions):

1. initialise stack- and heap-pointers and the base ruleset;
2. check input availability against possible matches;
3. copy data from input wires into the local heap for matching;
4. match available inputs against rules;
5. consume those inputs that have been matched and which are not ignored in the selected rule;
6. create a stack frame to hold local variables;

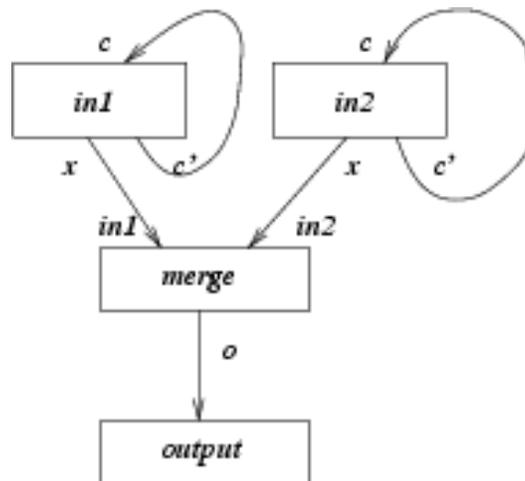


Figure 5: Hume example program as a network of boxes

constant	value (words)	constant	value (words)
$\mathcal{H}_{\text{Constr}}$	2	\mathcal{H}_{Fun}	4
$\mathcal{H}_{\text{Tuple}}$	2	\mathcal{H}_{Exn}	2
$\mathcal{H}_{\text{float}}$	2	\mathcal{H}_{Int}	2
\mathcal{H}_{int}	2	$\mathcal{H}_{\text{Float}}$	2
$\mathcal{H}_{\text{Bool}}$	2	$\mathcal{S}_{\text{saved}}$	4
$\mathcal{H}_{\text{Char}}$	2		
\mathcal{H}_{Str}	1		

Table 1: Sizes of headers for heap and stack objects in the prototype Hume Abstract Machine

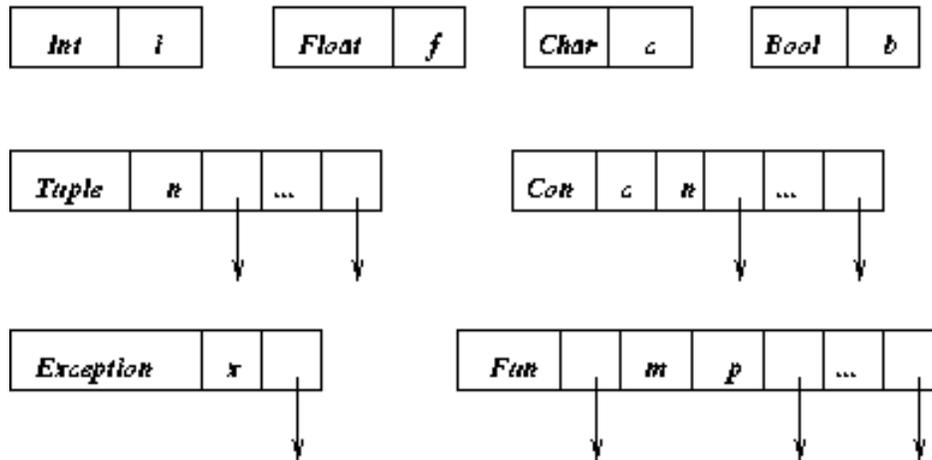


Figure 6: Heap representations in the Hume Abstract Machine

7. bind variables to input values;
8. evaluate the RHS of the selected rule;
9. check that outputs can be written to all output wires;
10. write non-ignored outputs to the corresponding wires;
11. reorder match rules according to the fairness criteria.

Note that all wires are single-buffered. A box will therefore block when writing to a wire which contains an output that has not yet been consumed. In order to ensure a consistent semantics, a single check is performed on all output wires just before any output is written. The check ignores * output positions. The box suspends if any of the needed output wires is occupied.

2.2 Heap Representations

In the prototype design, all heap cells are *boxed* with tags distinguishing different kinds of objects. Furthermore, tuple (**Tuple**), constructor (**Constr**) and exception structures (**Exn**) require *size* fields, and constructors also require a *constructor tag* field. Function closures (**Fun**) contain a pointer to an instruction sequence, the number of needed (*m*) and of already provided (*p*) arguments, and a list of pointers to those provided arguments. All data objects in a structure are referenced by pointer. There is one special representation: strings are represented as a tagged sequence of bytes. Clearly,

heap usage could be reduced using a more compact representation such as that used by modern high-performance implementations such as SML-NJ [Mac] or the STG-Machine [Pey92]. For now, we are, however, primarily concerned with bounding and predicting memory usage. Small changes to data representations can be easily incorporated into both models and implementations at a future date without affecting the fundamental results of cost modelling, except by reducing absolute costs in both cases.

2.2.1 Heap and Stack Sizes.

The sizes of the stack and heap spaces associated with each box are fixed at compile-time using a static analysis to compute the upper bound of stack and heap usage [HM02]. This analysis is already defined for flat Hume programs containing no recursion (FSM-Hume), and is currently being extended to more general cases. For the examples that have been tested to date, the analysis calculates space usage to within 15% of the actual dynamic usage. Since this will include situations where, for example, only one input of a 2-input box is ever available at run-time, or where a conditional branch is never chosen dynamically, this represents a good upper-bound estimate of the space usage. The simple analysis has been integrated into our prototype compiler (`phamc`), and results from the analysis are used in this paper where appropriate.

2.2.2 Exceptions.

Two forms of exceptions can occur during the execution of HAM code: synchronous and asynchronous exception. While synchronous exceptions are related to the execution of a particular piece of HAM code (e.g. division-by-zero), asynchronous exceptions can occur at any point and need to be checked by an external system.

The most important synchronous exceptions are stack- and heap-overflow. Before each (block of) instruction, that involves an increase of the stack- or heap-pointer, a stack- or heap-check is necessary, which makes sure that enough stack- or heap space is available, and if not raises an exception. For the pseudo-code in the reference implementation we assume that such stack- and heap-checks are added for any HAM instruction.

The most important asynchronous exception is a timeout. We assume that the HAM is embedded into a system that provides facilities for setting a timer and for checking when such a timer expires. In this case, the handler code for timeouts in the current box must be executed.

2.3 The Abstract Machine Instructions

The abstract machine instructions implement the abstract machine design described above. These instructions are shown in Figures 7–11. They are classified into stack, heap and control-flow operations, which are fairly standard, and matching and scheduling operations, which reflect the specific nature of the Hume design. The description of the instructions uses two auxiliary functions: *maxVars* calculates the maximum number of variables in a list of patterns; and *labels* generates new labels for a set of function/box rules. Where labels *lt*, *ln* etc. are used, these are assumed to be unique.

2.3.1 Heap Object Creation (Fig 7).

Tagged objects are created in the heap, and pointers to the new object stored on the top of the stack. For scalar values (booleans, characters, integers, floats and strings) the actual value is taken directly from the instruction stream (in the case of a string, this is a pointer into the global string table). The corresponding constructor instructions are `MkBool`, `MkChar`, `MkInt*`, `MkFloat*` and `MkString`. The instruction `MkNone` creates a special `None` value in the heap. This is used to prevent the writing of a particular output under dynamic programmer control.

MkBool b	$H[hp] = \text{Bool } b; S[sp] := hp; ++sp; hp := hp + \mathcal{H}_{\text{Bool}}; ++pc$
MkChar c	$H[hp] = \text{Char } c; S[sp] := hp; ++sp; hp := hp + \mathcal{H}_{\text{Char}}; ++pc$
MkInt i	$H[hp] = \text{Int } i; S[sp] := hp; ++sp; hp := hp + \mathcal{H}_{\text{Int}}; ++pc$
MkFloat f	$H[hp] = \text{Float } f; S[sp] := hp; ++sp; hp := hp + \mathcal{H}_{\text{Float}}; ++pc$
MkString s	$H[hp] = \text{Str } s; S[sp] := hp; ++sp; hp := hp + \mathcal{H}_{\text{Str}} + \text{ssize}(s); ++pc$
MkNone	$H[hp] = \text{None}; S[sp] := hp; ++sp; hp := hp + \mathcal{H}_{\text{None}}; ++pc$
MkCon $c \ n$	$H[hp] = \text{Constr } c \ n \ (S[sp-1]) \ \dots \ (S[sp-n-1]); sp := sp - n;$ $S[sp] := hp; ++sp; hp := hp + \mathcal{H}_{\text{Constr}} + n; ++pc$
MkTuple n	$H[hp] = \text{Tuple } n \ (S[sp-1]) \ \dots \ (S[sp-n-1]); sp := sp - n;$ $S[sp] := hp; ++sp; hp := hp + \mathcal{H}_{\text{Tuple}} + n; ++pc$
MkFun $l \ m \ p$	$H[hp] = \text{Fun } l \ m \ p \ (S[sp-1]) \ \dots \ (S[sp-p-1]);$ $sp := sp - p; S[sp] := hp; ++sp; hp := hp + \mathcal{H}_{\text{Fun}} + p; ++pc$

Figure 7: Abstract Machine Instructions: Heap operations

Push n	$sp := sp + n; ++pc$
Pop n	$sp := sp - n; ++pc$
Slide n	$\left. \begin{array}{l} \\ \\ \\ \end{array} \right\} \begin{cases} m := n; \\ m := H[S[fp+v]+1]; \\ fp' := fp; \text{while } d \neq 0 \text{ do } fp' := S[fp'-2]; --d \text{ endwhile}; m := H[S[fp'+v]+1] \\ S[sp-m-1] := S[sp-1]; sp := sp - m; ++pc \end{cases}$
SlideVar v	
SlideVarF $d \ v$	
Copy n	$S[sp] := S[sp-n-1]; ++sp; ++pc$
CopyArg n	$S[sp] := S[fp - \mathcal{S}_{\text{saved}} - n - 1]; ++sp; ++pc$
CreateFrame n	$S[sp] := fp; fp := sp + \mathcal{S}_{\text{saved}}; sp := sp + \mathcal{S}_{\text{saved}} + n; lp := sp; ++pc$
PushVar v	$S[sp] := S[fp+v]; ++sp; ++pc$
PushVarF $d \ v$	$fp' := fp; \text{while } d \neq 0 \text{ do } fp' := S[fp'-2]; --d \text{ endwhile};$ $S[sp] := S[fp'+v]; ++sp; ++pc$
MakeVar v	$S[fp+v] := S[sp-1]; --sp; ++pc$

Figure 8: Abstract Machine Instructions: Stack operations

Finally, two instructions build structured values. **MkCon** takes two arguments from the instruction stream, a constructor tag and a number of arguments, and builds the corresponding constructor in the heap using the relevant number of arguments from the stack. **MkTuple** is similar except that it has no constructor tag parameter, and builds a tuple rather than a constructor. **MkFun** takes the top p elements from the stack and builds a function closure in the heap. The construction of an exception closure is encoded in the **Raise** instruction, together with the transfer of control to the exception handler of the current box.

2.3.2 Stack Operations (Fig 8).

The abstract machine uses a number of simple stack manipulation operations. **Push** increments the stack pointer by a constant. This is used to create fixed space on the stack. **Pop** decrements the stack pointer. **Slide** pops the stack by a fixed amount, and ensures that the top of stack after the stack is the same as before the pop. **SlideVar** and **SlideVarF** perform the same operation but read the amount for the slide from a local or non-local variable, respectively. This is used, for example, to remove the arguments to a function when it returns. **Copy** duplicates the contents of a stack location relative to the current top of stack. Finally, **CopyArg** copies the specified box or function argument to the top of stack.

Three operations are provided on variables. **PushVar** copies the specified variable to the top of stack. **PushVarF** does the same, but from a non-local target stack frame (specified as the depth d in the frame list and the relative offset v , both given in the instruction). **MakeVar** assigns the value on the top of the stack to the corresponding local variable v .

2.3.3 Control-flow Operations (Fig 9).

The abstract machine control instructions are conventional. **Goto** sets the pc to the appropriate instruction. **If** does the same conditionally on the value on the top of the stack. **Call** calls the specified function, saving the current ruleset on the stack for future use. The new ruleset and program counter are derived from the label for the function that is called. **CallVar** and **CallVarF** call the function which is pointed to by the a local or non-local variable, respectively. This variable must point to a **Fun** closure in the heap, which contains a pointer to the code to be executed, as well as a list of already provided arguments. **Ap** calls the function pointed to by the top-of-stack element. **CallPrim1** and **CallPrim2** call primitive functions with one or two arguments, respectively.

2.3.4 Matching Operations (Fig 10).

Matching is initiated by the **StartMatches** instruction, which sets the program counter to the first rule in the base rule set.

The same matching operations are used both for box inputs and for function arguments. The operations are divided into three sets: the **MatchRule** operation, which initialises the matching for a rule, and matches with a closing **MatchedRule** indicating the end of a matching block; the **MatchAvailable** and **MatchNone** operations which check box input availability; and the value matching operations such as **MatchBool** etc.

Matching takes place against a special stack pointer, the match pointer mp , which records the current match position. This is initialised in **MatchRule** to be just above the first argument to the box or function. The input pointer inp is also initialised to support input availability checking. Finally, the program counter is initialised to the start of the next rule.

The availability checking operations are used only for box inputs. **MatchAvailable** checks whether the next box input is available. If not, then the entire rule match fails, and the next rule is tried.

