

Certifying Low-Level Programs with Hardware Interrupts and Preemptive Threads

Xinyu Feng^{†‡} Zhong Shao[†] Yuan Dong^{§†} Yu Guo^{*}

[†]Yale University

[‡]Toyota Technological Institute at Chicago

[§]Tsinghua University

^{*}University of Science and Technology of China

Abstract

Hardware interrupts are widely used in the world’s critical software systems to support preemptive threads, device drivers, operating system kernels, and hypervisors. Handling interrupts properly is an essential component of low-level system programming. Unfortunately, interrupts are also extremely hard to reason about: they dramatically alter the program control flow and complicate the invariants in low-level concurrent code (e.g., implementation of synchronization primitives). Existing formal verification techniques—including Hoare logic, typed assembly language, concurrent separation logic, and the assume-guarantee method—have consistently ignored the issues of interrupts; this severely limits the applicability and power of today’s program verification systems.

In this paper we present a novel Hoare-logic-like framework for certifying low-level system programs involving both hardware interrupts and preemptive threads. We show that enabling and disabling interrupts can be formalized precisely using simple ownership-transfer semantics, and the same technique also extends to the concurrent setting. By carefully reasoning about the interaction among interrupt handlers, context switching, and synchronization libraries, we are able to—for the first time—successfully certify a preemptive thread implementation and a large number of common synchronization primitives. Our work provides a foundation for reasoning about interrupt-based kernel programs and makes an important advance toward building fully certified operating system kernels and hypervisors.

1. Introduction

Low-level system programs (e.g., thread implementations, device drivers, OS kernels, and hypervisors) form the backbone of almost every safety-critical software system in the world. It is thus highly desirable to formally certify the correctness of these programs. Indeed, there have been several new projects launched recently (e.g., Verisoft/XT [8, 19], L4.verified [23], Singularity [12]), all aiming to build certified OS kernels and/or hypervisors. With formal specifications and provably safe components, certified system software can provide a trustworthy computing platform and enable anticipatory statement about system configuration and behavior [12].

Unfortunately, system programs—especially those involving both interrupts and concurrency—are extremely hard to reason about. In Fig 1, we divide programs in a typical OS kernel into two layers. At the “higher” abstraction level, we have threads that follow the standard concurrent programming model [10]: interrupts are invisible, but the execution of a thread can be preempted by other threads; synchronization operations are treated as primitives.

Below this layer (see the shaded box), we have more subtle “lower-level” code involving both interrupts and concurrency. The implementation of many synchronization primitives and input/output operations requires explicit manipulation of interrupts; they behave concurrently in a preemptive way (if interrupt is enabled) or a non-preemptive way (if interrupt is disabled). When

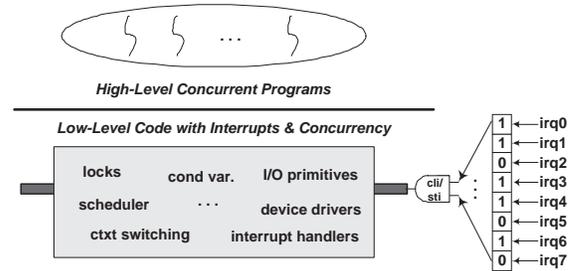


Figure 1. “High-Level” vs. “Low-Level” System Programs

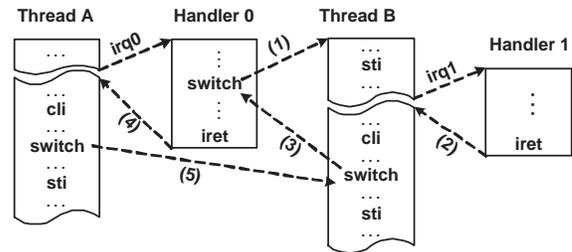


Figure 2. Interaction between Threads and Interrupts

execution of a thread is interrupted, control is transferred to an interrupt handler, which may call the thread scheduler and switch the control to another thread. Some of the code in the shaded box (e.g., the scheduler and context switching routine) may behave sequentially since they are always executed with interrupt disabled.

