

1 Introduction

This is a specification of a simple scheduler and assembler. The system contains a set of registers and a block of memory. Processes can be created, with each containing a sequence on instructions that are executed on the system. The instruction format is a simplified format of the Intel x86 architecture. Processes are scheduled based on the credit system that is found in the Linux 2.0 kernel.

2 Stack

This specification was written as a test spec for the CZT project. As a result, there are parts that may appear to be specified in a strange way - this is to test out the tools on a large set of Z.

section Stack parents standard_toolkit

A generic stack.

$Stack[X] == [stack : seq X]$
 $InitStack[X] == [Stack[X] \mid stack = \emptyset]$

$PushStack[X]$ $\Delta Stack[X]$ $x? : X$
$stack' = stack \frown \langle x? \rangle$

$PopStack[X]$ $\Delta Stack[X]$ $x! : X$
$stack' \frown \langle x! \rangle = stack$

Ok, lets see the value of 3 unboxed items in Section 2!

3 Definitions

section Definitions parents standard_toolkit

Declarations	This Section	Globally
Unboxed items	3	3
Axiomatic definitions	0	0
Generic axiomatic defs.	0	0
Schemas	0	0
Generic schemas	2	2
Total	5	5

Table 1: Summary of Z declarations for Section 2.

Firstly, we define some basic types and functions that are used throughout the specification.

singleton is the set of all sets whose size is less than or equal to 1. This is included only to have a generic axiom definition.

relation(*singleton* _)

[X]
<i>singleton</i> _ : $\mathbb{P}(\mathbb{P} X)$
$\forall s : \mathbb{P} X \bullet \textit{singleton } s \Leftrightarrow \# s \leq 1$

The basic type of this system is a word, which specifically, is an unsigned octet. An unsigned word is used so references to memory etc a 1-relative.

WORD == 0 .. 255

Then, we define the size of the memory block, and give it a value for animation purposes.

<i>mem_size</i> : WORD
<i>mem_size</i> = 100

A *LABEL* is used to label instructions for *jump* instructions etc, although 'jump' hasn't been specified yet.

[*LABEL*]

Now we define the different instructions, as well as their operands. A *CONSTANT* is used both as a constant value, as well as a memory reference for load and store instructions.

$INST_NAME ::=$
 $add \mid sub \mid divide \mid mult \mid push \mid pop \mid load \mid store \mid loadConst \mid print$
 $OPERAND ::= AX \mid BX \mid CX \mid DX \mid constant\langle\langle WORD \rangle\rangle$
 $REGISTER == \{AX, BX, CX, DX\}$
 $CONSTANT == OPERAND \setminus REGISTER$

An instruction is specified as a instruction name, a sequence of operands, and optionally, a label.

Instruction

$label : \mathbb{P} LABEL$
 $name : INST_NAME$
 $params : seq OPERAND$
singleton label

Declarations	This Section	Globally
Unboxed items	9	12
Axiomatic definitions	1	1
Generic axiomatic defs.	1	1
Schemas	1	1
Generic schemas	0	2
Total	12	17

Table 2: Summary of Z declarations for Section 3.

4 System

section System parents Definitions, Stack

The system consists of a set of registers, and a block of memory. There is also a buffer for displaying output.

$REGISTERS == REGISTER \rightarrow OPERAND$
 $MEMORY == 1 .. mem_size \leftrightarrow WORD$

System

registers : *REGISTERS*
memory : *MEMORY*
output : seq *WORD*

Initially, all registers and memory hold the minimum *WORD* value. The output buffer is empty.

InitSystem

System

registers = {*r* : *REGISTER* • *r* ↦ *constant*(*min*(*WORD*))}
memory = {*m* : 1 .. *mem_size* • *m* ↦ *min*(*WORD*)}
output = ⟨⟩

The system can have arithmetic and memory instructions.