Goto l	$pc := l$	
If l	if $S[sp - 1] := true$ then $pc := l$ else $++pc$ endif ; $--sp$	
Call f	$S[sp++] := pc + 1$; $S[sp++] := fn$; $S[sp++] := slp$; $slp := sp + 1$; $fn := f.ruleset$; $fn.rp := 0$; $pc := fn.rules[0]$	
TailCall $f n d sz$	while $d \neq 0$ do $fp := S[fp - 2]$; $--d$ endwhile ; for i in $1..n$ do $S[fp - \mathcal{S}_{saved} - i] := S[sp - i]$ endfor $sp := fp + sz$; $lp := sp$; $fn := f.ruleset$; $fn.rp := 0$; $pc := fn.rules[0]$	
CallVarF $d v n$	$\left. \begin{array}{l} \text{CallVarF } d v n \\ \text{CallVar } v n \\ \text{Ap } n \end{array} \right\} \left\{ \begin{array}{l} \text{while } d \neq 0 \text{ do } fp := S[fp - 2]; \text{ } --d \text{ endwhile; } c := S[fp + v] \\ c := S[fp + v] \\ c := S[sp] \end{array} \right. \begin{array}{l} \text{if CallVarF} \\ \text{if CallVar} \\ \text{if Ap} \end{array}$ $f := H[c + 1]$; $m := H[c + 2]$; $p := H[c + 3]$ if $(n + p \geq m)$ then for i in $p - 1..0$ do $S[sp++] := H[c + 4 + i]$ endfor ; $S[sp++] := pc$; $S[sp++] := fn$; $S[sp++] := slp$; $fn.rp := 0$; $pc := H[c + 1]$; else $H[hp] := \text{Fun } f m (p + n) H[c + 4] \dots H[c + 4 + (p - 1)] S[sp - 1] \dots S[sp - n - 1]$; $hp := hp + p + n + 4$; endif ; 	
Return		$fp := S[fp - 1]$; $slp := S[fp - 2]$; $fn := S[fp - 3]$; $pc := S[fp - 4]$; $sp' := fp - \mathcal{S}_{saved}$; $S[sp'] := S[sp - 1]$; $sp := sp' + 1$
CallPrim1 p		$S[sp - 1] := p (S[sp]) (S[sp - 1])$; $++pc$
CallPrim2 p	$S[sp - 2] := p (S[sp]) (S[sp - 1]) (S[sp - 2])$; $--sp$; $++pc$	
Raise x	$H[hp] := \text{Exn } x S[sp]$; $--sp$; $sp := fp := slp := 0$; $S[sp++] := hp$; $hp := hp + \mathcal{H}_{\text{Exn}} + 1$; $sp := sp + \mathcal{S}_{saved}$; $pc := fn.rules[fn.handler]$	

Figure 9: Abstract Machine Instructions: Control-flow

MatchRule	$mp := fp - \mathcal{S}_{saved} + 1; inp := 0; pc := fn.rules[fn.rp]; ++fn.rp; sp := lp$
MatchedRule	$sp := lp$
MatchNone	$-- mp; ++inp; ++pc$
MatchAvailable	if $ins[inp].available$ then $++pc$ else $pc := fn.rules[fn.rp]$ endif ; $inp := inp + 1$
MatchBool b	$-- mp; \text{if } H[S[mp]] \neq \text{Bool } b \text{ then } pc := fn.rules[fn.rp] \text{ endif}$
MatchChar c	$-- mp; \text{if } H[S[mp]] \neq \text{Char } c \text{ then } pc := fn.rules[fn.rp] \text{ endif}$
MatchString s	$-- mp; \text{if } H[S[mp]] \neq \text{Str } s \text{ then } pc := fn.rules[fn.rp] \text{ endif}$
MatchInt i	$-- mp; \text{if } H[S[mp]] \neq \text{Int } i \text{ then } pc := fn.rules[fn.rp] \text{ endif}$
MatchFloat f	$-- mp; \text{if } H[S[mp]] \neq \text{Float } f \text{ then } pc := fn.rules[fn.rp] \text{ endif}$
MatchCon $c n$	$-- mp; \text{if } H[S[mp]] \neq \text{Constr } c n \text{ then } pc := fn.rules[fn.rp] \text{ endif}$
MatchTuple n	$-- mp; \text{if } H[S[mp]] \neq \text{Tuple } n \text{ then } pc := fn.rules[fn.rp] \text{ endif}$
MatchExn x	$-- mp; \text{if } H[S[mp]] \neq \text{Exn } x \text{ then } pc := fn.rules[fn.rp] \text{ endif}$
Unpack	$-- sp; \text{if } H[S[sp]] = \text{Tuple } n \text{ then } offset := 2$ $\text{else if } H[S[sp]] = \text{Constr } c n \text{ then } offset := 3;$ $\text{else if } H[S[sp]] = \text{Fun } f m n \text{ then } offset := 4; \text{endif};$ for $i := 0$ to $n - 1$ do $S[sp] := H[S[sp] + offset + i]; ++sp;$ endfor ; $++pc$
StartMatches	$pc := base.rules[0]$
Reorder	$n := fn.nrules - 1; r := fn.rules[fn.rp];$ for $i := fn.rp$ to n do $fn.rules[i] := fn.rules[i + 1]$ endfor ; $fn.rules[n] := r; ++pc$

Figure 10: Abstract Machine Instructions: Rule matching

CopyInput n	$S[sp] := copy(ins[n]); ++sp; ++pc$
Consume n	$ins[n].available := false; ++pc$
MaybeConsume n	if $ins[n].available$ then $ins[n].available := false$ endif ; $++pc$
CheckOutputs	for $i := 0$ to $nouts$ do if $H[S[sp - i - 1]] \neq None$ and $\neg outs[i].available$ then $blocked := true; bon := i; pcr := pc; reschedule;$ endif ; endfor ; $++pc$
Write n	$--sp;$ if $H[S[sp]] \neq None$ then $outs[n] := copy(H[S[sp]]);$ endif ; $++pc$
Input	$H[hp] := Char\ getchar; S[sp] := hp; ++sp;$
Output	$hp := hp + \mathcal{H}_{char}; ++pc$ $putvalue(S[sp - 1]); --sp; ++pc$
Within $h\ t$	$fn.rules[fn.within_handler] := h; S[sp] := lo(t); S[sp + 1] := hi(t); sp := sp + 2$
WithinStackSpace $h\ p$	$fn.rules[fn.withinspace_handler] := h; S[sp + 1] := splim; splim := sp + p; sp := sp + 1$
WithinHeapSpace $h\ p$	$fn.rules[fn.withinspace_handler] := h; S[sp + 1] := hplim; hplim := hp + p; sp := sp + 1$
DoneWithin	$unsetTimer()$
RaiseWithin $h\ t$	$setTimer(S[sp], S[sp - 1])$
DoneWithinStackSpace h	$splim := S[sp - 1]; S[sp - 1] := S[sp]; sp := sp - 1$
DoneWithinHeapSpace h	$hplim := S[sp - 1]; S[sp - 1] := S[sp]; sp := sp - 1$
Schedule	$reschedule$

Figure 11: Abstract Machine Instructions: Scheduling and Wire I/O

Otherwise the input pointer inp is incremented. **MatchNone** simply increments inp without checking input availability. This is used to implement ***** and **_*** in patterns.

The match operations match the heap value pointed to by mp against the specified value. Failure means the next rule is tried. Otherwise the next match position is checked. Constructors and tuple matches are applied only to the outer level. Nested matching is achieved by unpacking the arguments to the constructor or tuple onto the stack using the **Unpack** instruction.

Finally, after successful matching, rules may be reordered by **Reorder** if fair matching is required. This is ensured by moving the successful rule to the end of the ruleset. As a consequence, a least recently used policy is implemented.

2.3.5 Scheduling, input and output operations (Fig 11).

Box input and output is handled by two sets of operations. The **CopyInput** copies the specified input from the input wire into the heap and places it on the top of the stack prior to matching. If matching is successful, input is *consumed* using the **Consume** operation, which resets the availability semaphore for the appropriate input wire, thereby permitting new write operations on that wire.

Output is handled by two similar operations. The **Write** operation writes the value on the top of the stack to the specified output wire. Before this can be done, the **CheckOutputs** operation is used to ensure that all required Write operations will succeed. This is achieved by checking that all output wire buffers are empty (as indicated by the wire's *availability* semaphore). If not, then the box blocks

Label l	l labels the next instruction
Function $f l_1 \dots l_n$	Function f has rules at labels $l_1 \dots l_n$
Box $b h s i o r$	Box b has heap h , stack s , i inputs, o outputs and r rules
Rule $b l_1 \dots l_n$	Box b has rules at labels $l_1 \dots l_n$
Require $b x_1 \dots x_n$	Box b requires inputs $x_1 \dots x_n$
Stream s In/Out $h s$	Stream h has heap h and stack s
Wire $wi i wo o h$	Wire input $wi.i$ to output $wo.o$ with heap h

Figure 12: Abstract Machine Pseudo-Instructions

<code>sp := 0</code>	stack pointer
<code>hp := 0</code>	heap pointer
<code>fp := 0</code>	frame pointer
<code>slp := 0</code>	function frame pointer
<code>inp := 0</code>	in pointer
<code>fn := base</code>	code base
<code>pc := INITPC</code>	program counter
<code>blocked := false</code>	blocker flag
<code>splim := SPLIM</code>	end of stack area
<code>hplim := HPLIM</code>	end of heap area

Figure 13: Box initialisation

until the value on the wire has been consumed, and the *availability* semaphore has been cleared. If the heap value is **None** (corresponding to ***** on the output), then the **Write** will not actually write anything to the output wire, and the *availability* semaphore is ignored by **CheckOutputs**.

Two operations are provided to manage stream input/output. A special box is attached to each stream input and output device. Executing the **Input** operation blocks the box if no input is available. Otherwise the input is read into the box's heap, and can then be written to its output wire using normal abstract machine operations. The **Output** operation simply writes the value on the top of stack to the appropriate device.

Control is returned to the scheduler either when a box blocks as a consequence of being unable to write some output during the **CheckOutputs** operation, or explicitly when a box terminates as a consequence of exiting the **Schedule** operation. In either case, the scheduler will select a new runnable box to execute. If there is no runnable box, then in the concurrent implementation the system will terminate. In a distributed system, it would be necessary to check for global termination, including outstanding communications that could awaken some box.

2.3.6 Box initialisation (Fig 13).

When a box is scheduled, its registers are initialised as shown in Figure 13. The initialisation during the wire input initialisation phase is similar except that different base and INITPC values are used, corresponding to the *_init* code for the box. Registers other than *blocked* and *pc* are not initialised if the box is restarted after it is blocked. In this case, *blocked* is set to *false* and *pc* is set to *pcr*. The registers *splim* and *hplim* hold the upper bounds for the stack and the heap areas, initialised with the constants SPLIM and HPLIM. These registers can be temporarily modified by the family of **Within** instructions.

3 Cost Modelling via Resource Algebras

As the basis for reasoning about resource consumption in Hume programs, we now present cost models of the HAM for heap space, stack space and time consumption. One of our main design goals is to separate the tracking of these resources as far as possible from the description of the functional behaviour of the abstract machine. We therefore use the concept of *resource algebras* [ABM06] as introduced in the MRG project.

A *resource algebra* \mathcal{R} is a partially ordered monoid $(R, 0, +, \leq)$, i.e. $(R, 0, +)$ is a monoid and (R, \leq) is a partially ordered set, where 0 is the minimum element, and $+$ is order preserving on both sides. Moreover, \mathcal{R} has constants in R for each expression former of the language. These constants denote the costs of the corresponding expression. In general, these constants are parameterised with the current state of the computation, in particular the current stack and heap. These costs are combined by the rules of the operational semantics in Section 4 with the $+$ operator.

3.1 A Resource Algebra for Heap Space

In instantiating the concept of a resource algebra with heap costs, we use the natural numbers \mathbb{N} as domain R , with addition as $+$, the natural number zero as 0 and the less-or-equal relation as \leq . Table 2 defines the changes in heap size for all instructions in the HAM language. The tables $\mathit{Prim1}_{s,\eta,f}$ and $\mathit{Prim2}_{s,\eta,f}$ define the heap costs for unary and binary primitive operations, with the current stack s , heap η , and the primitive function f as arguments.

The heap consumption of the constructors is defined by the constants in Table 1. The functions *ssize* returns the length of a string. For tuples, constructors, exceptions and function closures, the number of arguments has to be added to the header size. In the case of an exception this is always 1. The stack operations do not consume any heap space. Matching operations only compare values, and therefore consume neither heap nor stack space. An input operation allocates a new character in the heap and therefore consumes one character cell. Raising an exception involves first the construction of an exception closure. The 1 reserves space for the runtime value passed via the exception.

For primitive operations we use a table mapping the operation, and its runtime arguments, to the corresponding heap consumption: $\mathit{Prim1}_{s,\eta,f}, \mathit{Prim2}_{s,\eta,f}$. The notation for manipulating stacks is explained in Figure 14.

3.2 A Resource Algebra for Stack Space

For modelling stack space we again use $(\mathbb{N}, +, \leq)$ as resource algebra. In this case the constants are defined in Table 3.

For all heap operations the stack size increases by 1, since a pointer to the newly allocated heap object is left on the stack. Pushing adds and popping subtracts 1 from the stack size. The slide operations reduce the stack size. In the case of `Slide` the amount is part of the instruction, in the case of `SlideVar` and `SlideVarF` the amount is the value of a local or non-local variable, respectively. For a non-local variable this means that i stack frames are popped before a variable lookup is made to the j -th variable in the *locals* component of the frame. The notation $|s^{vals}|$ is used to denote the length of the *vals* component in the topmost stack frame. In `MatchRule` and `MatchedRule` this component can be discarded, hence the negative value.

Copy operations add a pointer onto the top of the stack. A `CreateFrame` instructions pushes as many (dummy) elements onto the stack as prescribed by its argument i . Matching operations only compare values, and therefore consume neither heap nor stack space. A `Write` operation writes the heap element, pointed to by the top-of-stack element, to an output stream and then pops the pointer from the stack.