Existing program verification techniques (including Hoare logic, typed assembly language [15], concurrent separation logic [17, 3], and its assume-guarantee variant [24]) can probably handle those high-level concurrent programs, but they have consistently ignored the issues of interrupts thus cannot be used to certify concurrent code in the shaded box. Having both explicit interrupts and threads creates the following new challenges:

- *Asymmetric preemption relations.* Non-handler code may be preempted by an interrupt handler (and low-priority handlers can be preempted by higher-priority ones), but not vice versa. Interrupt handlers cannot be simply treated as threads [20].
- *Subtle intertwining between interrupts and threads.* In Fig 2, thread A is interrupted by irq0 (say, the timer). In the handler, the control is switched to thread B. From thread A’s point of view, the behavior of the handler 0 is complex: should the handler be responsible for the behavior of thread B?
- *Asymmetric synchronizations.* Synchronization between handler and non-handler code is achieved simply by enabling and disabling interrupts (via sti and cli instructions in x86). Unlike

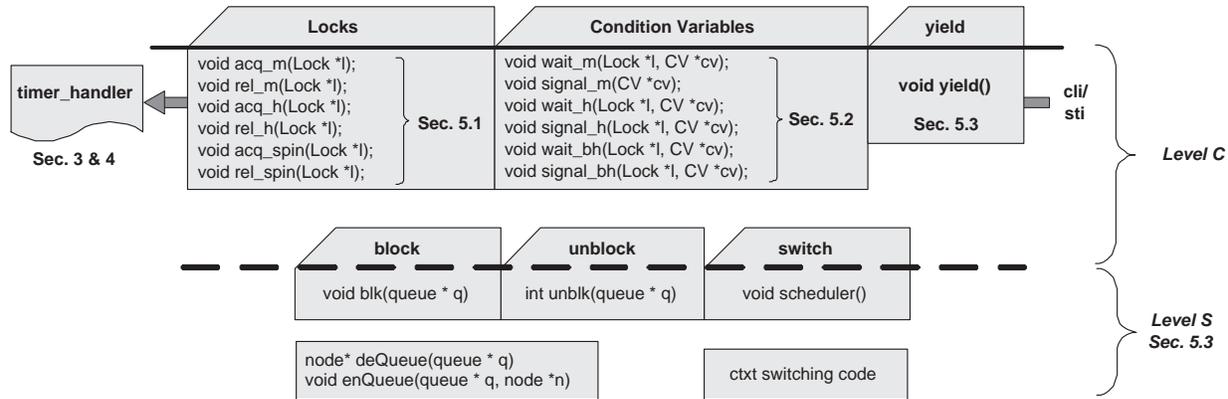


Figure 3. Structure of Our Certified Preemptive Thread Implementation

locks, interrupts can be disabled by one thread and enabled by another. In Fig 2, thread A disables interrupts and then switches control to thread B (step (5)), which will enable interrupts.

- Handler for higher-priority interrupts might be “interrupted” by lower-priority ones. In Fig 2, handler 0 switches the control to thread B at step (1); thread B enables interrupts and is interrupted by irq1, which may have a lower-priority than irq0.

In this paper we tackle these challenges directly and present a novel framework for certifying low-level programs involving both interrupts and preemptive threads. We introduce a new abstract interrupt machine (named AIM, see Sec 3 and the upper half of Fig 3) to capture “interrupt-aware” concurrency, and use simple ownership-transfer semantics to reason about the interaction among interrupt handlers, context switching, and synchronization libraries. Our paper makes the following new contributions:

- As far as we know, our work presents the first program logic (see Sec 4) that can successfully certify the correctness of low-level programs involving both interrupts and concurrency. Our idea of using ownership-transfer semantics to model interrupts is both novel and general (since it also works in the concurrent setting). Our logic supports modular verification: threads and handlers can be certified in the same way as we certify sequential code without worrying about possible interleaving. Soundness of our logic is formally proved in the Coq proof assistant.
- Following separation logic’s local-reasoning idea, our program logic also enforces partitions of resources between different threads and between threads and interrupt handlers. These logical partitions at different program points essentially give an abstract formalization of the semantics of interrupts and the interaction between handlers and threads.
- Our AIM machine (see Sec 3) unifies both the preemptive and non-preemptive threading models, and to our best knowledge, is the first to successfully formalize concurrency with explicit interrupt handlers. In AIM, operations that manipulates thread queues are treated as primitives; These operations, together with the scheduler and context-switching code (the low half of Fig 3), are strictly sequential thus can be certified in a simpler logic. Certified code at different levels are linked together using an OCAP-style framework [5].
- Synchronization operations can be implemented as subroutines in AIM. To demonstrate the power of our framework, we have certified, for the first time, various implementations of locks and condition variables (see Sec 5). Our specifications pinpoint precisely the differences between different implementations.

