


# An Esterel Virtual Machine

Stephen A. Edwards

Columbia University

Octopi Workshop  
Chalmers University of Technology  
Gothenburg, Sweden  
December 2018

# An Esterel Virtual Machine



Goal: Run big Esterel programs in memory-constrained settings.

Our target: the Hitachi H8-based RCX Microcontroller for Lego Mindstorms

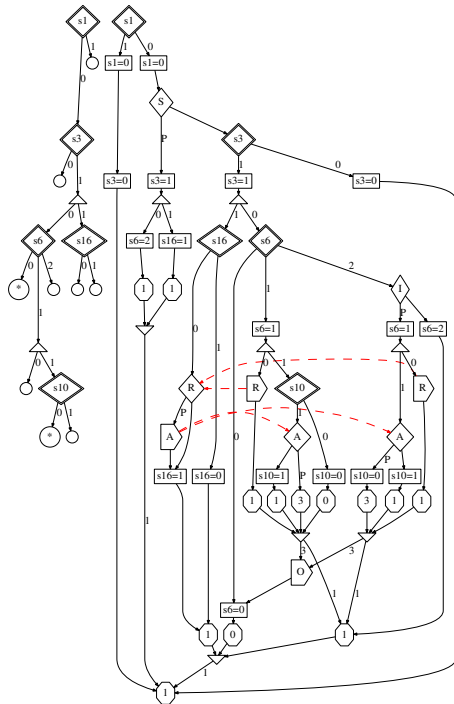
[SLAP 2006]

# An Example

```

module Example:
input I, S;
output O;
signal R,A in
  every S do
    await I;
    weak abort
      sustain R
    when immediate A;
    emit O
  ||
  loop
    pause; pause;
    present R then
      emit A
    end present
  end loop
end every
end signal
end module

```



# Challenges

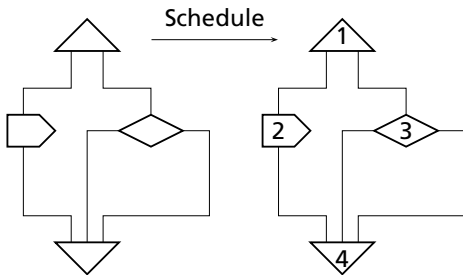
Esterel's semantics require any implementation to deal with three issues:

- ▶ Concurrent execution of sequential threads of control within a cycle
- ▶ The scheduling constraints among these threads due to communication dependencies
- ▶ How control state is updated between cycles


## How did we handle them?

- ▶ A virtual machine specifically designed to support Esterel features
- ▶ A sequentializing algorithm
- ▶ Conversion from GRC to BAL and then to a compact byte code