$\mathcal{R}_{s,\eta}^{\text{MkBool } b} = \mathcal{H}_{\text{Bool}}$ $\mathcal{R}_{s,\eta}^{\text{MkChar } x} = \mathcal{H}_{\text{Char}}$ $\mathcal{R}_{s,\eta}^{\text{MkString } x} = \mathcal{H}_{\text{Str}} + s_{\text{size } s}$ $\mathcal{R}_{s,\eta}^{\text{MkInt } i} = \mathcal{H}_{\text{Int}}$ $\mathcal{R}_{s,\eta}^{\text{MkFloat } f} = \mathcal{H}_{\text{Float}}$ $\mathcal{R}_{s,\eta}^{\text{MkTuple } i} = \mathcal{H}_{\text{Tuple}} + i$ $\mathcal{R}_{s,\eta}^{\text{MkCon } i j} = \mathcal{H}_{\text{Constr}} + j$ $\mathcal{R}_{s,\eta}^{\text{MkFun } f m p} = \mathcal{H}_{\text{Fun}} + p$ $\mathcal{R}_{s,\eta}^{\text{MkNone}} = 0$ $\mathcal{R}_{s,\eta}^{\text{Push } i} = 0$ $\mathcal{R}_{s,\eta}^{\text{Pop } i} = 0$ $\mathcal{R}_{s,\eta}^{\text{Slide } i} = 0$ $\mathcal{R}_{s,\eta}^{\text{SlideVar } i} = 0$ $\mathcal{R}_{s,\eta}^{\text{SlideVarF } i} = 0$	$\mathcal{R}_{s,\eta}^{\text{Copy } i} = 0$ $\mathcal{R}_{s,\eta}^{\text{CopyArg } i} = 0$ $\mathcal{R}_{s,\eta}^{\text{CreateFrame } i} = 0$ $\mathcal{R}_{s,\eta}^{\text{PushVar } i} = 0$ $\mathcal{R}_{s,\eta}^{\text{PushVarF } d i} = 0$ $\mathcal{R}_{s,\eta}^{\text{MakeVar } i} = 0$ $\mathcal{R}_{s,\eta}^{\text{Goto } \text{lbl}} = 0$ $\mathcal{R}_{s,\eta}^{\text{If } \text{lbl}} = 0$ $\mathcal{R}_{s,\eta}^{\text{Call } f} = 0$ $\mathcal{R}_{s,\eta}^{\text{Return}} = 0$	$\mathcal{R}_{s,\eta}^{\text{StartMatches}} = 0$ $\mathcal{R}_{s,\eta}^{\text{MatchRule}} = 0$ $\mathcal{R}_{s,\eta}^{\text{MatchedRule}} = 0$ $\mathcal{R}_{s,\eta}^{\text{MatchNone}} = 0$ $\mathcal{R}_{s,\eta}^{\text{MatchAvailable}} = 0$ $\mathcal{R}_{s,\eta}^{\text{AVAILSET}} = 0$ $\mathcal{R}_{s,\eta}^{\text{MatchAny}} = 0$ $\mathcal{R}_{s,\eta}^{\text{MatchBool } b} = 0$ $\mathcal{R}_{s,\eta}^{\text{MatchChar } x} = 0$ $\mathcal{R}_{s,\eta}^{\text{MatchString } x} = 0$ $\mathcal{R}_{s,\eta}^{\text{MatchInt } i} = 0$ $\mathcal{R}_{s,\eta}^{\text{MatchFloat } i} = 0$ $\mathcal{R}_{s,\eta}^{\text{MatchTuple } i} = 0$ $\mathcal{R}_{s,\eta}^{\text{MatchCon } i j} = 0$ $\mathcal{R}_{s,\eta}^{\text{MatchExn } x} = 0$	$\mathcal{R}_{s,\eta}^{\text{Unpack}} = 0$ $\mathcal{R}_{s,\eta}^{\text{Reorder}} = 0$ $\mathcal{R}_{s,\eta}^{\text{CheckOutputs}} = 0$ $\mathcal{R}_{s,\eta}^{\text{CopyInput } i} = 0$ $\mathcal{R}_{s,\eta}^{\text{COPYALLINPUTS}} = 0$ $\mathcal{R}_{s,\eta}^{\text{Consume } i} = 0$ $\mathcal{R}_{s,\eta}^{\text{MaybeConsume } i} = 0$ $\mathcal{R}_{s,\eta}^{\text{CONSUMESET}} = 0$ $\mathcal{R}_{s,\eta}^{\text{Write } i} = 0$ $\mathcal{R}_{s,\eta}^{\text{Input}} = \mathcal{H}_{\text{Char}}$ $\mathcal{R}_{s,\eta}^{\text{Output}} = 0$ $\mathcal{R}_{s,\eta}^{\text{Schedule}} = 0$ $\mathcal{R}_{s,\eta}^{\text{Raise } i} = \mathcal{H}_{\text{Exn}} + 1$
$\mathcal{R}_{s,\eta}^{\text{TailCall } f n d sz} = 0$ $\mathcal{R}_{s,\eta}^{\text{CallVar } f n} = \begin{cases} 0 & \text{if } \text{arity } s_f^{\text{locals}} \leq n + \text{providedArgs } s_f^{\text{locals}} \\ n + 4 & \text{otherwise} \end{cases}$ $\mathcal{R}_{s,\eta}^{\text{CallVarF } d n} = \begin{cases} 0 & \text{if } \text{arity } (s - \square^i)_j^{\text{locals}} \leq n + \text{providedArgs } (s - \square^i)_j^{\text{locals}} \\ n + 4 & \text{otherwise} \end{cases}$ $\mathcal{R}_{s,\eta}^{\text{CallPrim1 } f} = \mathcal{P}_{\text{prim1 } s,\eta,f} s_0$ $\mathcal{R}_{s,\eta}^{\text{CallPrim2 } f} = \mathcal{P}_{\text{prim2 } s,\eta,f} s_0(s - \top_0)$			

Table 2: Heap consumption of HAM instructions

An `Unpack` operation simply consumes the top-of-stack pointer unless it is a tuple or constructor. In the latter cases, it thereafter adds all the elements of the tuple or constructor onto the stack. In a `Raise` instruction the entire stack is discarded, a new frame is built on the empty stack and a pointer to an exception closure is pushed onto the stack. Therefore, we have to subtract the entire stack size, computed as a sum over the sizes of all stack frames, and then add the frame header size plus 1. In a `TailCall` instruction the top n frames are discarded, a new frame header created, and o dummy values (for local variables) added to the stack. A `Return` operation removes the topmost frame from the stack. The size of the stack frame is the frame header size plus the number of local variables plus the size of the expression stack. The stack consumption of a unary primitive operation is 0, because it overwrites its argument on the stack. For a binary primitive operation the stack consumption is -1 because both arguments can be discarded and only the result value is pushed.

3.3 A Resource Algebra for Time

For modelling time consumption we use rational numbers as the base domain with the usual addition and less-or-equal relation $(\mathbb{R}, +, \leq)$ as resource algebra.

In Table 4 we summarise upper bounds for execution time on our target processor, a Renesas M32C. We have derived these bounds by using AbsInt's aiT tool [?], which performs machine-code-level analysis

$\mathcal{R}_{s,\eta}^{\text{MkBool } b} =$ $\mathcal{R}_{s,\eta}^{\text{MkChar } x} =$ $\mathcal{R}_{s,\eta}^{\text{MkString } x} =$ $\mathcal{R}_{s,\eta}^{\text{MkInt } i} =$ $\mathcal{R}_{s,\eta}^{\text{MkFloat } f} =$ $\mathcal{R}_{s,\eta}^{\text{GETCONST}} = 0$ $\mathcal{R}_{s,\eta}^{\text{MkTuple } i} = -n + 1$ $\mathcal{R}_{s,\eta}^{\text{MkCon } i j} = -n + 1$ $\mathcal{R}_{s,\eta}^{\text{MkVector } i j} = -n + 1$ $\mathcal{R}_{s,\eta}^{\text{MkFun } f m p} = -p + 1$ $\mathcal{R}_{s,\eta}^{\text{MkNone}} = 1$ $\mathcal{R}_{s,\eta}^{\text{Push } i} = i$ $\mathcal{R}_{s,\eta}^{\text{PushVar } i} = 1$ $\mathcal{R}_{s,\eta}^{\text{Pop } i} = -i$ $\mathcal{R}_{s,\eta}^{\text{Slide } i} = -i$ $\mathcal{R}_{s,\eta}^{\text{SlideVar } i} = -(\eta s_i^{\text{locals}})$ $\mathcal{R}_{s,\eta}^{\text{SlideVarF } i j} = -((s - \square^i)_j^{\text{locals}})$	$\mathcal{R}_{s,\eta}^{\text{Copy } i} = 1$ $\mathcal{R}_{s,\eta}^{\text{CopyArg } i} = 1$ $\mathcal{R}_{s,\eta}^{\text{CreateFrame } i} = i + 1$ $\mathcal{R}_{s,\eta}^{\text{PushVar } i} = 1$ $\mathcal{R}_{s,\eta}^{\text{PushVarF } d i} = 1$ $\mathcal{R}_{s,\eta}^{\text{MakeVar } i} = -1$ $\mathcal{R}_{s,\eta}^{\text{Goto } \text{lbl}} = 0$ $\mathcal{R}_{s,\eta}^{\text{If } \text{lbl}} = -1$ $\mathcal{R}_{s,\eta}^{\text{Call } f} = 3$	$\mathcal{R}_{s,\eta}^{\text{StartMatches}} = 0$ $\mathcal{R}_{s,\eta}^{\text{MatchRule}} = - s^{\text{vals}} $ $\mathcal{R}_{s,\eta}^{\text{MatchedRule}} = - s^{\text{vals}} $ $\mathcal{R}_{s,\eta}^{\text{MatchNone}} = 0$ $\mathcal{R}_{s,\eta}^{\text{MatchAvailable}} = 0$ $\mathcal{R}_{s,\eta}^{\text{AVAILSET}} = 0$ $\mathcal{R}_{s,\eta}^{\text{MatchNone } b} = 0$ $\mathcal{R}_{s,\eta}^{\text{MatchBool } b} = 0$ $\mathcal{R}_{s,\eta}^{\text{MatchChar } x} = 0$ $\mathcal{R}_{s,\eta}^{\text{MatchString } x} = 0$ $\mathcal{R}_{s,\eta}^{\text{MatchInt } i} = 0$ $\mathcal{R}_{s,\eta}^{\text{MatchFloat } i} = 0$ $\mathcal{R}_{s,\eta}^{\text{MatchTuple } i} = 0$ $\mathcal{R}_{s,\eta}^{\text{MatchCon } i j} = 0$ $\mathcal{R}_{s,\eta}^{\text{MatchVector } i} = 0$ $\mathcal{R}_{s,\eta}^{\text{MatchExn } x} = 0$	$\mathcal{R}_{s,\eta}^{\text{Reorder}} = 0$ $\mathcal{R}_{s,\eta}^{\text{CheckOutputs}} = 0$ $\mathcal{R}_{s,\eta}^{\text{CopyInput } i} = 0$ $\mathcal{R}_{s,\eta}^{\text{COPYALLINPUTS } i} = i$ $\mathcal{R}_{s,\eta}^{\text{Consume } i} = 0$ $\mathcal{R}_{s,\eta}^{\text{MaybeConsume } i} = 0$ $\mathcal{R}_{s,\eta}^{\text{CONSUMESET } i} = 0$ $\mathcal{R}_{s,\eta}^{\text{Write } i} = -1$ $\mathcal{R}_{s,\eta}^{\text{Input}} = 1$ $\mathcal{R}_{s,\eta}^{\text{Output}} = -1$ $\mathcal{R}_{s,\eta}^{\text{Schedule}} = 0$
$\mathcal{R}_{s,\eta}^{\text{TailCall } f n sz} = - s^{\text{locals}} $ $\mathcal{R}_{s,\eta}^{\text{CallVar } f n} = \begin{cases} \text{providedArgs } s_f^{\text{locals}} + 3 & \text{if } \text{arity } s_f^{\text{locals}} \leq n + \text{providedArgs } s_f^{\text{locals}} \\ -n + 1 & \text{otherwise} \end{cases}$ $\mathcal{R}_{s,\eta}^{\text{CallVarF } f d n} = \begin{cases} \text{providedArgs } s_f^{\text{locals}} + 3 & \text{if } \text{arity}(s - \square^i)_j^{\text{locals}} \leq n + \text{providedArgs } (s - \square^i)_j^{\text{locals}} \\ -n + 1 & \text{otherwise} \end{cases}$ $\mathcal{R}_{s,\eta}^{\text{Unpack}} = \begin{cases} n - 1 & \text{if } \eta s_0 = \text{Tuple } n \ xs \\ n - 1 & \text{if } \eta s_0 = \text{Constr } c \ n \ xs \\ n - 1 & \text{if } \eta s_0 = \text{Fun } f \ m \ n \ xs \\ -1 & \text{otherwise} \end{cases}$ $\mathcal{R}_{s,\eta}^{\text{Return}} = -(\mathcal{S}_{\text{saved}} + s^{\text{locals}} + s^{\text{vals}}) + 1$	$\mathcal{R}_{s,\eta}^{\text{CallPrim1 } f} = 0$ $\mathcal{R}_{s,\eta}^{\text{CallPrim2 } f} = -1$		
$\mathcal{R}_{s,\eta}^{\text{Raise } x} = \mathcal{S}_{\text{saved}} + 1 - \sum_{0 \leq i < \infty} (\mathcal{S}_{\text{saved}} + s - \square^i _{\text{locals}} + s - \square^i _{\text{vals}})$ $\mathcal{R}_{s,\eta}^{\text{TailCall } f n d sz} = \mathcal{S}_{\text{saved}} + sz - \sum_{0 \leq i \leq d} (\mathcal{S}_{\text{saved}} + s - \square^i _{\text{locals}} + s - \square^i _{\text{vals}})$			

Table 3: Stack consumption of HAM instructions

of worst case execution time. Details on these measurements are reported in Deliverable D14 [?]. This tool accounts for low-level architecture issues such as the states of the pipeline and, if present, of the cache. However, comparisons of the sum of HAM instruction costs with analysing entire basic blocks have shown that these low-level architecture issues have little influence on the overall costs on the M32C processor. We therefore don't model these low-level issues in this resource algebra for time and compute the costs of a basic block as the sum of the costs of its HAM instructions. In the future, we might consider a system that uses the aiT tool on basic blocks of the program under consideration.