2. Informal Development

Before presenting our formal framework, we first informally explain the key ideas underlying our abstract machine and our ownership-transfer semantics for reasoning about interrupts.

2.1 Design of the Abstract Machine

In Figure 3 we outline the structure of a thread implementation taken from a simplified OS kernel. We split all “shaded” code into two layers: the upper level C (for “Concurrent”) and the low level S (for “Sequential”). Code at Level C is concurrent; it handles interrupt explicitly and implements the interrupt handler but abstracts away the implementation of threads. Code at Level S is sequential (always executed with interrupt disabled); functions that require to know the concrete representations of thread control blocks (TCBs) and thread queues are implemented at Level S; there is one queue for ready threads and multiple queues for blocked threads.

We implement three primitive thread operations at Level S: switch, block, and unblock. The switch primitive, shown as the scheduler() function in Fig 3, saves the execution context of the current thread into the ready queue, picks another one from the queue, and switches to the execution context of the new thread. The block primitive takes a pointer to a block queue as argument, puts the current thread into the block queue, and switches the control to a thread in the ready queue. The unblock primitive also takes a pointer to a block queue as argument; it moves a thread from the block queue to the ready queue but does not do context switching. Level S also contains code for queue operations and thread context switching, which are called by these thread primitives.

In the abstract machine at Level C, we use instructions cli/sti to enable/disable interrupts (same as on x86 processors); the primitives switch, block and unblock are also treated as instructions; thread queues are now abstract algebraic structures outside of data heap and can only be accessed via the thread primitives.

Although the abstract machine is based on our implementation of the preemptive thread library, the idea of interfacing using Levels C and S is very general, given that the thread primitives we use are very common in most implementations of thread libraries. Note that we are not trying to claim this is the only way or the best way to divide code into separate abstraction levels, but this design does give us a nice abstraction at Level C so that we can focus on the interaction between threads and interrupts.

2.2 Ownership-Transfer Semantics

Concurrent entities, *i.e.*, the handler code and the threads consisting of the non-handler code, all need to access memory. To guarantee the non-interference, we enforce the following invariant, inspired by recent work on Concurrent Separation Logic [3, 17]: *there al-*

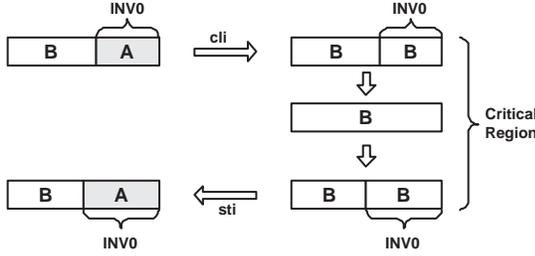


Figure 4. Memory Partition for Handler and Non-Handler

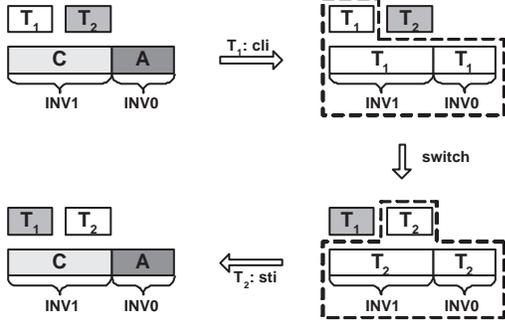


Figure 5. The Memory Model for Multi-Threaded Non-Handler

ways exists a partition of memory among these concurrent entities, and each entity can only access its own part of memory. There are two important points about this invariant:

- the partition is *logical*; we do not need to change our model of the physical machine, which only has one global shared data heap. The logical partition can be enforced following Separation Logic [13, 21], as we will explain below.
- the partition is not static; it can be dynamically adjusted during program execution, which is done by transferring the ownership of memory from one entity to the other.