Arith_Inst == [*Instruction* | # *params* = 2 ∧ *params*(1) ∈ *REGISTER*]
Add_Inst == [*Arith_Inst* | *name* = *add*]
Sub_Inst == [*Arith_Inst* | *name* = *sub*]
Mult_Inst == [*Arith_Inst* | *name* = *mult*]
Div_Inst == [*Arith_Inst* | *name* = *divide*]

Memory_Inst == [*Instruction* | # *params* = 2 ∧ *params*(1) ∈ *REGISTER*
∧ *params*(2) ∈ *CONSTANT*]
Load_Inst == [*Memory_Inst* | *name* = *load*]
LoadConst_Inst == [*Memory_Inst* | *name* = *loadConst*]
Store_Inst == [*Memory_Inst* | *name* = *store*]

A print instruction prints the value of a register.

Print_Inst == [*Instruction* | # *params* = 1]

val maps constants to their value, and *dereference* dereferences the value of a register, transitively if required.

val : *CONSTANT* → *WORD*
dereference : *OPERAND* × *REGISTERS* → *WORD*
∀ *c* : *CONSTANT* •
 (∃ *n* : *WORD* • *c* = *constant*(*n*) ∧ *val*(*c*) = *n*)
∀ *a* : *OPERAND*; *r* : *REGISTERS* •
 dereference(*a*, *r*) =
 if *a* ∈ *REGISTER* **then** *dereference*(*r*(*a*), *r*) **else** *val*(*a*)

The specification of the arithmetic instructions.

Add

Δ System
Add_Inst

$\exists o_1 == \text{dereference}(\text{params}(1), \text{registers});$
 $o_2 == \text{dereference}(\text{params}(2), \text{registers}) \bullet$
 $\text{registers}' = \text{registers} \oplus \{\text{params}(1) \mapsto \text{constant}(o_1 + o_2)\}$
 $\text{memory}' = \text{memory}$
 $\text{output}' = \text{output}$

Sub

Δ System
Sub_Inst

$\exists o_1 == \text{dereference}(\text{params}(1), \text{registers});$
 $o_2 == \text{dereference}(\text{params}(2), \text{registers}) \bullet$
 $\text{registers}' = \text{registers} \oplus \{\text{params}(1) \mapsto \text{constant}(o_1 - o_2)\}$
 $\text{memory}' = \text{memory}$
 $\text{output}' = \text{output}$

Mult

Δ System
Mult_Inst

$\exists o_1 == \text{dereference}(\text{params}(1), \text{registers});$
 $o_2 == \text{dereference}(\text{params}(2), \text{registers}) \bullet$
 $\text{registers}' = \text{registers} \oplus \{\text{params}(1) \mapsto \text{constant}(o_1 * o_2)\}$
 $\text{memory}' = \text{memory}$
 $\text{output}' = \text{output}$

Div

Δ System
Div_Inst

$\exists o_1 == \text{dereference}(\text{params}(1), \text{registers});$
 $o_2 == \text{dereference}(\text{params}(2), \text{registers}) \bullet$
 $\text{registers}' = \text{registers} \oplus \{\text{params}(1) \mapsto \text{constant}(o_1 \text{ div } o_2)\}$
 $\text{memory}' = \text{memory}$
 $\text{output}' = \text{output}$

The `load` operation loads a constant from memory. The second parameter is an index to the memory location from which the constant is loaded.

<i>Load</i>
$\Delta System$
<i>Load_Inst</i>
$\exists o_2 == \text{val}(\text{params}(2)) \bullet$ $\text{registers}' = \text{registers} \oplus$ $\{\text{params}(1) \mapsto \text{constant}(\text{memory}(o_2))\}$ $\text{memory}' = \text{memory}$ $\text{output}' = \text{output}$

`loadConst` loads a constant into a register. The second parameter the constant to be loaded.

<i>Load_Const</i>
$\Delta System$
<i>LoadConst_Inst</i>
$\exists o_2 == \text{val}(\text{params}(2)) \bullet$ $\text{registers}' = \text{registers} \oplus \{\text{params}(1) \mapsto \text{constant}(o_2)\}$ $\text{memory}' = \text{memory}$ $\text{output}' = \text{output}$

Store the value of a register in memory.