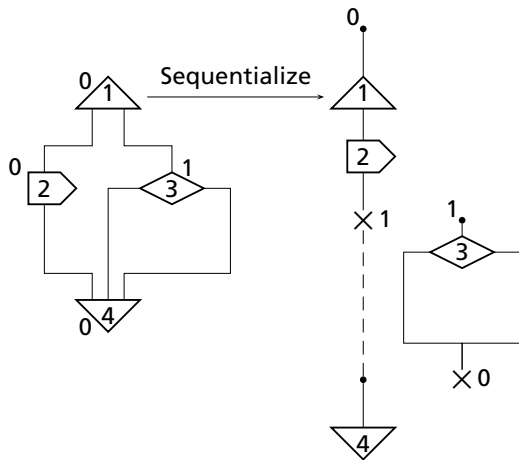
# Phase 1: Schedule



## Phase 2: Assign Threads

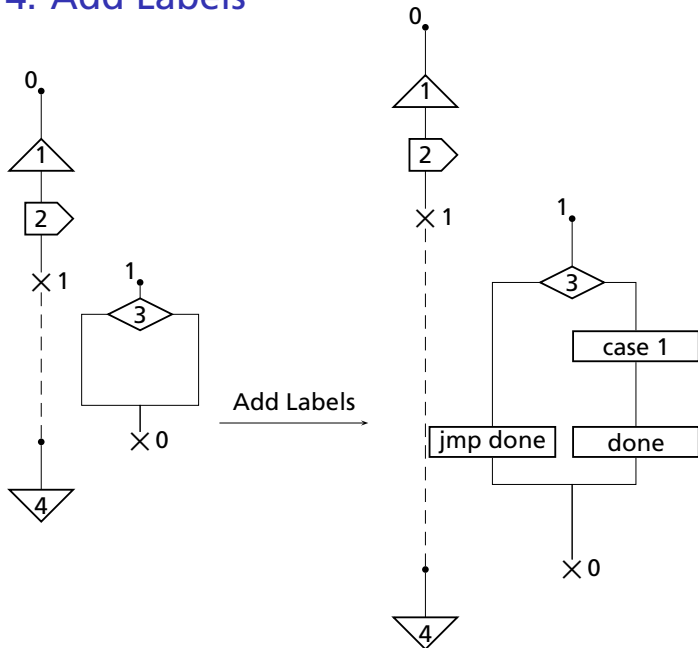


## Phase 3: Sequentialize

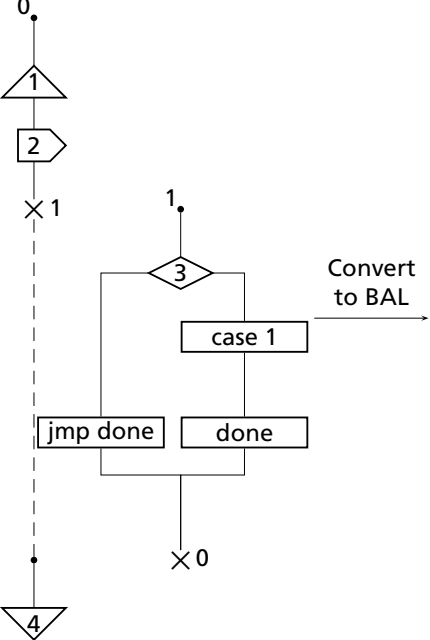




## Phase 4: Add Labels



# Phase 5: Convert to BAL



```

t0
  STHR 1 t1
  EMT 1
  SWC 1
  STHR 1 NR1
  END
NR1
  SWCU
t1
  TWB 2 2 case_1
  JMP done
case_1
done
  SWC 0
  
```

Convert to BAL

## Phase 6: Convert to Byte Code

t0			
	STHR 1 t1	Convert	07 01 00 0e
	EMT 1	to	04 01
	SWC 1	byte code	05 01
	STHR 1 NR1	→	07 01 00 0d
	END		03
NR1			
	SWCU		0c
t1			
	TWB 2 2 case_1		49 02 00 15
	JMP done		06 00 15
case_1			
done			
	SWC 0		05 00


# Sequential Code Generation

1. Schedule the nodes in the graph
2. Assign thread numbers
3. Sequentialize the graph
4. Set the execution path by adding labels
5. Convert to BAL
6. Assemble to produce bytecode

# Sequentialization




# Sequentialization




The dotted line labeled F represents the frontier. The frontier starts at the top of the graph.

# Sequentialization



The frontier moves down a node at a time in scheduled order.


# Sequentialization



When a node is in the same thread as the most recently moved one, it is simply moved above the frontier.




# Sequentialization



However, when the next node is from a different thread, a switch is added to the previous thread and an active point is added to the new thread just above the just-moved node.

# Sequentialization



The algorithm is complete when the frontier has swept across all nodes in scheduled order.


# Sequentializing Algorithm

```
1: for each thread  $t$  in  $G$  do
2:   create new active point  $p$ 
3:   copy first node  $n$  of  $t$  in  $G$  to  $n'$  new node in  $G'$ 
4:   connect  $p$  and  $n'$ 
5:   add  $p$  to  $P[t]$  and add  $n'$  to  $A[t]$ 
6:  $t' =$  the first thread
7: for each node  $n$  in scheduled order do
8:    $t$  is thread of  $n$ 
9:   if  $t \neq t'$  then
10:    for each parent  $p$  in  $P[t']$  do
11:      for each successor  $c$  of  $p$  in  $A[t']$  do
12:        create switch node  $s$  from  $t'$  to  $t$  and connect  $s$  between  $p$  and  $c$ 
13:      replace  $P[t']$  with the set of new switch nodes
14:    move  $n$  to  $P[t]$  and remove it from  $A[t]$ 
15:    for each unreached successor  $c$  of  $n$  do
16:      copy  $c$  to  $c'$  new node in  $G'$ 
17:      if  $n$  is a fork then
18:        add child to  $A[\text{thread of } c]$ 
19:      else
20:        add child to  $A[t]$ 
21:     $t' = t$  {remember the last thread}
```

## Why VM?

- ▶ Goal: constrained-memory environment
- ▶ Instruction set has direct support for Esterel constructs like concurrency, preemption, and signals
- ▶ E.g., a context switch can be specified in just two bytes

# VM Details



# VM Details

- ▶ Signal status registers
- ▶ Completion code registers
- ▶ Per-thread program counters
- ▶ Inter-instant state-holding registers