Instead of using the operational semantics of `cli`, `sti` and thread primitives described above to reason about programs, we model their semantics in terms of memory ownership transfer. This semantics completely hides thread queues and thus the complex interleaving between the non-handler threads and the handler code.

We first study the semantics of `cli` and `sti` assuming that the non-handler code is always sequential. Since the interrupt handler can preempt the non-handler code but not vice versa, we reserve the part of memory used by the handler from the global memory, shown as block A in Fig 4. Block A needs to be well-formed with respect to the precondition of the handler, which ensures safe execution of the handler code. We call the precondition an invariant $INV0$, since the interrupt may come at any program point (as long as it is enabled) and this precondition needs to always hold. If the interrupt is enabled, the non-handler code can only access the rest part of memory, called block B. If it needs to access block A, it has to first disable the interrupt by `cli`. Therefore we can model the semantics of `cli` as a transfer of ownership of the well-formed block A, as shown in Fig 4. The non-handler code does not need to preserve the invariant $INV0$ if the interrupt is disabled, but it needs to ensure $INV0$ holds before it enables the interrupt again using `sti`. The `sti` instruction returns the well-formed block A to the interrupt handler.

If the non-handler code is multi-threaded, we also need to guarantee non-interference between these threads. Fig 5 refines the memory model. The block A is still dedicated to the interrupt handler. The memory block B is split into three parts (assuming there

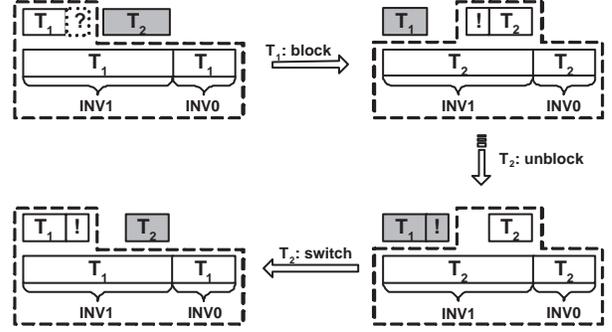


Figure 6. Block and Unblock

are only two threads): each thread has its own private memory, and both threads share the block C. When block C is available for share, it needs to be well-formed with some specification $INV1$. However, a thread cannot directly access block C if the interrupt is enabled, even if the handler does not access block C. That is because the handler may switch to another thread, as shown in Fig 2 (step (1)). To access block A and C, the current thread, say T_1 , needs to disable the interrupt; so `cli` grants T_1 the ownership of well-formed blocks A and C. If T_1 wants to switch control to T_2 , it first makes sure that $INV0$ and $INV1$ hold over A and C respectively. When T_2 takes control, it can either access A and C, or enable the interrupt and release their ownership (knowing they are well-formed).

Blocking thread queues are used to implement synchronization primitives, such as locks or condition variables. When the lock is not available, or the condition associated with the condition variable does not hold, the current thread is put into the corresponding block queue. We can also model the semantics of block and unblock as resource ownership transfer: a blocked thread is essentially waiting for the availability of some resource, *e.g.*, the lock and the resource protected by the lock, or the resource over which the condition associated with the condition variable holds. As shown in Fig 6, thread T_1 executes `block` when it waits for some resource (represented as the dashed box containing “?”). Since block switches control to other threads, T_1 needs to ensure that $INV0$ and $INV1$ hold over A and C, which is the same requirement as `switch`. When T_2 makes the resource available, it executes `unblock` to release a thread in the corresponding block queue, and transfers the ownership of the resource to the released thread. Note that `unblock` itself does not do context switching. When T_1 takes control again, it’ll own the resource. From T_1 ’s point of view, the `block` operation acquires the resource associated with the corresponding block queue. This view of `block` and `unblock` is very flexible: by choosing whether the resource is empty or not, we can certify implementations of Mesa- and Hoare-style condition variables (see Sec 5).

3. The Abstract Interrupt Machine (AIM)

In this section, we present our Abstract Interrupt Machine (AIM) in two steps. AIM-1 shows the interaction between the handler and sequential non-handler code. AIM-2, the final definition of AIM, extends AIM-1 with multi-threaded non-handler code.