<i>Store</i>
$\Delta System$
<i>Store_Inst</i>
$\exists o_1 == \text{dereference}(\text{params}(1), \text{registers});$ $o_2 == \text{val}(\text{params}(2)) \bullet$ $\text{memory}' = \text{memory} \oplus \{o_2 \mapsto o_1\}$ $\text{registers}' = \text{registers}$ $\text{output}' = \text{output}$

<i>Print</i>
$\Xi System$
<i>Print_Inst</i>
$\text{output}' = \text{output} \hat{\ } \langle \text{dereference}(\text{params}(1), \text{registers}) \rangle$ $\text{registers}' = \text{registers}$ $\text{memory}' = \text{memory}$

$$\begin{aligned}
Stack_Inst &== [Instruction \mid \# \text{ params} = 1] \\
Push_Inst &== [Stack_Inst \mid \text{name} = \text{push}] \\
Pop_Inst &== [Stack_Inst \mid \text{name} = \text{pop}]
\end{aligned}$$

The specification of the stack instructions on the system.

$ \begin{aligned} &Push0 \\ &\exists System \\ &PushStack[WORD] \\ &Push_Inst \end{aligned} $
$x? = \text{dereference}(\text{params}(1), \text{registers})$

$ \begin{aligned} &Pop0 \\ &\Delta System \\ &PopStack[WORD] \\ &Pop_Inst \end{aligned} $
$ \begin{aligned} \text{registers}' &= \text{registers} \oplus \{\text{params}(1) \mapsto \text{constant}(x!)\} \\ \text{memory}' &= \text{memory} \\ \text{output}' &= \text{output} \end{aligned} $

$$\begin{aligned}
Push &== Push0 \upharpoonright [System; Stack[WORD]] \\
Pop &== Pop0 \upharpoonright [System; Stack[WORD]]
\end{aligned}$$

This executes an instruction on the on the system. $inst?$ is the instruction to execute, and $base?$ is the base memory value of the executing process. If the instruction is a **load** or **store** instruction, the memory reference must offset using the base value.

$exec_inst$ $\Delta System$ $inst? : Instruction$ $base? : 1 .. mem_size$
$\exists label : \mathbb{P} LABEL; name : INST_NAME; params : seq OPERAND \mid$ $label = inst?.label \wedge name = inst?.name \wedge$ $params = inst?.params \bullet$ $Add \vee Sub \vee Mult \vee Div \vee$ $Print \vee Load_Const \vee$ $name \in \{load, store\} \Rightarrow (\exists p : seq OPERAND \mid$ $p = \langle params(1),$ $constant(val(params(2)) + base?) \rangle \bullet$ $Load[p/params] \vee Store[p/params])$

Declarations	This Section	Globally
Unboxed items	19	31
Axiomatic definitions	1	2
Generic axiomatic defs.	0	1
Schemas	13	14
Generic schemas	0	2
Total	33	50

Table 3: Summary of Z declarations for Section 4.

5 Scheduler

section *Scheduler* parents *System*

This part of the specification is the scheduler.

Here, we declare the set of process IDs, the priority values, and the default number of credits a process receives when it is created.

$Pid == \mathbb{N}$
 $Priority == -19 .. 19$
 $Default_Credits == 10$

The possible status that a process can hold.

$Status ::= pWaiting \mid pReady \mid pRunning$

A process consists of a process ID, a status, a number of credits, and a priority. Each process has a sequence of instructions to be executed on the assembler, with a pointer to the current instruction. The memory that a process can occupy is between a base and limit value. Instructions must only access memory with a value less than the limit, but they know nothing about the base value - this is added onto the memory index provided by the instruction when an instruction is executed. Each process also contains a stack and values for all registers, which are used to store values when the process is suspended.