The current version of the Hume native code-generator pre-allocates constants in the text area of the compiled code. Therefore, all `Mk...` instructions in the HAM code are replaced by the new `GETCONSTANT` HAM instruction, which only fetches a pointer to the pre-allocated constant. In the costs

for the `Return` instruction, f indicates that the current function, from which to return, is the f -th function in the program. This is due to the fact that the `Return` on the M32C uses a switch statement over all possible return points. The HAM instruction `AVAILSET n` is a new instruction that checks the availability of n input wires. The HAM instruction `CONSUMESET n` is a new instruction that consumes the values from n input wires after a successful match.

Table 5 summarises the bounds on execution time for primitive functions used in the Hume compiler. They have been derived by using the aiT tool.

$\mathcal{R}_{s,\eta}^{\text{MkBool } b} = 85$		$\mathcal{R}_{s,\eta}^{\text{StartMatches}} = 111$	
$\mathcal{R}_{s,\eta}^{\text{MkChar } x} = 84$		$\mathcal{R}_{s,\eta}^{\text{MatchRule}} = 20$	$\mathcal{R}_{s,\eta}^{\text{Reorder}} =$
$\mathcal{R}_{s,\eta}^{\text{MkInt } i} = 83$		$\mathcal{R}_{s,\eta}^{\text{MatchedRule}} = 10$	$\mathcal{R}_{s,\eta}^{\text{CheckOutputs}} = 602$
$\mathcal{R}_{s,\eta}^{\text{MkFloat } f} = 91$	$\mathcal{R}_{s,\eta}^{\text{Copy } i} = 31$	$\mathcal{R}_{s,\eta}^{\text{MatchNone}} = 11$	$\mathcal{R}_{s,\eta}^{\text{CopyInput } i} = 78$
$\mathcal{R}_{s,\eta}^{\text{GETCONST } i} = 35$	$\mathcal{R}_{s,\eta}^{\text{CopyArg } i} = 35$	$\mathcal{R}_{s,\eta}^{\text{MatchAny}} = 11$	$\mathcal{R}_{s,\eta}^{\text{COPYALLINPUTS } i} = 4 + 74i$
$\mathcal{R}_{s,\eta}^{\text{MkTuple } i} = 52i + 78$	$\mathcal{R}_{s,\eta}^{\text{CreateFrame } i} = 72$	$\mathcal{R}_{s,\eta}^{\text{MatchAvailable}} = 13$	$\mathcal{R}_{s,\eta}^{\text{Consume } i} = 31$
$\mathcal{R}_{s,\eta}^{\text{MkCon } i j} = 54i + 107$	$\mathcal{R}_{s,\eta}^{\text{PushVar } i} = 39$	$\mathcal{R}_{s,\eta}^{\text{AVAILSET } n} = 6 + 7n$	$\mathcal{R}_{s,\eta}^{\text{MaybeConsume } i} = 28$
$\mathcal{R}_{s,\eta}^{\text{MkFun } f m p} = 60p + 166$	$\mathcal{R}_{s,\eta}^{\text{PushVarF } d i} = 11d + 35$	$\mathcal{R}_{s,\eta}^{\text{MatchBool } b} = 35$	$\mathcal{R}_{s,\eta}^{\text{CONSUMESET } i} = 32i$
$\mathcal{R}_{s,\eta}^{\text{MkNone}} = 25$	$\mathcal{R}_{s,\eta}^{\text{MakeVar } i} = 35$	$\mathcal{R}_{s,\eta}^{\text{MatchChar } x} = 35$	$\mathcal{R}_{s,\eta}^{\text{Write } i} =$
$\mathcal{R}_{s,\eta}^{\text{Push } i} = 9$	$\mathcal{R}_{s,\eta}^{\text{Goto } \text{lbl}} = 3$	$\mathcal{R}_{s,\eta}^{\text{MatchString } x} = 35$	$\mathcal{R}_{s,\eta}^{\text{Input}} =$
$\mathcal{R}_{s,\eta}^{\text{Pop } i} = 9$	$\mathcal{R}_{s,\eta}^{\text{Call } f} = 70$	$\mathcal{R}_{s,\eta}^{\text{MatchInt } i} = 32$	$\mathcal{R}_{s,\eta}^{\text{Output}} =$
$\mathcal{R}_{s,\eta}^{\text{Slide } i} = 53$	$\mathcal{R}_{s,\eta}^{\text{Return}^1} = 51 + 15f$	$\mathcal{R}_{s,\eta}^{\text{MatchFloat } i} = 35$	$\mathcal{R}_{s,\eta}^{\text{Schedule}} = 602$
$\mathcal{R}_{s,\eta}^{\text{SlideVar } i} = 76$		$\mathcal{R}_{s,\eta}^{\text{MatchTuple } i} = 11$	$\mathcal{R}_{s,\eta}^{\text{Raise } i} = 377$
$\mathcal{R}_{s,\eta}^{\text{SlideVarF } i j} = 11j + 79$		$\mathcal{R}_{s,\eta}^{\text{MatchCon } i j} = 30$	
		$\mathcal{R}_{s,\eta}^{\text{MatchExn } x} = 30$	
		$\mathcal{R}_{s,\eta}^{\text{MatchVector } i} = 11$	
$\mathcal{R}_{s,\eta}^{\text{TailCall } f n d sz} = 22 + 36d + 27n$			
$\mathcal{R}_{s,\eta}^{\text{CallVar } n} = 13$			
$\mathcal{R}_{s,\eta}^{\text{CallVarF } d n} = 13d$			
$\mathcal{R}_{s,\eta}^{\text{MkString } x} = 13 x + 140$			
$\mathcal{R}_{s,\eta}^{\text{Unpack } x} = 44 x + 51$		$\mathcal{R}_{s,\eta}^{\text{If } \text{lbl}} = 30$	
$\mathcal{R}_{s,\eta}^{\text{CopyInput } i} = \begin{cases} 78 \text{ if Int} \\ 74 s + 4 \text{ if Str} \end{cases}$		$\mathcal{R}_{s,\eta}^{\text{CallPrim1 } f} = \text{Prim1}_{s,\eta,f} s_0$	
		$\mathcal{R}_{s,\eta}^{\text{CallPrim2 } f} = \text{Prim2}_{s,\eta,f} s_0(s - \top_0)$	

Table 4: WCET of HAM instructions on a M32C (in cycles)

4 Formalisation of the HAM

The goal of this section is to give a formal specification of the HAM semantics, corresponding to the behaviour of the references implementation in the previous section. This formalisation is the basis of an HAM operational semantics that is encoded in Isabelle. This encoding is developed alongside this document and will be the basis for automated certification of Hume code as needed in WP4.

==.c = 125							
==.n = 125	>.c = 124		/.n = 206			@ = 89	
==.w = 125	>.n = 124		/.w = 206			vecdef =	
==.i = 125	>.w = 124		/.i = 206			vecmap =	
==.f = 125	>.i = 124	+.n = 116	/.f = 959	toInt.i = 2973		update =	
==.s = 125	>.f = 209	+.w = 116	/ = 959	toInt.n = 2973		vecmake =	
==.e = 125	>.s = 209	+.i = 116	mod = 1294	toInt.w = 2973		sqrt =	
== = 125	> = 209	+.f = 1106	mod.n = 1294	toInt.f = 2973		ln =	
!=.c = 243	<=.c = 122	+ = 1106	mod.w = 1294	toInt = 2973		log10 =	
!=.n = 243	<=.n = 122	-.n = 116	mod.i = 1294	toFloat.n = 699		sin =	
!=.w = 243	<=.w = 122	-.w = 116	\% = 1294	toFloat.w = 699		cos =	
!=.i = 243	<=.i = 122	-.i = 116	div.n = 206	toFloat.i = 699		tan =	
!=.f = 243	<=.f = 211	-.f = 1112	div.w = 206	toFloat.f = 699		asin =	
!=.s = 243	<=.s = 211	- = 1112	div.i = 206	toFloat = 699		acos =	
!=.e = 243	<= = 211	*.n = 124	div = 206	toChar =		atan =	
!= = 243	<.c = 124	*.w = 124	&& = 153	show =		sinh =	
>=.c = 122	<.n = 124	*.i = 124	= 156	length =		cosh =	
>=.n = 122	<.w = 124	*.f = 356	^& = 151	++ =		tanh =	
>=.w = 122	<.i = 124	* = 356	^ = 151	++.s =		atan2 = 14305	
>=.i = 122	<.f = 209		^ = 151			** =	
>=.f = 211	<.s = 209		~ = 151			exp =	
>=.s = 211	< = 209		not = 118				
>= = 211							

Table 5: WCET of primitive operations on a M32C (in cycles)

4.1 Basic Definitions

As basic types we use *Locn* as locations into the heap, *Ref* as references (either a proper location or a null-pointer *Nullref*). Heap values (*HVal*) are a tagged union of basic types, compound types of tuples or constructors, exceptions or function closures. We use disjoint name spaces for *Label*, *Function* and box names (*BName*).

<i>Locn</i>	\equiv	\mathbb{N}															
<i>Ref</i>	\equiv	<i>Nullref</i> <i>Ref Locn</i>															
<i>HVal</i>	\equiv	<table border="0"> <tbody> <tr> <td>Int \mathbb{I}</td> <td> </td> <td>Tuple \mathbb{N} (<i>Ref list</i>)</td> </tr> <tr> <td>Bool \mathbb{B}</td> <td> </td> <td>Constr \mathbb{N} \mathbb{N} (<i>Ref list</i>)</td> </tr> <tr> <td>Char \mathbb{R}</td> <td> </td> <td>Fun <i>Function</i> \mathbb{N} \mathbb{N} (<i>Ref list</i>)</td> </tr> <tr> <td>Str \mathbb{R}</td> <td> </td> <td>Exn \mathbb{N} <i>Ref</i></td> </tr> <tr> <td><i>R Ref</i></td> <td> </td> <td>None</td> </tr> </tbody> </table>	Int \mathbb{I}		Tuple \mathbb{N} (<i>Ref list</i>)	Bool \mathbb{B}		Constr \mathbb{N} \mathbb{N} (<i>Ref list</i>)	Char \mathbb{R}		Fun <i>Function</i> \mathbb{N} \mathbb{N} (<i>Ref list</i>)	Str \mathbb{R}		Exn \mathbb{N} <i>Ref</i>	<i>R Ref</i>		None
Int \mathbb{I}		Tuple \mathbb{N} (<i>Ref list</i>)															
Bool \mathbb{B}		Constr \mathbb{N} \mathbb{N} (<i>Ref list</i>)															
Char \mathbb{R}		Fun <i>Function</i> \mathbb{N} \mathbb{N} (<i>Ref list</i>)															
Str \mathbb{R}		Exn \mathbb{N} <i>Ref</i>															
<i>R Ref</i>		None															

Stack values are references into the heap. A stack is a list of frame records. A frame record (*Frame*) contains the return address (*ret*), the arguments (*args*), the local variables (*locals*) and the expression stack (*vals*). By default, stack operations such as push and pop work on the value stack. See Figure 14 for the notation used to access and modify stack entries.

The heap is a finite map of locations to heap values, i.e. a function from locations to either the value *None* or *Some x*, where *x* is an *HVal*. An IO record collects all state info, not recorded in stack or heap. These are mainly related to wire input/output and rule matching. In particular, name of the box (*b*), a flag whether the box is blocked (*blocked*), a pointer to the next rule in the rule set of the box

(*rp*), and pointers to the next wire (*inp*) and the next stack value (*mp*) to be checked in a rule match.

<i>SVal</i>	\equiv	<i>Ref</i>
<i>Frame</i>	\equiv	(<i>ret</i> :: <i>Label</i> , <i>args</i> :: <i>SVal list</i> , <i>locals</i> :: <i>SVal list</i> , <i>vals</i> :: <i>SVal list</i>)
<i>Stack</i>	\equiv	<i>Frame list</i>
<i>Heap</i>	\equiv	<i>Locn</i> \rightsquigarrow_f <i>HVal</i>
<i>Wire</i>	\equiv	\mathbb{N}
<i>IO</i>	\equiv	(<i>b</i> :: <i>BName</i> , <i>blocked</i> :: <i>bool</i> , <i>crp</i> :: \mathbb{N} , <i>rp</i> :: \mathbb{N} , <i>inp</i> :: <i>Wire</i> , <i>mp</i> :: \mathbb{N})

4.2 A Roadmap through the Operational Semantics

The initial set of rules deals with the construction of heap values. The next set encodes basic operations on the stack components. Note that we abstract over several pointers present in the reference implementation, by using a structured stack representation as a list of frames.

The main control-flow operations are (conditional) jumps and calls covered in the rules STEP-GOTO, STEPIFTRUE, STEPIFFALSE, STEPTAILCALL, STEPCALL, STEPRETURN, STEPCALLPRIM1, STEPCALLPRIM2 at the end of the list of rules. To determine the target of a (conditional) jump the table Σ^{Lab} , mapping labels to instructions sequences, is used. For function calls a similar table Σ^{Fun} and for exceptions Σ^{Exn} is used.