3.1 AIM-1

AIM-1 is defined in Figure 7. The whole machine configuration \mathbb{W} consists of a code heap \mathbb{C} , a mutable program state \mathbb{S} , a control stack \mathbb{K} , and a program counter pc . The code heap \mathbb{C} is a finite partial mapping from code labels to commands c . Each command c is either a sequential or branch instruction ι , or jump or return instructions. The state \mathbb{S} contains the data heap \mathbb{H} , the register file \mathbb{R} , and flags ie and is . Data heap is modeled as a *finite partial* mapping from labels to integers. The register file is a total function

(World)	\mathbb{W}	::= (C, S, K, pc)
(CodeHeap)	\mathbb{C}	::= {f \rightsquigarrow c} [*]
(State)	\mathbb{S}	::= (H, R, ie, is)
(Heap)	\mathbb{H}	::= {1 \rightsquigarrow w} [*]
(RegFile)	\mathbb{R}	::= {r ₀ \rightsquigarrow w ₀ , ..., r _k \rightsquigarrow w _k }
(Stack)	\mathbb{K}	::= nil f :: K (f, R) :: K
(Bit)	\mathbf{b}	::= 0 1
(Flags)	ie, is	::= b
(Labels)	l, f, pc	::= n (nat nums)
(Word)	w	::= i (integers)
(Register)	r	::= r ₀ r ₁ ...
(Instr)	\mathbf{i}	::= mov r _d , r _s movi r _d , w add r _d , r _s sub r _d , r _s ld r _d , w(r _s) st w(r _t), r _s beq r _s , r _t , f call f cli sti
(Commd)	c	::= \mathbf{i} j f jr r _s ret ired
(InstrSeq)	\mathbb{I}	::= \mathbf{i} ; \mathbb{I} j f jr r _s ret ired

Figure 7. Definition of AIM-1

$$\mathbb{C}[\mathbf{f}] \triangleq \begin{cases} \mathbf{c} & \mathbf{c} = \mathbb{C}(\mathbf{f}) \text{ and } \mathbf{c} = \mathbf{j} \mathbf{f}, \mathbf{j}r \ r_s, \text{ ret, or } \text{ired} \\ \mathbf{i}; \mathbb{I} & \mathbf{i} = \mathbb{C}(\mathbf{f}) \text{ and } \mathbb{I} = \mathbb{C}[\mathbf{f} + 1] \end{cases}$$

$$(F\{a \rightsquigarrow b\})(x) \triangleq \begin{cases} b & \text{if } x = a \\ F(x) & \text{otherwise.} \end{cases}$$

$$\mathbb{S}|_{\mathbb{H}'} \triangleq (\mathbb{H}', \mathbb{S}, \mathbb{R}, \mathbf{S}, \mathbf{ie}, \mathbf{is})$$

$$\mathbb{S}|_{\mathbb{R}'} \triangleq (\mathbb{S}, \mathbb{H}, \mathbb{R}, \mathbf{S}, \mathbf{ie}, \mathbf{is})$$

$$\mathbb{S}|_{\{\mathbf{ie}=\mathbf{b}\}} \triangleq (\mathbb{S}, \mathbb{H}, \mathbb{S}, \mathbb{R}, \mathbf{b}, \mathbf{S}, \mathbf{is})$$

$$\mathbb{S}|_{\{\mathbf{is}=\mathbf{b}\}} \triangleq (\mathbb{S}, \mathbb{H}, \mathbb{S}, \mathbb{R}, \mathbf{S}, \mathbf{ie}, \mathbf{b})$$

Figure 8. Definition of Representations

which maps register names to integers. The binary flags **ie** and **is** record whether the interrupt is disabled, and whether it is currently being serviced, respectively. The abstract control stack \mathbb{K} saves the return address of the current function or the interrupt handler. Each stack frame either contains a code label **f** or a pair (f, R). The frame (f, R) is pushed when the interrupt is processed, which will be explained below. An empty stack is represented as nil. The program counter **pc** points to the current command in \mathbb{C} . We also define the instruction sequence \mathbb{I} as a sequence of sequential instructions ending with jump or return commands. $\mathbb{C}[\mathbf{f}]$ extracts an instruction sequence starting from **f** in \mathbb{C} , as defined in Figure 8. We use dot notation to represent a component in a tuple, e.g., $\mathbb{S}.\mathbb{H}$ means the data heap in state \mathbb{S} . More representations are defined in Figure 8.