Processes

$$\begin{array}{l}
pids : \mathbb{P} Pid \\
status : Pid \mapsto Status \\
credits : Pid \mapsto \mathbb{N} \\
priority : Pid \mapsto Priority \\
instructions : Pid \mapsto (\text{seq } Instruction) \\
inst_pointer : Pid \mapsto \mathbb{N}_1 \\
base, limit : Pid \mapsto WORD \\
pregisters : Pid \mapsto REGISTERS \\
pstack : Pid \mapsto Stack[WORD] \\
\hline
pids = \text{dom}(status) = \text{dom}(credits) = \text{dom}(priority) = \\
\text{dom}(instructions) = \text{dom}(inst_pointer) = \text{dom}(base) = \\
\text{dom}(limit) = \text{dom}(pstack) \\
\forall pid : pids \bullet inst_pointer(pid) \leq \#(instructions(pid)) \\
\forall pid : pids \bullet base(pid) + limit(pid) \leq mem_size
\end{array}$$

The *sort* function takes the credits and priorities of all processes, and returns a sequence of process IDS sorted firstly by their credits (the more credits a process has, the higher preference they get), and if the credits are equal, then their priority. If the priority is equal, then the order is non-deterministic.

$$\begin{array}{l}
sort : (Pid \mapsto \mathbb{N}) \times (Pid \mapsto Priority) \mapsto \text{iseq } Pid \\
\hline
sort = (\lambda credits : (Pid \mapsto \mathbb{N}); priority : (Pid \mapsto Priority) \mid \\
\text{dom}(credits) = \text{dom}(priority) \bullet \\
(\mu s : \text{iseq } Pid \mid \text{ran}(s) = \text{dom}(credits) \wedge \\
(\forall i : 1 \dots \#s - 1 \bullet \\
credits(s(i)) > credits(s(i+1)) \vee \\
(credits(s(i)) = credits(s(i+1)) \wedge \\
priority(s(i)) > priority(s(i+1)))) \bullet s))
\end{array}$$

To interrupt a process during execution, the kernel must be in *kernel* mode.

Mode ::= user | kernel

For the scheduler, we track which mode the operating system is in, as well as declaring three “secondary” variables, *waiting*, *running*, and *ready*, to keep the sets of waiting running, and ready variables respectively. In fact, *ready* is a sequence, and is ordered based on the credits that each process has. A process with more credits will have a higher priority. This is fair scheduling, because at each timer interrupt (the *tick* operation specified below), the current process loses one credit, therefore, process spending a lot of time executing will eventually have a low priority.

```

Scheduler
  Processes
  System
  Stack[WORD]
  mode : Mode
  waiting, running :  $\mathbb{P}$  Pid
  ready : iseq Pid

  # running  $\leq$  1
  waiting  $\cap$  running  $\cap$  ran ready =  $\emptyset$ 
  waiting  $\cup$  running  $\cup$  ran ready = pids
  waiting = {p : pids | (status  $\sim$ )(pWaiting) = p}
  running = {p : pids | (status  $\sim$ )(pRunning) = p}
  ready = sort((waiting  $\cup$  running)  $\triangleleft$  credits,
    (waiting  $\cup$  running)  $\triangleleft$  priority)
   $\forall r$  : ran(ready)  $\bullet$  status(r) = pReady
   $\forall r$  : running  $\bullet$  credits(r) > 0

```

This uses semicolons as conjunctions for predicates, which conforms to the grammar in the ISO standard, but according to the list of differences between ZRM and ISO Z on Ian Toyn’s website, semicolons can no longer be used to conjoin predicates.

```

InitScheduler
  Scheduler
  InitStack[WORD]
  InitSystem

  pids =  $\emptyset$ ; status =  $\emptyset$ ; priority =  $\emptyset$ 
  credits =  $\emptyset$ ; instructions =  $\emptyset$ ; inst_pointer =  $\emptyset$ 
  waiting =  $\emptyset$ ; running =  $\emptyset$ ; ready =  $\langle \rangle$ 
  base = {}; limit = {}; pregisters = {}
  mode = user

```

newProcess creates a new process with a unique process ID and a specified priority, and places this new process on the ready queue.