## VM: Signal, State, and Thread

<b>Opcode</b>	<b>Description</b>	<b>Encoding</b>
EMT	Emit a Signal	04 RR
SSIG	Clear Signal	0A RR
SSTT	Set State	0B RR VV
STHR	Set Thread	07 TT HH LL

## VM: Control Flow Instructions

<b>Opcode</b>	<b>Description</b>	<b>Encoding</b>
END	Tick End	03
JMP	Jump	06 HH LL
NOP	No Operation	01



## VM: Branch, Switch, Terminate

Opcode	Description	Encoding
MWB	Multiway Branch (State)	2D NL RR HH2 LL2 ...
	Multiway Branch (Comp.)	4D NL RR HH2 LL2 ...
TWB	Two Way Branch (State)	29 RR HH LL
	Two Way Branch (Signal)	49 RR HH LL
	Two Way Branch (Comp.)	69 RR HH LL
SWC	Switch Thread	05 TT
SWCU	Switch Unknown	0C
TRM	Set Completion Code for Join	08 RR VV


# VM: Context Switch

```
...
switch(opcode & 0x1F){
    ...
    case SWC:
        // Increment the program counter
        ++pc;
        // Store the current thread as the last thread
        last_thread = current_thread;
        // Get the next thread
        current_thread = *pc;
        // Increment the program counter
        ++pc;
        // Store old pc associated with the old thread
        threads[last_thread] = pc;
        // Load the pc associated with the new thread
        pc = threads[current_thread];
        break;
    ...
}
```

# VM: Switch Unknown


```
...
case SWCU:
    // Make the thread stored in last_thread, the current thread
    temp = current_thread;
    current_thread = last_thread;
    last_thread = temp;
    // Store old pc
    threads[last_thread] = pc;
    // Load new pc
    pc = threads[current_thread];
    break;
...
```

## VM in action




# VM in action

```
    t0
00:  STHR 1 t1
04:  EMT 1
06:  SWC 1
08:  STHR 1 NR1
12:  SWC 1
14:  END
    NR1
15:  SWCU
    t1
16:  TWB 2 1 case_1
19:  JMP done
    case_1
    done
22:  SWC 0
```



# VM in action

```
t0
00:  STHR 1 t1
04:  EMT 1
06:  SWC 1
08:  STHR 1 NR1
12:  SWC 1
14:  END
NR1
15:  SWCU
t1
16:  TWB 2 1 case_1
19:  JMP done
case_1
done
22:  SWC 0
```



# VM in action

```
t0
00:  STHR 1 t1
04:  EMT 1
06:  SWC 1
08:  STHR 1 NR1
12:  SWC 1
14:  END
NR1
15:  SWCU
t1
16:  TWB 2 1 case_1
19:  JMP done
case_1
done
22:  SWC 0
```

pc = 6  
last\_thread = 0

Threads	Signals				
<table border="1"><tr><td>0</td></tr><tr><td>16</td></tr></table>	0	16	<table border="1"><tr><td>0</td></tr><tr><td>1</td></tr></table>	0	1
0					
16					
0					
1					
States	Joins				
<table border="1"><tr><td>..</td></tr></table>	..	<table border="1"><tr><td>..</td></tr></table>	..		
..					
..					

# VM in action

```
t0
00:  STHR 1 t1
04:  EMT 1
06:  SWC 1
08:  STHR 1 NR1
12:  SWC 1
14:  END
NR1
15:  SWCU
t1
16:  TWB 2 1 case_1
19:  JMP done
case_1
done
22:  SWC 0
```

pc = 16  
last\_thread = 0

Threads	Signals				
<table border="1"><tr><td>8</td></tr><tr><td>16</td></tr></table>	8	16	<table border="1"><tr><td>0</td></tr><tr><td>1</td></tr></table>	0	1
8					
16					
0					
1					
States	Joins				
<table border="1"><tr><td>..</td></tr></table>	..	<table border="1"><tr><td>..</td></tr></table>	..		
..					
..					



# VM in action

```
t0
00:  STHR 1 t1
04:  EMT 1
06:  SWC 1
08:  STHR 1 NR1
12:  SWC 1
14:  END
NR1
15:  SWCU
t1
16:  TWB 2 1 case_1
19:  JMP done
case_1
done
22:  SWC 0
```

pc = 19  
last\_thread = 0

Threads	Signals
8	0
16	1

States	Joins
..	..