The rules relevant for higher-order functions are STEPCALLVARSATURATED, STEPCALLVARUNSATURATED, STEPCALLVARFSATURATED, STEPCALLVARFUNSATURATED, STEPAPSATURATED, STEPAPUNSATURATED, STEPSLIDEVAR, STEPSLIDEVARF. A higher-order function can be called from a local variable **CallVar** a non-local variable **CallVarF** or from the top of the stack **Ap**. In each case we distinguish between the case where enough arguments are provided to perform the call, i.e. the call is saturated, or whether another function closure needs to be allocated. The only rule needed for exceptions is STEPRaise.

The interface between expression level and system level is defined by the rule STEPCHECKOUTPUTSFALSE. Note that no rules for a **Schedule** as first instruction or an empty instruction list exists. Thus, in this case, no reduction can be done on expression level, which operationally amounts to yielding control to the scheduler and performing a rule on the system level.

4.3 Rules of the Operational Semantics

The following judgement of the small-step semantics on Hume expression level

$$s, \eta, io, \theta \vdash cs \Downarrow_n^m (s', \eta', io', cs', p) \theta'$$

is read as: with an initial stack *s*, initial heap η , IO record *io* and with the wire environment θ , the HAM instruction sequence (a list) *cs* evaluates in $n - m$ steps to a final stack *s'*, final heap η' , a final IO record *io'*, using *p* resources and leaving the code sequence *cs'* as continuation. The new wire environment is θ' . The semantics of resources is intentionally left open in this semantics. It can be instantiated to every structure corresponding to a resource algebra as discussed in Section 3. Being a small-step semantics, the result of one step will be a state with a continuation code sequence.

While the usage of a stack and heap are standard, the IO record collects additional information on the state of a box execution, that is primarily needed during box input/output and matching operations. The pointers *mp* and *inp* are indices into the arguments component of the top stack frame (*s^{args}*) and into the set of in-wires for the current box. The **MatchBool** etc. instructions check whether the current argument is a boolean value, and if not continue with the next rule. The **CopyInput** instruction is the only one that uses the *inp* pointer in order to copy the value from the current wire into the argument portion of the stack frame.

l	a reference to location l in the heap
$fresh\ S$	returns an element x , s.t. $x \notin S$
$dom\ f$	the domain of a partial function f , i.e. the values x for which $f\ x$ is defined
s_0	get the top-of-stack element from the stack s
$s + v$	push value v on top of the stack s
$s + xs$	push all elements in the list xs on top of the stack s
$s - \top$	pop the top element from the stack s
$s - \top^n$	pop the n topmost elements from the stack s
$s - \square^i$	pop the top i frames from the stack s
$s + \perp^n$	push n dummy values onto the stack s
$s - \top^n + s_0$	pop elements 1 to n from the stack but keep the top-of-stack
$s_{0..n-1}$	get the elements 0 to $n-1$ from the stack s
$s_n^{locals} := s_0$	assign top-of-stack to n -th local variable
$s + \boxed{c, f}$	allocate a new frame (for function f) on the stack, with return address c
$s - \square$	remove the topmost frame from the stack s
$\boxed{f=l} = s \boxed{i}$	binds l to the value of field f in the i -th frame on stack s
io_f	access the value of field f in record io
$io(f := x)$	set the value of field f in record io to x
$f(x \mapsto y)$	update the function f to map x to y
$wire^{b,m}$	the name of the m -th in-wire of box b
$wire_{b,m}$	the name of the m -th out-wire of box b

Figure 14: Notation used in the operational semantics

The notation used in the rules of the operational semantics are summarised in Figure 14. We use $\#$ for list cons, $@$ for list append, $_++$ for incrementing an integer variable. The construct (\cdot) constructs a record, $r(x := \cdot)$ modifies the record r by updating the field x . We use $f\ x$ for applying a function f to argument x , and $f(x \mapsto \cdot)$ for updating the function f with a mapping of x to \cdot .

4.3.1 Heap operations

The heap operations are fairly standard. They allocate a fresh location, by calling $fresh\ (dom\ \eta)$. It is guaranteed by definition that $fresh\ S \notin S$. A pointer to the newly allocated heap object is left on the stack. In the case of tuples and constructors, the top n elements are taken from the stack and put into the heap closure.

$$\frac{l = fresh\ (dom\ \eta)}{s, \eta, io, \theta \vdash (\text{MkBool } x) \#cs \Downarrow_{n+1}^n (s + l, \eta(l \mapsto (\text{Bool } x)), io, cs, \mathcal{R}_{s,\eta}^{\text{MkBool } x}) \theta} \quad (\text{STEPMkBool})$$

$$\frac{l = fresh\ (dom\ \eta)}{s, \eta, io, \theta \vdash (\text{MkChar } x) \#cs \Downarrow_{n+1}^n (s + l, \eta(l \mapsto (\text{Char } x)), io, cs, \mathcal{R}_{s,\eta}^{\text{MkChar } x}) \theta} \quad (\text{STEPMkChar})$$

$$\frac{l = fresh\ (dom\ \eta)}{s, \eta, io, \theta \vdash (\text{MkString } x) \#cs \Downarrow_{n+1}^n (s + l, \eta(l \mapsto (\text{Str } x)), io, cs, \mathcal{R}_{s,\eta}^{\text{MkString } x}) \theta} \quad (\text{STEPMkString})$$

$$\frac{l = \text{fresh } (\text{dom } \eta)}{s, \eta, io, \theta \vdash (\text{MkInt } x) \# \text{cs } \Downarrow_{n+1}^n ((s + l, \eta(l \mapsto (\text{Int } x))), io, \text{cs}, \mathcal{R}_{s,\eta}^{\text{MkInt } x}) \theta} \quad (\text{STEPMKINT})$$

$$\frac{l = \text{fresh } (\text{dom } \eta)}{s, \eta, io, \theta \vdash (\text{MkFloat } x) \# \text{cs } \Downarrow_{n+1}^n ((s + l, \eta(l \mapsto (\text{Float } x))), io, \text{cs}, \mathcal{R}_{s,\eta}^{\text{MkFloat } x}) \theta} \quad (\text{STEPMKFLOAT})$$

$$\frac{l = \text{fresh } (\text{dom } \eta)}{s, \eta, io, \theta \vdash (\text{MkTuple } j) \# \text{cs } \Downarrow_{n+1}^n (s - \top^j + l, \eta(l \mapsto (\text{Tuple } x \ s_{0..j-1}))), io, \text{cs}, \mathcal{R}_{s,\eta}^{\text{MkTuple } j}) \theta} \quad (\text{STEPMKTUPLE})$$

$$\frac{l = \text{fresh } (\text{dom } \eta)}{s, \eta, io, \theta \vdash (\text{MkCon } i \ j) \# \text{cs } \Downarrow_{n+1}^n (s - \top^j + l, \eta(l \mapsto (\text{Constr}_i \ j \ s_{0..j-1}))), io, \text{cs}, \mathcal{R}_{s,\eta}^{\text{MkCon } i \ j}) \theta} \quad (\text{STEPMKCON})$$

$$\frac{l = \text{fresh } (\text{dom } \eta)}{s, \eta, io, \theta \vdash (\text{MkFun } f \ m \ p) \# \text{cs } \Downarrow_{n+1}^n (s - \top^p + l, \eta(l \mapsto (\text{Fun } f \ m \ p \ (s_{0..p-1}))), io, \text{cs}, \mathcal{R}_{s,\eta}^{\text{MkFun } f \ m \ p}) \theta} \quad (\text{STEPMKFUN})$$

$$\frac{l = \text{fresh } (\text{dom } \eta)}{s, \eta, io, \theta \vdash (\text{MkNone}) \# \text{cs } \Downarrow_{n+1}^n (s + l, \eta(l \mapsto \text{None})), io, \text{cs}, \mathcal{R}_{s,\eta}^{\text{MkNone}}) \theta} \quad (\text{STEPMKNONE})$$

4.3.2 Stack operations

The stack operations usually modify the expression stack, i.e. the *vals* component of the top stack frame. Superscripting accesses a particular component of the top stack frame; subscripting such a component selects an element; e.g. s_x^{args} accesses the argument x of the top stack frame. The notation $s - \square^i$ is used for popping the top i frames from stack s . For more details on notation see Figure 14.

$$\frac{}{s, \eta, io, \theta \vdash (\text{Push } i) \# \text{cs } \Downarrow_{n+1}^n (s + \perp^i, \eta, io, \text{cs}, \mathcal{R}_{s,\eta}^{\text{Push } i}) \theta} \quad (\text{STEPUSH})$$

$$\frac{}{s, \eta, io, \theta \vdash (\text{Pop } i) \# \text{cs } \Downarrow_{n+1}^n (s - \top^i, \eta, io, \text{cs}, \mathcal{R}_{s,\eta}^{\text{Pop } i}) \theta} \quad (\text{STEPPOP})$$

$$\frac{}{s, \eta, io, \theta \vdash (\text{Slide } i) \# \text{cs } \Downarrow_{n+1}^n ((s - \top^{i+1} + s_0, \eta, io, \text{cs}, \mathcal{R}_{s,\eta}^{\text{Slide } i}) \theta} \quad (\text{STEPSLIDE})$$

$$\frac{\text{Int } i = \eta \ s_x^{\text{locals}}}{s, \eta, io, \theta \vdash (\text{SlideVar } x) \# \text{cs } \Downarrow_{n+1}^n ((s - \top^{i+1} + s_0, \eta, io, \text{cs}, \mathcal{R}_{s,\eta}^{\text{SlideVar } x}) \theta} \quad (\text{STEPSLIDEVAR})$$

$$\frac{\text{Int } i = \eta (s - \square_x^y)^{locals}}{s, \eta, io, \theta \vdash (\text{SlideVarF } x \ y) \# \text{cs} \Downarrow_{n+1}^n ((s - \top^{i+1} + s_0, \eta, io, \text{cs}, \mathcal{R}_{s,\eta}^{\text{SlideVarF } x \ y}) \theta)} \quad (\text{STEPSLIDEVARF})$$

$$\frac{}{s, \eta, io, \theta \vdash (\text{Copy } i) \# \text{cs} \Downarrow_{n+1}^n (s + s_i, \eta, io, \text{cs}, \mathcal{R}_{s,\eta}^{\text{Copy } i}) \theta} \quad (\text{STEPCOPY})$$

$$\frac{}{s, \eta, io, \theta \vdash (\text{CopyArg } i) \# \text{cs} \Downarrow_{n+1}^n ((s + s_i^{args}, \eta, io, \text{cs}, \mathcal{R}_{s,\eta}^{\text{CopyArg } i}) \theta)} \quad (\text{STEPCOPYARG})$$

$$\frac{}{s, \eta, io, \theta \vdash (\text{CreateFrame } i) \# \text{cs} \Downarrow_{n+1}^n (s + \boxed{\text{c,f}} + \perp^i, \eta, io, \text{cs}, \mathcal{R}_{s,\eta}^{\text{CreateFrame } i}) \theta} \quad (\text{STEPCREATEFRAME})$$

$$\frac{}{s, \eta, io, \theta \vdash (\text{PushVar } i) \# \text{cs} \Downarrow_{n+1}^n (s + s_i^{locals}, \eta, io, \text{cs}, \mathcal{R}_{s,\eta}^{\text{PushVar } i}) \theta} \quad (\text{STEP PUSHVAR})$$

$$\frac{}{s, \eta, io, \theta \vdash (\text{PushVarF } i \ j) \# \text{cs} \Downarrow_{n+1}^n (s + (s - \square^i)^j^{locals}, \eta, io, \text{cs}, \mathcal{R}_{s,\eta}^{\text{PushVarF } i \ j}) \theta} \quad (\text{STEP PUSHVARF})$$

$$\frac{}{s, \eta, io, \theta \vdash (\text{MakeVar } i) \# \text{cs} \Downarrow_{n+1}^n ((s_i^{locals} := s_0) - \top, \eta, io, \text{cs}, \mathcal{R}_{s,\eta}^{\text{MakeVar } i}) \theta} \quad (\text{STEPMAKEVAR})$$

4.3.3 Control-flow operations

The following set of rules describes the control flow on the expression level of Hume. Here we assume that the compiler has generated for each box several tables, mapping labels, including function names, and exceptions to their corresponding instruction sequences: $\Sigma^{Lab}, \Sigma^{Exn}$. For convenience we drop the box name from the table lookup. Note, that the box name is always available as io_b .