<i>create_new_process</i>
Δ <i>Scheduler</i> \exists <i>System</i> <i>priority?</i> : <i>Priority</i> <i>instructions?</i> : seq <i>Instruction</i> <i>base?, limit?</i> : <i>WORD</i> <i>pid!</i> : <i>Pid</i>
<i>pid!</i> \notin <i>pids</i> <i>status'</i> = <i>status</i> \cup { <i>pid!</i> \mapsto <i>pReady</i> } <i>credits'</i> = <i>credits</i> \cup { <i>pid!</i> \mapsto <i>Default_Credits</i> } <i>priority'</i> = <i>priority</i> \cup { <i>pid!</i> \mapsto <i>priority?</i> } <i>instructions'</i> = <i>instructions</i> \cup { <i>pid!</i> \mapsto <i>instructions?</i> } <i>inst_pointer'</i> = <i>inst_pointer</i> \cup { <i>pid!</i> \mapsto 1} <i>base'</i> = <i>base</i> \cup { <i>pid!</i> \mapsto <i>base?</i> } <i>limit'</i> = <i>limit</i> \cup { <i>pid!</i> \mapsto <i>limit?</i> } <i>pregisters'</i> = <i>pregisters</i> \cup { <i>pid!</i> \mapsto { <i>r</i> : <i>REGISTER</i> • <i>r</i> \mapsto constant(<i>min</i> (<i>WORD</i>))}} <i>pstack'</i> = <i>pstack</i> \cup { <i>pid!</i> \mapsto (\langle <i>stack</i> \rangle)} <i>pids'</i> = <i>pids</i> \cup { <i>pid!</i> }

We define a schema that contains only the variables that do not change when a reschedule occurs.

$$\text{RescheduleChange} == \text{Scheduler} \setminus (\text{status}, \text{running}, \text{ready}, \text{waiting}, \text{credits})$$

A reschedule occurs when all ready processes have no credits. Every process, not just the ready processes, have their credits re-calculated using the formula $\text{credits} = \text{credits}/2 + \text{priority}$. This implies that the ready process with the highest priority will be the next process executed.

<i>reschedule</i>
Δ <i>Scheduler</i> \exists <i>RescheduleChange</i>
<i>ready</i> $\neq \emptyset$ $\forall r : \text{ran}(\text{ready}) \bullet \text{credits}(r) = 0 \Rightarrow$ <i>credits'</i> = { <i>p</i> : <i>pids</i> • <i>p</i> \mapsto (<i>credits</i> (<i>p</i>) div 2) + <i>priority</i> (<i>p</i>)} \wedge <i>status'</i> = <i>status</i> $\neg (\forall r : \text{ran}(\text{ready}) \bullet \text{credits}(r) = 0) \Rightarrow$ <i>status'</i> = <i>status</i> \oplus { <i>head</i> (<i>ready</i>) \mapsto <i>pRunning</i> } \wedge <i>credits'</i> = <i>credits</i>

We declare a new schema that contains only the state variables that do not change when a status change occurs.

$$\text{StatusChange} == \text{Scheduler} \setminus (\text{status}, \text{running}, \text{waiting}, \text{ready}, \text{registers}, \text{pregisters}, \text{pstack})$$

Interrupts the currently executing process if the new process is of a higher priority than the current process and the kernel is in *kernel* mode.

interrupt
$\Delta \text{Scheduler}$
$\Xi \text{StatusChange}$
$\text{create_new_process}$
$\text{mode} = \text{kernel}$
$\text{running} = \emptyset \vee (\exists p : \text{running} \bullet \text{priority?} \geq \text{priority}(p))$
$\exists r : \text{running} \bullet$
$\text{status}' = \text{status} \oplus \{pid! \mapsto p\text{Running}, r \mapsto p\text{Ready}\} \wedge$
$\text{pregisters}' = \text{pregisters} \oplus \{r \mapsto \text{registers}\} \wedge$
$\theta \text{Stack}' = \text{pstack}(r)$
$\text{registers}' = \text{pregisters}(pid!)$

Remove the currently running process and put it back in the ready queue.