Operational Semantics At each step, the machine either executes the next instruction at **pc** or jumps to handle the incoming interrupt. To simplify the presentation, the machine supports only one interrupt, with a global interrupt handler entry **h_entry**. Support of multi-level interrupts is discussed in the Section 4.6. An incoming interrupt is processed only if the **ie** bit is set, and no interrupt is currently being serviced (i.e., **is** = 0). The processor handles the interrupt in the following steps:

- pushes the current **pc** and register file \mathbb{R} onto the stack \mathbb{K} ;
- clears the **ie** bit and sets the **is** bit;
- sets the **pc** to **h_entry**.

The state transition ($\mathbb{W} \xrightarrow{\mathbf{c}} \mathbb{W}'$) is defined in the **IRQ** rule in Fig 9.

NextS _(c,K) S S' where S = (H, R, ie, is)	
if c =	S' =
mov r _d , r _s	(H, R {r _d \rightsquigarrow R(r _s)}, ie, is)
movi r _d , w	(H, R {r _d \rightsquigarrow w}, ie, is)
add r _d , r _s	(H, R {r _d \rightsquigarrow (R(r _s) + R(r _d))}, ie, is)
sub r _d , r _s	(H, R {r _d \rightsquigarrow (R(r _s) - R(r _s))}, ie, is)
ld r _d , w(r _s)	(H, R {r _d \rightsquigarrow H(R(r _s) + w)}, ie, is) if (R(r _s) + w) \in dom(H)
st w(r _t), r _s	(H { (R(r _t) + w) \rightsquigarrow R(r _s) }, R, ie, is) if (R(r _t) + w) \in dom(H)
cli	S _{ie=0}
sti	S _{ie=1}
ired	(H, R', 1, 0) if is = 1, K = (f, R') :: K' for some f and K'
other cases	S

NextK _(pc,c) K K'	
if c =	S' =
call f	(pc + 1) :: K
ret	K'' if K = f :: K'' for some f
ired	K'' if K = (f, R) :: K'' for some f and R
other cases	K

NextPC _(c,R,K) pc pc'	
if c =	pc' =
beq r _s , r _t , f	f if R(r _s) = R(r _t)
beq r _s , r _t , f	pc + 1 if R(r _s) \neq R(r _t)
call f	f
j f	f
jr r _s	R(r _s)
ret	f if K = f :: K' for some K'
ired	f if K = (f, R') :: K' for some K' and R'
other cases	pc + 1

$$\frac{c = \mathbb{C}(\text{pc}) \quad \text{NextS}_{(c,K)} \mathbb{S} \mathbb{S}' \quad \text{NextK}_{(pc,c)} \mathbb{K} \mathbb{K}' \quad \text{NextPC}_{(c,S,R,K)} \text{pc} \text{pc}'}{(C, S, K, \text{pc}) \mapsto (C, S', K', \text{pc}')} \quad (\text{PC})$$

$$\frac{\text{ie} = 1 \quad \text{is} = 0}{(C, (H, R, \text{ie}, \text{is}), K, \text{pc}) \not\leq (C, (H, R, 0, 1), (\text{pc}, R) :: K, \text{h_entry})} \quad (\text{IRQ})$$

$$\mathbb{W} \mapsto \mathbb{W}' \triangleq (\mathbb{W} \mapsto \mathbb{W}') \vee (\mathbb{W} \not\leq \mathbb{W}')$$

Figure 9. Operational Semantics

The operational semantics of each instruction is defined in Figure 9. The relation $\text{NextS}_{(c,K)}$ shows the transition of states by executing **c** with stack \mathbb{K} ; $\text{NextK}_{(pc,c)}$ describes the change of stacks made by **c** at the program counter **pc**; while $\text{NextPC}_{(c,R,K)}$ shows how **pc** changes after **c** is executed with \mathbb{R} and \mathbb{K} . Semantics of most instructions are straightforward, except **ired** which runs at the end of each interrupt handler and does the following:

- pops the stack frame on the top of the stack \mathbb{K} ; the frame must be in the form of (f, R'), which is saved when the interrupt is handled (see the **IRQ** rule);
- restores **ie** and **is** with the value when the interrupt occurs, which must be 1 and 0 respectively (otherwise the interrupt cannot have been handled);
- resets the **pc** and the register file \mathbb{R} with **f** and \mathbb{R}' , respectively.

```

inclleft:   $\neg\{(p_0, \text{NoG})\}$       h_entry:   $\neg\{(p_i, g_i)\}$ 
movi $r1, RIGHT                    movi $r1, LEFT
movi $r2, LEFT                      movi $r2, RIGHT
l_loop:    $\neg\{(p_1, \text{NoG})\}$       movi $r3, 0
movi $r3, 0                          ld $r4, 0($r1)
cli                                       beq $r3, $r4, r_win
 $\neg\{(p_2, \text{NoG})\}$               movi $r3, 1
ld $r4, 0($r1)                        sub $r4, $r3
beq $r3, $r4, l_loop                  st 0($r1), $r4
movi $r3, 1                            ld $r4, 0($r2)
sub $r4, $r3                           add $r4, $r3
st 0($r1), $r4                         st 0($r2), $r4
ld $r4, 0($r2)                          irect
add $r4, $r3      r_win:  $\neg\{(p_4, g_{id})\}$  irect
st 0($r2), $r4
sti
 $\neg\{(p_1, \text{NoG})\}$ 
j l_loop
l_win:   $\neg\{(p_3, \text{NoG})\}$ 
sti
j l_loop

h_entry:   $\neg\{(p_i, g_i)\}$ 
j h_timer
h_timer:   $\neg\{(p_i, g_i)\}$ 
movi $r1, CNT
ld $r2, 0($r1) ; $r2 <- [CNT]
movi $r3, 100
beq $r2, $r3, schd ; if ([CNT]=100)
movi $r3, 1 ; goto schd
add $r2, $r3
st 0($r1), $r2 ; [CNT]++
irect
schd:   $\neg\{(p_0, g_0)\}$ 
movi $r2, 0
st 0($r1), $r2 ; [CNT] := 0
switch
irect

 $p_0 \triangleq \text{enable}_{irect} \wedge (r_1 = \text{CNT})$ 
 $g_0 \triangleq \left\{ \begin{array}{l} \text{CNT} \mapsto - \\ \text{INVO} \end{array} \right\} \wedge (ie = ie') \wedge (is = is')$ 

```

Figure 13. A Preemptive Timer Handler

queues are used to implement synchronization primitives such as locks and condition variables. \mathbb{Q} is a set of thread ids pointing to thread contexts in \mathbb{T} . Note here we do not need a separate \mathbb{Q} for ready threads, which are threads in \mathbb{T} but not blocked:

$$\text{readyQ}(\mathbb{T}, \mathbb{B}) \triangleq \{\text{tid} \mid \text{tid} \in \text{dom}(\mathbb{T}) \wedge \neg \exists w. \text{tid} \in \mathbb{B}(w)\}.$$

We also add three primitive instructions: switch, block and unblock. They correspond to system calls to low-level thread implementations in the real machine (see Fig 3).

The step relation ($\mathbb{W} \mapsto \mathbb{W}'$) of AIM-2 is defined in Fig 12. The switch instruction saves the execution context of the current thread into the thread queue \mathbb{T} , and randomly picks a new thread from $\text{readyQ}(\mathbb{T}, \mathbb{B})$. To let our abstraction fit into the interfaces shown in Fig 3, we require that the interrupt be disabled before switch. This also explains why *ie* is not saved in the thread context, and why it is set to 0 when a new thread is scheduled from \mathbb{T} : the only way to switch control from one thread to the other is to execute switch, which can be executed only if the interrupt is disabled. The “block r_t ” instruction puts the current thread id into the block queue $\mathbb{B}(r_t)$, and switches the control to another thread in $\text{readyQ}(\mathbb{T}, \mathbb{B})$. If there are no other threads in readyQ , the machine stutters (in our x86 implementation, this would never happen because there is an idle thread and our program logic prohibits it from executing block). The “unblock r_t, r_d ” instruction removes a thread from $\mathbb{B}(r_t)$ and puts its *tid* into r_d if the queue is not empty; otherwise r_d contains 0. By the definition of readyQ , we know *tid* will be in readyQ after being unblocked. unblock does not switch controls. Like switch, block and unblock can be executed only if the interrupt is disabled. The effects of other instructions over \mathbb{S} , \mathbb{K} and *pc* are the same as in AIM-1. They do not change \mathbb{T} , \mathbb{B} and *tid*. The transition ($\mathbb{W} \not\mapsto \mathbb{W}'$) for AIM-2 is almost the same as the one for AIM-1 defined by the *IRQ* rule. It does not change \mathbb{T} and *tid* either. The definition of ($\mathbb{W} \mapsto \mathbb{W}'$) is unchanged.