$$\frac{}{s, \eta, io, \theta \vdash (\text{Goto } l) \# \text{cs} \Downarrow_{n+1}^n (s, \eta, io, \Sigma_1^{Lab}, \mathcal{R}_{s,\eta}^{\text{Goto } l}) \theta} \quad (\text{STEPGOTO})$$

$$\frac{l = s_0 \quad \eta \ l = \text{Bool } true}{s, \eta, io, \theta \vdash (\text{If } l \ b1) \# \text{cs} \Downarrow_{n+1}^n (s - \top, \eta, io, \Sigma_{l \ b1}^{Lab}, \mathcal{R}_{s,\eta}^{\text{If } l \ b1}) \theta} \quad (\text{STEPIFTRUE})$$

$$\frac{l = s_0 \quad \eta \ l = \text{Bool } false}{s, \eta, io, \theta \vdash (\text{If } l \ b1) \# \text{cs} \Downarrow_{n+1}^n (s - \top, \eta, io, \text{cs}, \mathcal{R}_{s,\eta}^{\text{If } l \ b1}) \theta} \quad (\text{STEPIFFALSE})$$

$$\frac{}{s, \eta, io, \theta \vdash (\text{Call } f) \# (\text{Label } c) \# \text{cs} \Downarrow_{n+1}^n (s + \boxed{\text{c,f}}, \eta, io, \Sigma_f^{Fun}, \mathcal{R}_{s,\eta}^{\text{Call } f}) \theta} \quad (\text{STEPCALL})$$

$$\frac{s' = ((s - \square^d) + \boxed{\text{c,f}})^{\text{args}} := s_{0..j-1} + \top^z}{s, \eta, io, \theta \vdash (\text{TailCall } f \ j \ d \ z) \# (\text{Label } c) \# \text{cs} \Downarrow_{n+1}^n (s', \eta, io, \Sigma_f^{\text{Fun}}, \mathcal{R}_{s,\eta}^{\text{TailCall } f \ j \ d \ z}) \theta} \quad (\text{STEPTAILCALL})$$

$$\frac{\boxed{\text{ret=1}} = s \ \boxed{0}}{s, \eta, io, \theta \vdash (\text{Return}) \# \text{cs} \Downarrow_{n+1}^n ((s - \square + s_0, \eta, io, \Sigma_l^{\text{Lab}}, \mathcal{R}_{s,\eta}^{\text{Return}}) \theta)} \quad (\text{STEPRETURN})$$

$$\frac{s, \eta, io, \theta \vdash (\text{CallPrim1 } f) \# \text{cs} \Downarrow_{n+1}^n (s - \top^1 + (f \ s_0), \eta, io, \text{cs}, \mathcal{R}_{s,\eta}^{\text{CallPrim1 } f}) \theta}{s, \eta, io, \theta \vdash (\text{CallPrim1 } f) \# \text{cs} \Downarrow_{n+1}^n (s - \top^1 + (f \ s_0), \eta, io, \text{cs}, \mathcal{R}_{s,\eta}^{\text{CallPrim1 } f}) \theta} \quad (\text{STEPCALLPRIM1})$$

$$\frac{s, \eta, io, \theta \vdash (\text{CallPrim2 } f) \# \text{cs} \Downarrow_{n+1}^n (s - \top^2 + (f \ s_0 \ s_1), \eta, io, \text{cs}, \mathcal{R}_{s,\eta}^{\text{CallPrim2 } f}) \theta}{s, \eta, io, \theta \vdash (\text{CallPrim2 } f) \# \text{cs} \Downarrow_{n+1}^n (s - \top^2 + (f \ s_0 \ s_1), \eta, io, \text{cs}, \mathcal{R}_{s,\eta}^{\text{CallPrim2 } f}) \theta} \quad (\text{STEPCALLPRIM2})$$

$$\frac{l = \text{fresh } (\text{dom } \eta)}{s, \eta, io, \theta \vdash (\text{Raise } x) \# \text{cs} \Downarrow_{n+1}^n (\emptyset + l, \eta(l \mapsto (\text{Exn } x \ s_0)), io, \Sigma_x^{\text{Exn}}, \mathcal{R}_{s,\eta}^{\text{Raise } x}) \theta} \quad (\text{STEPRAISE})$$

Three HAM instructions can be used for applying a higher-order function to some arguments: **CallVar** for a local variable, **CallVarF** for a non-local variable, and **Ap** for a function closure pointed to by the top-of-stack element. Each of these 3 come in a version for unsaturated application, i.e. the number of arguments provided is smaller than the arity of the function, and in a saturated version. In the former case a new function closure is built in the heap. In the latter case the function is applied to all its arguments.

$$\frac{\text{Fun } f \ m \ p \ xs = \eta \ s_i^{\text{locals}} \quad p + j \geq m}{s, \eta, io, \theta \vdash (\text{CallVar } i \ j) \# (\text{Label } c) \# \text{cs} \Downarrow_{n+1}^n ((s - s_{0..j-1} + \boxed{\text{c,f}})^{\text{args}} := xs@s_{0..j-1}, \eta, io, \Sigma_f^{\text{Fun}}, \mathcal{R}_{s,\eta}^{\text{CallVar } i \ j}) \theta} \quad (\text{STEPCALLVARSATURATED})$$

$$\frac{\text{Fun } f \ m \ p \ xs = \eta \ s_i^{\text{locals}} \quad p + j < m \quad l = \text{fresh } \eta}{s, \eta, io, \theta \vdash (\text{CallVar } i \ j) \# (\text{Label } c) \# \text{cs} \Downarrow_{n+1}^n (s - s_{0..j-1}, \eta(l \mapsto \text{Fun } f \ m \ (p + j) \ (xs@s_{0..j-1})), io, \text{cs}, \mathcal{R}_{s,\eta}^{\text{CallVar } i \ j}) \theta} \quad (\text{STEPCALLVARUNSATURATED})$$

$$\frac{\text{Fun } f \ m \ p \ xs = \eta \ (s - \square^d)_i^{\text{locals}} \quad p + j \geq m}{s, \eta, io, \theta \vdash (\text{CallVarF } d \ i \ j) \# (\text{Label } c) \# \text{cs} \Downarrow_{n+1}^n ((s - s_{0..j-1} + \boxed{\text{c,f}})^{\text{args}} := xs@s_{0..j-1}, \eta, io, \Sigma_f^{\text{Fun}}, \mathcal{R}_{s,\eta}^{\text{CallVarF } d \ i \ j}) \theta} \quad (\text{STEPCALLVARFSATURATED})$$

$$\frac{\text{Fun } f \ m \ p \ xs = \eta \ (s - \square^d)_i^{\text{locals}} \quad p + j < m \quad l = \text{fresh } \eta}{s, \eta, io, \theta \vdash (\text{CallVarF } d \ i \ j) \# (\text{Label } c) \# \text{cs} \Downarrow_{n+1}^n (s - s_{0..j-1}, \eta(l \mapsto \text{Fun } f \ m \ (p + j) \ (xs@s_{0..j-1})), io, \text{cs}, \mathcal{R}_{s,\eta}^{\text{CallVarF } d \ i \ j}) \theta} \quad (\text{STEPCALLVARFUNSATURATED})$$

$$\frac{\text{Fun } f \ m \ p \ xs = \eta \ s_0 \quad p + j \geq m}{s, \eta, io, \theta \vdash (\text{Ap } j) \# (\text{Label } c) \# \text{cs} \Downarrow_{n+1}^n ((s - s_{0..j-1} + \boxed{c, f})^{args} := xs@s_{0..j-1}, \eta, io, \Sigma_f^{Fun}, \mathcal{R}_{s, \eta}^{Ap \ j}) \theta}$$

(STEPAPSATURATED)

$$\frac{\text{Fun } f \ m \ p \ xs = \eta \ s_0 \quad p + j < m \quad l = \text{fresh } \eta}{s, \eta, io, \theta \vdash (\text{Ap } j) \# (\text{Label } c) \# \text{cs} \Downarrow_{n+1}^n (s - s_{0..j-1}, \eta(l \mapsto \text{Fun } f \ m \ (p + j) (xs@s_{0..j-1})), io, \text{cs}, \mathcal{R}_{s, \eta}^{Ap \ j}) \theta}$$

(STEPAPUNSATURATED)

4.3.4 Matching, scheduling and I/O operations

The following functions are used to retrieve wire names: $wire^{b,i}$ is the name of the i -th input wire of box b ; $wire_{b,i}$ is the name of the i -th output wire of box b . These can be calculated from the static information in *Box*. Note that box name and wire are usually retrieved from the *io* record. The function *outputable* checks whether all needed output wires are free, corresponding to the innermost loop in Figure 4.

The rules for matching always compare the heap cell, pointed to by the top-of-stack element, with a value encoded in the operation itself. If the expected kind of heap cell is found, evaluation continues with the next instruction, otherwise the next rule is tried.

In these rules several fields of the IO record are used to perform matching. The field io_{rp} points to the next rule. If a match fails with the current rule, execution will continue at $\Sigma_{io_{rp}}^{RuleSet}$, where $\Sigma^{RuleSet}$ is the rule-set of the current box. The field io_{inp} points to the wire that is examined for input. Before performing matching the HAM code must check availability of a data item on this wire. The field io_{mp} points to a position on the stack representing the current value. It is copied from the wire to the stack before a match can be performed.

$$\frac{}{s, \eta, io, \theta \vdash (\text{StartMatches}) \# \text{cs} \Downarrow_{n+1}^n (s, \eta, io \langle io_{rp} ++ \rangle, \Sigma_{io_{rp}}^{RuleSet}, \mathcal{R}_{s, \eta}^{\text{StartMatches}}) \theta}$$

(STEPSTARTMATCHES)

$$\frac{}{s, \eta, io, \theta \vdash (\text{MatchRule}) \# \text{cs} \Downarrow_{n+1}^n (s^{vals} := [], \eta, io \langle io_{inp} := 0, io_{mp} := 0, io_{rp} ++ \rangle, \Sigma_{io_{rp}}^{RuleSet}, \mathcal{R}_{s, \eta}^{\text{MatchRule}}) \theta}$$

(STEPMATCHRULE)

$$\frac{}{s, \eta, io, \theta \vdash (\text{MatchedRule}) \# \text{cs} \Downarrow_{n+1}^n (s^{vals} := [], \eta, io, \Sigma_{io_{rp}}^{RuleSet}, \mathcal{R}_{s, \eta}^{\text{MatchedRule}}) \theta}$$

(STEPMATCHEDRULE)

$$\frac{}{s, \eta, io, \theta \vdash (\text{MatchNone}) \# \text{cs} \Downarrow_{n+1}^n (s, \eta, io \langle io_{inp} ++, io_{mp} ++ \rangle, \text{cs}, \mathcal{R}_{s, \eta}^{\text{MatchNone}}) \theta}$$

(STEPMATCHNONE)

$$\frac{\exists v. \theta \ wire^{io_b, io_{inp}} = \text{Some } v}{s, \eta, io, \theta \vdash (\text{MatchAvailable}) \# \text{cs} \Downarrow_{n+1}^n (s, \eta, io \langle io_{inp} ++ \rangle, \text{cs}, \mathcal{R}_{s, \eta}^{\text{MatchAvailable}}) \theta}$$

(STEPMATCHAVAILABLETRUE)

$$\begin{array}{c}
\frac{\neg \exists v. \theta \text{ wire}^{io_b, io_{inp}} = \text{Some } v}{s, \eta, io, \theta \vdash (\text{MatchAvailable}) \# \text{cs} \Downarrow_{n+1}^n (s, \eta, io \langle io_{inp} ++ \rangle, \Sigma_{io_{rp}}^{\text{RuleSet}}, \mathcal{R}_{s, \eta}^{\text{MatchAvailable}}) \theta} \\
\text{(STEPMATCHAVAILABLEFALSE)} \\
\\
\frac{\eta s_{io_{mp}}^{\text{args}} = \text{Some } (\text{Bool } b)}{s, \eta, io, \theta \vdash (\text{MatchBool } b) \# \text{cs} \Downarrow_{n+1}^n (s, \eta, io \langle io_{mp} ++ \rangle, \text{cs}, \mathcal{R}_{s, \eta}^{\text{MatchBool } b}) \theta} \\
\text{(STEPMATCHBOOLTTRUE)} \\
\\
\frac{\neg(\eta s_{io_{mp}}^{\text{args}} = \text{Some } (\text{Bool } b))}{s, \eta, io, \theta \vdash (\text{MatchBool } b) \# \text{cs} \Downarrow_{n+1}^n (s, \eta, io \langle io_{mp} ++ \rangle, \Sigma_{io_{rp}}^{\text{RuleSet}}, \mathcal{R}_{s, \eta}^{\text{MatchBool } b}) \theta} \\
\text{(STEPMATCHBOOLFALSE)} \\
\\
\frac{\eta s_{io_{mp}}^{\text{args}} = \text{Some } (\text{Char } c)}{s, \eta, io, \theta \vdash (\text{MatchChar } c) \# \text{cs} \Downarrow_{n+1}^n (s, \eta, io \langle io_{mp} ++ \rangle, \text{cs}, \mathcal{R}_{s, \eta}^{\text{MatchChar } c}) \theta} \\
\text{(STEPMATCHCHARTRUE)} \\
\\
\frac{\neg(\eta s_{io_{mp}}^{\text{args}} = \text{Some } (\text{Char } c))}{s, \eta, io, \theta \vdash (\text{MatchChar } c) \# \text{cs} \Downarrow_{n+1}^n (s, \eta, io \langle io_{mp} ++ \rangle, \Sigma_{io_{rp}}^{\text{RuleSet}}, \mathcal{R}_{s, \eta}^{\text{MatchChar } c}) \theta} \\
\text{(STEPMATCHCHARFALSE)} \\
\\
\frac{\eta s_{io_{mp}}^{\text{args}} = \text{Some } (\text{Str } x)}{s, \eta, io, \theta \vdash (\text{MatchString } x) \# \text{cs} \Downarrow_{n+1}^n (s, \eta, io \langle io_{mp} ++ \rangle, \text{cs}, \mathcal{R}_{s, \eta}^{\text{MatchString } x}) \theta} \\
\text{(STEPMATCHSTRINGTRUE)} \\
\\
\frac{\neg(\eta s_{io_{mp}}^{\text{args}} = \text{Some } (\text{Str } x))}{s, \eta, io, \theta \vdash (\text{MatchString } x) \# \text{cs} \Downarrow_{n+1}^n (s, \eta, io \langle io_{mp} ++ \rangle, \Sigma_{io_{rp}}^{\text{RuleSet}}, \mathcal{R}_{s, \eta}^{\text{MatchString } x}) \theta} \\
\text{(STEPMATCHSTRINGFALSE)} \\
\\
\frac{\eta s_{io_{mp}}^{\text{args}} = \text{Some } (\text{Int } i)}{s, \eta, io, \theta \vdash (\text{MatchInt } i) \# \text{cs} \Downarrow_{n+1}^n (s, \eta, io \langle io_{mp} ++ \rangle, \text{cs}, \mathcal{R}_{s, \eta}^{\text{MatchInt } i}) \theta} \\
\text{(STEPMATCHINTTRUE)} \\
\\
\frac{\neg(\eta s_{io_{mp}}^{\text{args}} = \text{Some } (\text{Int } i))}{s, \eta, io, \theta \vdash (\text{MatchInt } i) \# \text{cs} \Downarrow_{n+1}^n (s, \eta, io \langle io_{mp} ++ \rangle, \Sigma_{io_{rp}}^{\text{RuleSet}}, \mathcal{R}_{s, \eta}^{\text{MatchInt } i}) \theta} \\
\text{(STEPMATCHINTFALSE)} \\
\\
\frac{\exists xs. \eta s_{io_{mp}}^{\text{args}} = \text{Some } (\text{Tuple } m \text{ xs})}{s, \eta, io, \theta \vdash (\text{MatchTuple } m) \# \text{cs} \Downarrow_{n+1}^n (s, \eta, io \langle io_{mp} ++ \rangle, \text{cs}, \mathcal{R}_{s, \eta}^{\text{MatchTuple } m}) \theta} \\
\text{(STEPMATCHTUPLETRUE)}
\end{array}$$