$\text{remove_running_process}$
$\Delta \text{Scheduler}$
$\Xi \text{StatusChange}$
$\exists pid == (\mu r : \text{running}) \bullet$
$\text{status}' = \text{status} \oplus \{pid \mapsto p\text{Ready}\} \wedge$
$\text{pregisters}' = \text{pregisters} \oplus \{pid \mapsto \text{registers}\} \wedge$
$\text{pstack}' = \text{pstack} \oplus \{pid \mapsto \theta \text{Stack}'\}$

A process becomes blocked if it is waiting on a resource such as an IO device, or waiting on another process

$$\text{block_process} == \text{remove_running_process} \circledast \text{reschedule}$$

We declare a schema containing only the variables that change for an unblock.

$$\text{UnblockProcessChange} == \text{Scheduler} \setminus (\text{status}, \text{running}, \text{ready}, \text{waiting})$$

Unblocks a process that is blocked by another process.

unblock_process
$\Delta \text{Scheduler}$
$\exists \text{UnblockProcessChange}$
$pid? : \text{Pid}$
<hr/>
$pid? \in \text{pids}$
$\text{status}(pid?) = p \text{Waiting}$
$\text{running} = \emptyset \Leftrightarrow \text{status}' = \text{status} \oplus \{pid? \mapsto p \text{Running}\}$
$\text{running} \neq \emptyset \Leftrightarrow \text{status}' = \text{status} \oplus \{pid? \mapsto p \text{Ready}\}$

Remove a process from the system

remove_process
$\Delta \text{Scheduler}$
$\exists \text{Stack}[\text{WORD}]$
$\exists \text{System}$
$pid? : \text{Pid}$
<hr/>
$pid? \in \text{pids}$
$\text{pids}' = \text{pids} \setminus \{pid?\}$
$\text{status}' = \{pid?\} \triangleleft \text{status}$
$\text{credits}' = \{pid?\} \triangleleft \text{credits}$
$\text{priority}' = \{pid?\} \triangleleft \text{priority}$
$\text{instructions}' = \{pid?\} \triangleleft \text{instructions}$
$\text{inst_pointer}' = \{pid?\} \triangleleft \text{inst_pointer}$
$\text{base}' = \{pid?\} \triangleleft \text{base}$
$\text{limit}' = \{pid?\} \triangleleft \text{limit}$
$\text{pregisters}' = \{pid?\} \triangleleft \text{pregisters}$
$\text{pstack}' = \{pid?\} \triangleleft \text{pstack}$

Update the details in the process table when each instruction is executed, as well as communicate the current instruction and the base value for the current process.

$$\text{ChangeInstPointer} == \text{Scheduler} \setminus (\text{inst_pointer})$$

update_process_table

Δ Scheduler

inst! : Instruction

base! : WORD

running $\neq \emptyset$

$(\exists pid == (\mu r : running) \bullet$

$inst! = head(instructions(pid)) \wedge$

$base! = base(pid) \wedge$

$(inst_pointer(pid) = \#(instructions(pid)) \Rightarrow$

$remove_process[pid/pid?] \wedge$

$inst_pointer(pid) < \#(instructions(pid)) \Rightarrow$

$inst_pointer' =$

$inst_pointer \oplus \{pid \mapsto inst_pointer(pid) + 1\}$)

θ ChangeInstPointer = θ ChangeInstPointer'

$next == exec_inst \gg update_process_table \circ$
 $([\Delta Scheduler \mid running = \emptyset] \wedge reschedule) \vee$
 $([\Xi Scheduler \mid running \neq \emptyset])$

$idle0 == \neg \mathbf{pre} \ next$

idle

Ξ Scheduler

inst? : Instruction

base? : WORD

idle0

$tick == next \vee idle$

$\vdash? (\forall n : \mathbb{N}_1 \bullet n > 0)$

$[X] \vdash? \forall x : \mathbb{P} X \bullet \# x \leq 1 \Leftrightarrow singleton \ x$

theorem PreconditionCheck

$\forall Scheduler \bullet \mathbf{pre} \ update_process_table$

Declarations	This Section	Globally
Unboxed items	21	52
Axiomatic definitions	1	3
Generic axiomatic defs.	0	1
Schemas	11	25
Generic schemas	0	2
Total	33	83

Table 4: Summary of Z declarations for Section 5.