# VM in action

```
t0
00:  STHR 1 t1
04:  EMT 1
06:  SWC 1
08:  STHR 1 NR1
12:  SWC 1
14:  END
NR1
15:  SWCU
t1
16:  TWB 2 1 case_1
19:  JMP done
case_1
done
22:  SWC 0
```

pc = 22  
last\_thread = 0

Threads	Signals
8	0
16	1

States	Joins
..	..

# VM in action

```
t0
00:  STHR 1 t1
04:  EMT 1
06:  SWC 1
08:  STHR 1 NR1
12:  SWC 1
14:  END
NR1
15:  SWCU
t1
16:  TWB 2 1 case_1
19:  JMP done
case_1
done
22:  SWC 0
```

pc = 8  
last\_thread = 1

Threads	Signals				
<table border="1"><tr><td>8</td></tr><tr><td>24</td></tr></table>	8	24	<table border="1"><tr><td>0</td></tr><tr><td>1</td></tr></table>	0	1
8					
24					
0					
1					
States	Joins				
<table border="1"><tr><td>..</td></tr></table>	..	<table border="1"><tr><td>..</td></tr></table>	..		
..					
..					

# VM in action

```
t0
00:  STHR 1 t1
04:  EMT 1
06:  SWC 1
08:  STHR 1 NR1
12:  SWC 1
14:  END
NR1
15:  SWCU
t1
16:  TWB 2 1 case_1
19:  JMP done
case_1
done
22:  SWC 0
```

pc = 12  
last\_thread = 1

Threads	Signals				
<table border="1"><tr><td>8</td></tr><tr><td>15</td></tr></table>	8	15	<table border="1"><tr><td>0</td></tr><tr><td>1</td></tr></table>	0	1
8					
15					
0					
1					
States	Joins				
<table border="1"><tr><td>..</td></tr></table>	..	<table border="1"><tr><td>..</td></tr></table>	..		
..					
..					

# VM in action

```
t0
00:  STHR 1 t1
04:  EMT 1
06:  SWC 1
08:  STHR 1 NR1
12:  SWC 1
14:  END
NR1
15:  SWCU
t1
16:  TWB 2 1 case_1
19:  JMP done
case_1
done
22:  SWC 0
```

pc = 15  
last\_thread = 0

Threads	Signals
14	0
15	1

States	Joins
..	..

# VM in action

```
t0
00:  STHR 1 t1
04:  EMT 1
06:  SWC 1
08:  STHR 1 NR1
12:  SWC 1
14:  END
NR1
15:  SWCU
t1
16:  TWB 2 1 case_1
19:  JMP done
case_1
done
22:  SWC 0
```

pc = 14  
last\_thread = 1

Threads	Signals
14	0
15	1

States	Joins
..	..

# VM in action

```
t0
00:  STHR 1 t1
04:  EMT 1
06:  SWC 1
08:  STHR 1 NR1
12:  SWC 1
14:  END
NR1
15:  SWCU
t1
16:  TWB 2 1 case_1
19:  JMP done
case_1
done
22:  SWC 0
```

pc = 15  
last\_thread = 0

Threads	Signals
15	0
15	1

States      Joins

..	..
----	----

## The engineering details

- ▶ brickOS 2.6.10 on Redhat Linux
- ▶ gcc cross compiler 4.0.2. for H8300
- ▶ Download lx files to the lego RCX via USB IR tower



## Code Sizes

<b>Example</b>	<b>BAL</b>	<b>x86</b>		<b>H8</b>	
dacexample	369	917	60%	842	57%
abcd	870	2988	71%	2648	68%
greycounter	1289	3571	64%	2836	55%
tcint	5667	11486	51%	10074	51%
atds-100	10481	38165	73%	26334	60%

BAL: the size of our bytecode (in bytes)

x86: the size of optimized C code for an x86

H8: the size of optimized C code for an Hitachi H8

Percentages represent the size savings of using bytecode.

## Execution Times

<b>Example</b>	<b>x86</b>	<b>BAL</b>	
dacexample	0.06 $\mu$ s	1.1 $\mu$ s	18×
tcint	0.28 $\mu$ s	1.1 $\mu$ s	4×
atds-100	0.20 $\mu$ s	1.4 $\mu$ s	7×

# Conclusions

- ▶ Simple Virtual Machine
- ▶ Compilation scheme statically schedules the concurrent behavior and generates straight-line code for each thread
- ▶ VM supports context-switching well
- ▶ Bytecode for our virtual machine is roughly half the size of optimized native assembly code generated from C
- ▶ Speed tradeoff not that bad! Between 4 and 7 times slower than optimized C code