$$\frac{\neg(\exists xs. \eta s_{io_{mp}}^{args} = \text{Some} (\text{Tuple } m \ xs))}{s, \eta, io, \theta \vdash (\text{MatchTuple } m) \# \text{cs} \Downarrow_{n+1}^n (s, \eta, io \langle io_{mp} \ ++ \ \rangle, \Sigma_{io_{rp}}^{RuleSet}, \mathcal{R}_{s, \eta}^{\text{MatchTuple } m}) \ \theta} \quad (\text{STEPMATCHTUPLEFALSE})$$

$$\frac{\exists xs. \eta s_{io_{mp}}^{args} = \text{Some}(\text{Constr } i \ j \ xs)}{s, \eta, io, \theta \vdash (\text{MatchCon } i \ j) \# \text{cs} \Downarrow_{n+1}^n (s, \eta, io \langle io_{mp} \ ++ \ \rangle, \text{cs}, \mathcal{R}_{s, \eta}^{\text{MatchCon } i \ j}) \ \theta} \quad (\text{STEPMATCHCONTRUE})$$

$$\frac{\neg(\exists xs. \eta s_{io_{mp}}^{args} = \text{Some} (\text{Constr } i \ j \ xs))}{s, \eta, io, \theta \vdash (\text{MatchCon } i \ j) \# \text{cs} \Downarrow_{n+1}^n (s, \eta, io \langle io_{mp} \ ++ \ \rangle, \Sigma_{io_{rp}}^{RuleSet}, \mathcal{R}_{s, \eta}^{\text{MatchCon } i \ j}) \ \theta} \quad (\text{STEPMATCHCONFALSE})$$

$$\frac{\exists l. \eta s_{io_{mp}}^{args} = \text{Some}(\text{Exn } x \ l)}{s, \eta, io, \theta \vdash (\text{MatchExn } x) \# \text{cs} \Downarrow_{n+1}^n (s, \eta, io \langle io_{mp} \ ++ \ \rangle, \text{cs}, \mathcal{R}_{s, \eta}^{\text{MatchExn } x}) \ \theta} \quad (\text{STEPMATCHEXNTRUE})$$

$$\frac{\neg(\exists l. \eta s_{io_{mp}}^{args} = \text{Some} (\text{Exn } x \ l))}{s, \eta, io, \theta \vdash (\text{MatchExn } x) \# \text{cs} \Downarrow_{n+1}^n (s, \eta, io \langle io_{mp} \ ++ \ \rangle, \Sigma_{io_{rp}}^{RuleSet}, \mathcal{R}_{s, \eta}^{\text{MatchExn } x}) \ \theta} \quad (\text{STEPMATCHEXNFALSE})$$

The following rules unpack structured data, i.e. they bring the contents of a tuple, constructor or function closure onto the stack.

$$\frac{\exists l \ m \ xs. s_0 = l \ \wedge \ \eta \ l = \text{Some} (\text{Tuple } m \ xs) \ \wedge \ s' = (s - \top) + xs}{s, \eta, io, \theta \vdash (\text{Unpack}) \# \text{cs} \Downarrow_{n+1}^n (s', \eta, io, \text{cs}, \mathcal{R}_{s, \eta}^{\text{Unpack}}) \ \theta} \quad (\text{STEPUNPACKTUPLE})$$

$$\frac{\exists l \ t \ m \ xs. s_0 = l \ \wedge \ \eta \ l = \text{Some} (\text{Constr}_t \ m \ xs) \ \wedge \ s' = (s - \top) + xs}{s, \eta, io, \theta \vdash (\text{Unpack}) \# \text{cs} \Downarrow_{n+1}^n (s', \eta, io, \text{cs}, \mathcal{R}_{s, \eta}^{\text{Unpack}}) \ \theta} \quad (\text{STEPUNPACKCON})$$

$$\frac{\exists l \ f \ m \ p \ xs. s_0 = l \ \wedge \ \eta \ l = \text{Some} (\text{Fun } f \ m \ p \ xs) \ \wedge \ s' = (s - \top) + xs}{s, \eta, io, \theta \vdash (\text{Unpack}) \# \text{cs} \Downarrow_{n+1}^n (s', \eta, io, \text{cs}, \mathcal{R}_{s, \eta}^{\text{Unpack}}) \ \theta} \quad (\text{STEPUNPACKFUN})$$

The following operations copy box input into the heap, check whether a box can write to its output wires, and perform the actual write operation. Note that the compiler must ensure that all `Write` operations are guarded by `CheckOutput` operations.

$$\frac{l = \text{fresh} (\text{dom } \eta)}{s, \eta, io, \theta \vdash (\text{CopyInput } m) \# \text{cs} \Downarrow_{n+1}^n (s + l, \eta(l \mapsto (\theta \ \text{wire}^{io_b, m})), io, \text{cs}, (\mathcal{R}_{s, \eta}^{\text{CopyInput } m})) \ \theta} \quad (\text{STEPCOPYINPUT})$$

$$\frac{s, \eta, io, \theta \vdash (\text{Consume } m) \# \text{cs} \Downarrow_{n+1}^n (s, \eta, io, \text{cs}, \mathcal{R}_{s,\eta}^{\text{Consume } m}) \theta(\text{wire}^{io_b, m} \mapsto \text{None})}{(\text{STEPCONSUME})}$$

$$\frac{\exists v. \theta \text{ wire}^{io_b, m} = \text{Some } v}{s, \eta, io, \theta \vdash (\text{MaybeConsume } m) \# \text{cs} \Downarrow_{n+1}^n (s, \eta, io, \text{cs}, \mathcal{R}_{s,\eta}^{\text{MaybeConsume } m}) \theta(\text{wire}^{io_b, m} \mapsto \text{None})} (\text{STEPMAYBECONSUMETRUE})$$

$$\frac{\neg \exists v. \theta \text{ wire}^{io_b, m} = \text{Some } v}{s, \eta, io, \theta \vdash (\text{MaybeConsume } m) \# \text{cs} \Downarrow_{n+1}^n (s, \eta, io, \text{cs}, \mathcal{R}_{s,\eta}^{\text{MaybeConsume } m}) \theta} (\text{STEPMAYBECONSUMEFALSE})$$

$$\frac{\text{outputable } io \ \theta}{s, \eta, io, \theta \vdash (\text{CheckOutputs}) \# \text{cs} \Downarrow_{n+1}^n (s, \eta, io, \text{cs}, \mathcal{R}_{s,\eta}^{\text{CheckOutputs}}) \theta} (\text{STEPCHECKOUTPUTSTRUE})$$

$$\frac{\neg (\text{outputable } io \ \theta)}{s, \eta, io, \theta \vdash (\text{CheckOutputs}) \# \text{cs} \Downarrow_{n+1}^n (s, \eta, io \mid \text{blocked} := \text{true} \mid), (\text{Schedule}) \# (\text{CheckOutputs}) \# \text{cs}, \mathcal{R}_{s,\eta}^{\text{CheckOutputs}}) \theta} (\text{STEPCHECKOUTPUTSFALSE})$$

$$\frac{l = s_0 \quad \eta \ l = \text{Some } x}{s, \eta, io, \theta \vdash (\text{Write } m) \# \text{cs} \Downarrow_{n+1}^n ((s - \top), \eta, io, \text{cs}, \mathcal{R}_{s,\eta}^{\text{Write } m}) \theta(\text{wire}_{io_b, m} \mapsto \text{Some } x)} (\text{STEPWRITE})$$

The operations **Input** and **Output** interact with the outside world, enabling character input and output. Since we are mainly interested in the internal behaviour of the system, in particular its resource consumption, we do not model the exact values and use a dummy character value \perp instead. Costs are still accurately modelled, provided they are constant for all possible character values.

$$\frac{l = \text{fresh } (\text{dom } \eta)}{s, \eta, io, \theta \vdash (\text{Input}) \# \text{cs} \Downarrow_{n+1}^n (s + l, \eta(l \mapsto (\text{Char } \perp)), io, \text{cs}, \mathcal{R}_{s,\eta}^{\text{Input}}) \theta} (\text{STEPINPUT})$$

$$\frac{l = \text{fresh } (\text{dom } \eta)}{s, \eta, io, \theta \vdash (\text{Output}) \# \text{cs} \Downarrow_{n+1}^n (s - \top, \eta, io, \text{cs}, \mathcal{R}_{s,\eta}^{\text{Output}}) \theta} (\text{STEPOUTPUT})$$

$$\frac{s, \eta, io, \theta \vdash (\text{WithinStack } h \ p) \# \text{cs} \Downarrow_{n+1}^n (s + io_{\text{splim}}, \eta, io \mid \text{splim} := |s| + p \mid), \text{cs}, \mathcal{R}_{s,\eta}^{\text{WithinStack } h \ p}) \theta}{(\text{STEPWITHINSTACKSPACE})}$$

$$\frac{s, \eta, io, \theta \vdash (\text{WithinHeap } h \ p) \# \text{cs} \Downarrow_{n+1}^n (s + io_{\text{hplim}}, \eta, io \mid \text{hplim} := | \text{dom } \eta | + p \mid), \text{cs}, \mathcal{R}_{s,\eta}^{\text{WithinHeap } h \ p}) \theta}{(\text{STEPWITHINHEAPSPACE})}$$

$$\frac{}{s, \eta, io, \theta \vdash (\text{DoneWithinStack } h) \# \text{cs} \Downarrow_{n+1}^n (s - \top, \eta, io \mid \text{splim} := s_0 \mid), \text{cs}, \mathcal{R}_{s,\eta}^{\text{DoneWithinStack } h}) \theta} \quad (\text{STEPDONEYWITHINSTACK})$$

$$\frac{}{s, \eta, io, \theta \vdash (\text{DoneWithinHeap } h) \# \text{cs} \Downarrow_{n+1}^n (s - \top, \eta, io \mid \text{hplim} := s_0 \mid), \text{cs}, \mathcal{R}_{s,\eta}^{\text{DoneWithinHeap } h}) \theta} \quad (\text{STEPDONEYWITHINHEAP})$$

$$\frac{}{s, \eta, io, \theta \vdash (\text{Within } h \text{ t}) \# \text{cs} \Downarrow_{n+1}^n (s + t, io, \text{cs}, \mathcal{R}_{s,\eta}^{\text{Within } h \text{ t}}) \theta} \quad (\text{STEPWITHIN})$$

$$\frac{}{s, \eta, io, \theta \vdash (\text{DoneWithin}) \# \text{cs} \Downarrow_{n+1}^n (s, \text{unsetTimer } io, \text{cs}, \mathcal{R}_{s,\eta}^{\text{DoneWithin}}) \theta} \quad (\text{STEPDONEYWITHIN})$$

$$\frac{}{s, \eta, io, \theta \vdash (\text{RaiseWithin } h \text{ t}) \# \text{cs} \Downarrow_{n+1}^n (s - s_{0..2-1}, \text{setTimer } io \ s_0 \ s_1, \text{cs}, \mathcal{R}_{s,\eta}^{\text{RaiseWithin } h \text{ t}}) \theta} \quad (\text{STEPRAISEWITHIN})$$

4.3.5 Expression-level semantics

As a small-step semantics the above rules specify one step in the evaluation, depending on the next element in the instruction list cs . The semantics on expression level is the transitive closure over this relation in the following sense.

$$\frac{c \neq \text{Schedule} \quad s, \eta, io, \theta \vdash (c \# \text{cs}) \Downarrow_{m+1}^m (s'', \eta'', io'', \text{cs}'', p'') \theta'' \quad s'', \eta'', io'', \theta'' \vdash \text{cs}'' \rightsquigarrow_n (s', \eta', io', \text{cs}', p') \theta'}{s, \eta, io, \theta \vdash (c \# \text{cs}) \rightsquigarrow_{m+n+1} (s', \eta', io', \text{cs}', p'' + p') \theta'} \quad (\text{SSEM})$$

The continuation cs' is part of the result, because the computation may block in the CHECKOUTPUTS rule. The costs of the entire computation is the sum of the costs of the components, using the $+$ operation as defined by the relevant resource algebra. The step counters m, n are only needed for technical reasons (to enable induction over this step counter) and do not represent resource consumption in any way. Note that the rule is restricted to the case where the first instruction is different from **Schedule**. This assures that in the case of a **Schedule** control is yielded to the scheduler, which then picks the next runnable box, removes the initial **Schedule** and continues with evaluating the box.

4.3.6 Exceptions

The above rule for the expression-level semantics is slightly simplified in the sense that it doesn't check for (synchronous) exceptions. Such checks can be added to the rule, provided that the resource algebra used in the semantics models the resource that might become empty. For example, in the case of heap space a pre-condition of the form

$$| \text{hplim} - \text{hp} | < \mathcal{R}_{s,\eta}^c$$

can be added, if \mathcal{R} models heap consumption. If this condition is false, another rule can be added for raising a heap-overflow exception.

4.4 System level

As static information we use for each box a mapping *Wiring* that maps a wire index to the corresponding box name and wire index it is connected to. The tables Σ^{Fun} , Σ^{Lab} and Σ^{Exn} are used to map function names, labels and exceptions to their corresponding instruction sequences. The $\Sigma^{RuleSet}$ is a partial mapping from natural numbers to instruction sequences, containing the code for the rules in the box. Here the domain needs to be a total order to allow reordering. A *Box* is then represented as a rule-set (a mapping from natural numbers to instruction sequences) and such a wiring.

$$\begin{aligned}
\Sigma^{RuleSet} &\equiv \mathbb{N} \rightsquigarrow_f Instr\ list \\
\Sigma^{Fun} &\equiv Function \Rightarrow Instr\ list \\
\Sigma^{Lab} &\equiv Label \Rightarrow Instr\ list \\
\Sigma^{Exn} &\equiv Exception \Rightarrow Instr\ list \\
Wiring &\equiv Wire \Rightarrow (BName \times Wire) \\
Box &\equiv (\Sigma^{RuleSet} \times \Sigma^{Lab} \times \Sigma^{Exn} \times Wiring)
\end{aligned}$$

We model the dynamic state of the entire system as follows. The entire system (*System*) is a triple (bs, β, θ) , where *bs* is a collection of box names still to be processed in the current iteration, β is a (total) mapping of box-names to box states, and θ is a (total) mapping of wire-names to wire states. A *box state* (*BState*) is a quadrupel (s, η, io, cs) , where *s* is the stack, η is the heap, *io* is the IO record, and *cs* is the continuation, i.e. the code still to be processed by the box. The *wire state* (*WState*) is simply an optional heap value, i.e. either a proper heap value (*Some x*) or an empty flag (*None*).

$$\begin{aligned}
WState &\equiv HVal\ option \\
BState &\equiv Stack \times Heap \times IO \times Instr\ list \\
System &\equiv BName\ set \times (BName \Rightarrow BState) \times (WName \Rightarrow WState)
\end{aligned}$$

NeededIn maps for each box, each rule in the ruleset to the set of needed inputs, i.e. a set of wires on which input must be available for the rule to fire. *NeededOut* maps for each box, each rule in the ruleset to the set of needed outputs, i.e. a set of wires to which this rule will write. These tables are filled with the compiler directives shown in Figure 12.

$$\begin{aligned}
NeededIn &:: Box \Rightarrow \mathbb{N} \Rightarrow Wire\ set \\
NeededOut &:: Box \Rightarrow \mathbb{N} \Rightarrow Wire\ set
\end{aligned}$$

The global table *RuleTab* maps a box name to a *Box* structure, containing its static information. Here we are only interested in the rule-set component of the box structure and we use the shorthand *RuleTab₁* for the first projection on the corresponding rule tab entry.

$$\begin{aligned}
RuleTab &:: BName \Rightarrow Box \\
RuleTab_1 &:: BName \Rightarrow RuleSet
\end{aligned}$$

We need the following additional definitions to specify the semantics on system level:

$$\begin{aligned}
\text{outputable} &:: IO \Rightarrow (WName \Rightarrow WState) \Rightarrow bool \\
\text{outputable } io \theta &\equiv (\forall w. w \in (NeededOut (RuleTab io_b) io_{rp}) \longrightarrow \\
&\quad \exists v. \theta \text{ wire}_{io_b, m} = Some v) \\
\\
\text{runnable} &:: (WName \Rightarrow WState) \Rightarrow BState \Rightarrow bool \\
\text{runnable } \theta (s, \eta, io, cs) &\equiv \neg(io_{blocked}) \wedge \\
&\quad (\exists r. r \in dom (RuleTab_1 io_b) \wedge \\
&\quad (\forall w. w \in (NeededIn (RuleTab io_b) r) \longrightarrow \exists v. \theta \text{ wire}^{io_b, w} = Some v)) \\
\\
\text{restartable} &:: (WName \Rightarrow WState) \Rightarrow BState \Rightarrow bool \\
\text{restartable } \theta (s, \eta, io, cs) &\equiv io_{blocked} \wedge \\
&\quad (\forall w. w \in (NeededOut (RuleTab io_b) io_{rp}) \longrightarrow \theta \text{ wire}^{io_b, w} = None)
\end{aligned}$$

The transition relation on system level has the form $s \rightarrow s'$ meaning that in one step, the system state s proceeds to state s' . The transitive closure over the small-step semantics on expression level, written as \rightsquigarrow_m , is defined in the previous section. We use \emptyset to denote the empty heap, mapping all locations to *None*, and the empty stack, an empty list of frames.

The rules defining the HAM semantics on system level are as follows. We use S as a short-hand for bs, β, θ, p . We distinguish between three cases: if the box, selected from the scheduling queue, is a runnable box, then the continuation of the box must be empty, and execution continues with the next rule in the box; note that in this case the box starts with an empty heap and stack; if the box, selected from the scheduling queue, is a restartable box, then the continuation must start with a **Schedule**, which is removed from the code and the remaining code is executed on expression level; if the scheduling queue is empty, it is re-filled using the \mathcal{U}_S operation. The costs recorded in the last field of the state, combine the costs of the individual operations on scheduling level referring to constants such as \mathcal{R}_S^\ominus as explained in the following sub-section.

$$\frac{
\begin{array}{l}
bn \in_{bs, \beta, \theta, p} bs \quad (s, \eta, io, []) = \beta bn \quad \text{runnable } \theta (s, \eta, io, []) \\
cs = RuleTab_1 io_b io_{rp} \quad \emptyset, \emptyset, io, \theta \vdash cs \rightsquigarrow_m (s', \eta', io', cs', p') \theta'
\end{array}
}{
(bs, \beta, \theta, p) \rightarrow ((bs \ominus_{bs, \beta, \theta, p} bn) \oplus_{bs, \beta, \theta, p} bn, \beta(bn \mapsto (s', \eta', io', cs')), \theta', p + \mathcal{R}_S^\ominus + \mathcal{R}_S^\oplus + p' + \mathcal{R}_S^\oplus)
} \text{(SCHEDULERUNNABLE)}$$

$$\frac{
\begin{array}{l}
bn \in_{bs, \beta, \theta, p} bs \quad (s, \eta, io, cs'') = \beta bn \quad \text{restartable } \theta (s, \eta, io, cs'') \\
(\text{Schedule}\#cs) = cs'' \quad s, \eta, io, \theta \vdash cs \rightsquigarrow_m (s', \eta', io', cs', p') \theta'
\end{array}
}{
(bs, \beta, \theta, p) \rightarrow ((bs \ominus_{bs, \beta, \theta, p} bn) \oplus_{bs, \beta, \theta, p} bn, \beta(bn \mapsto (s', \eta', io', cs')), \theta', p + \mathcal{R}_S^\ominus + \mathcal{R}_S^\oplus + p' + \mathcal{R}_S^\oplus)
} \text{(SCHEDULERESTARTABLE)}$$

$$\frac{
\neg \exists bn \in_{bs, \beta, \theta, p} bs
}{
(bs, \beta, \theta, p) \rightarrow (\mathcal{U}_{bs, \beta, \theta, p}, \mathcal{C}_B \beta \theta, \mathcal{C}_W \beta \theta, p + \mathcal{R}_S^\mathcal{U} + \mathcal{R}_S^{\mathcal{C}_B} + \mathcal{R}_S^{\mathcal{C}_W})
} \text{(SCHEDULENIL)}$$

Note that this specification of the system level is parameterised by the concrete membership relation \in_S and by the operations for picking the next box to be executed \ominus_S and putting it back into the system \oplus_S . In the case where there is no next box available for scheduling, we assume an operation for generating a new collection of schedulable boxes, \mathcal{U}_S , an operation for modifying the box mapping, \mathcal{C}_B , and an operation for modifying the wire mapping, \mathcal{C}_W .

We expect implementations of Hume to use well-known scheduling policies such as round-robin scheduling based on some total ordering of the boxes derived from the source code (this is what's

used in the reference implementation of the HAM). However, we leave the concrete realisation of the scheduling operations ($\in_S, \ominus_S, \oplus_S$) and of the synchronisation operations ($\mathcal{U}_S, \mathcal{C}_B, \mathcal{C}_W$) to the implementor. Naturally, the overall costs of the system will be depend on the scheduling policy. Thus, all these parametric operations have to come with corresponding cost functions, describing how the costs coming out of the expression level shall be combined.

The following operations implement round-robin scheduling. We initially sort the boxes by a global ordering derived from the source code. The membership relation only examines the head of the list. Picking the next element means taking the head of the list. Putting a box back means adding it to the end of the list. Since the length of the scheduling queue never decreases no global synchronisation is needed, and the definition of $\mathcal{U}_S, \mathcal{C}_B, \mathcal{C}_W$ is arbitrary.

$$\begin{aligned}
\mathcal{D} &\equiv BName\ list \\
b \in_S bs &\equiv b = hd\ bs \\
\ominus_S &\equiv \lambda bs\ b.tl\ bs \\
\oplus_S &\equiv \lambda bs\ b.bs@[b] \\
\mathcal{U}_S &\equiv [] \\
\mathcal{C}_B\ \beta\ \theta &\equiv \beta \\
\mathcal{C}_W\ \beta\ \theta &\equiv \theta
\end{aligned}$$

4.4.1 Cost modelling on system level

The cost model on the system level is realised similarly to the expression level. We use a resource algebra $(R, +, \leq)$, where R must be the same data structure as on the expression level, but $+$ and \leq may be different operations. For example, in a parallel implementation of the HAM the overall time consumption is not necessarily the sum over all the time consumptions of all boxes plus system time. The constants in the system-level resource algebra must correspond to the basic scheduling operations: \mathcal{R}_S^\ominus for getting a box name from the scheduling queue, \mathcal{R}_S^\oplus for putting a box back into the scheduling queue, \mathcal{R}_S^\in for testing whether a box name is in the scheduling queue; \mathcal{R}_S^u for re-filling the scheduling queue; $\mathcal{R}_S^{\mathcal{C}_B}$ for synchronising all boxes, and $\mathcal{R}_S^{\mathcal{C}_W}$ for synchronising all wires. Note, that in general all these costs depend on the entire system state. In the case of modelling time consumption in a round-robin scheduler, the last three constants would be 0; \mathcal{R}_S^\in , \mathcal{R}_S^\ominus , \mathcal{R}_S^\oplus would be the costs for testing for list membership, extracting an element from the list and adding an element to the end of the list.

5 Summary

As specification of the HAM behaviour we presented a 2 level, small-step operational semantics as well as a reference implementation in pseudo C. The former will act as the basis for reasoning about resource consumption. The latter is more concrete and thus more clearly exhibits the costs involved in executing HAM instructions. These costs are captured in the operational semantics via operations of a resource algebra applied to counters that are part of the system state. Since the rules of the semantics only refer to these abstract operations of the resource algebra, they can be chosen independently from the rules of the semantics, and we have shown instances for heap space, stack space and time consumption, thus giving cost models of the HAM for these resources. On top of the small-step semantics we have defined a system-level semantics, that deals with the scheduling of multiple boxes in Hume. Here we parameterise the semantics with concrete scheduling operations, accompanied again by a system-level resource algebra for composing costs on system level. Thus, we arrive at a simple, yet flexible, formalisation of HAM behaviour and resource consumption on system level.

References

- [ABM06] D. Aspinall, L. Beringer, and A. Momigliano. Optimisation Validation. In *COCV06 — Fifth Workshop on Compiler Optimization Meets Compiler Verification*, Vienna, Austria, April 2, 2006. To appear in ENTCS.
- [Aug87] L. Augustsson. *Compiling Lazy Functional Languages, Part II*. PhD thesis, Dept. of Computer Science, Chalmers University of Technology, Göteborg, 1987.
- [BLOP96] S. Breitinger, R. Loogen, Y. Ortéga Mallén, and Ricardo Peña Marí. Eden — the Paradise of Concurrent Functional Programming. In *Europar'96*, volume 1123 of *LNCS*, 1996.
- [Ham03] K. Hammond. An Abstract Machine for Resource Bounded Computations in Hume. In *IFL03 — Intl Workshop on Implementation of Functional Languages*, pages 79–94, Edinburgh, Scotland, September 2003. Draft proceedings.
- [HM] K. Hammond and G. Michaelson. The Hume Report. Available from: <http://www-fp.dcs.st-and.ac.uk/hume/report/hume-report.ps>. Version 0.3.
- [HM02] K. Hammond and G.J. Michaelson. Predictable Space Behaviour in FSM-Hume. In *14th Intl Workshop on Implementation of Functional Languages*, pages 386–403, Madrid, Spain, September 2002.
- [JLH07] S. Jost, H-W. Loidl, and K. Hammond. Report on WCET Analysis. EmBounded Project Deliverable, February 2007. Deliverable D14.
- [Lan64] P. Landin. The Mechanical Evaluation of Expressions. *The Computer Journal*, 6(4):308–320, January 1964.
- [LY99] T. Lindholm and F. Yellin. *The Java Virtual Machine Specification*. Addison-Wesley, second edition edition, April 1999.
- [Mac] D. MacQueen. Standard ML of New Jersey. Web page. <http://www.smlnj.org/>.
- [Pey92] S. L. Peyton Jones. Implementing Lazy Functional Languages on Stock Hardware: the Spineless Tagless G-Machine. *J. of Functional Programming*, 2(2):127–202, 1992.
- [PL92] S. L. Peyton Jones and D. Lester. *Implementing Functional Languages: a Tutorial*. Prentice-Hall, 1992.