A preemptive timer interrupt handlers The design of the low-level AIM machine is very interesting in that it provides different choices of thread models for high-level concurrent programs (Figure 1). If high-level threads assumes the interrupt is always disabled, they work in the non-preemptive model and each thread voluntarily gives up the control by switch. However, threads cannot handle interrupts from I/O devices either in this case. An alternative approach to the non-preemptive model is to install an interrupt handler implemented in AIM that does no context switching. To

Figure 10. Sample AIM-1 Program: Teeter-Totter

```

(World)  $\mathbb{W} ::= (\mathbb{C}, \mathbb{S}, \mathbb{K}, \text{pc}, \text{tid}, \mathbb{T}, \mathbb{B})$ 
(ThrdSet)  $\mathbb{T} ::= \{\text{tid} \rightsquigarrow (\mathbb{R}, \mathbb{K}, \text{is}, \text{pc})\}^*$ 
(BlkQSet)  $\mathbb{B} ::= \{w \rightsquigarrow \mathbb{Q}\}^*$ 
(ThrdQ)  $\mathbb{Q} ::= \{\text{tid}_1, \dots, \text{tid}_n\}$ 
(ThrdID)  $\text{tid} ::= n$  (nat nums, and  $n > 0$ )
(qID)  $w ::= n$  (nat nums, and  $n > 0$ )
(Instr)  $\mathfrak{I} ::= \dots \mid \text{switch} \mid \text{block } r_t \mid \text{unblock } r_t, r_d \mid \dots$ 

```

Figure 11. AIM-2 defined as an Extension of AIM-1

In AIM, the register file \mathbb{R} is automatically saved and restored at the entry and exit point of the interrupt handler. This is a simplification of the x86 interrupt mechanism for a cleaner presentation and is not essential for the technical development. In our implementation for x86, the handler code needs to save and restore the registers.

Fig 9 also defines ($\mathbb{W} \mapsto \mathbb{W}'$) for executing the instruction at the current *pc*; program execution is then modeled as $\mathbb{W} \mapsto \mathbb{W}'$. Note that our semantics of AIM-1 programs is not deterministic: the state transition may be made either by executing the next instruction or by handling an incoming interrupt. Also, given a \mathbb{W} , there may not always exist a \mathbb{W}' such that ($\mathbb{W} \mapsto \mathbb{W}'$) holds. If there is no such \mathbb{W}' , we say the program gets stuck at \mathbb{W} . One important goal of our program logic is to show that certified programs never get stuck.

Fig 10 shows a sample AIM-1 program. The program specifications in shadowed boxes are explained in Section 4. Initially LEFT and RIGHT point to memory cells containing the same value (say, 50). The non-handler increases the value stored at LEFT and decrease the value at RIGHT. The interrupt handler code does the reverse. Which side wins depends on how frequent the interrupt comes. To avoid races, the non-handler code always disables interrupt before it accesses LEFT and RIGHT.

3.2 AIM-2

Fig 11 defines AIM-2 as an extension over AIM-1. We extend World \mathbb{W} with an abstract thread queue \mathbb{T} , a set of block queues \mathbb{B} , and the id *tid* for the current thread. \mathbb{T} maps a thread id to a thread execution context, which contains the register file, stack, the *is* flag and *pc*. \mathbb{B} maps block queue ids *w* to block queues \mathbb{Q} . These block

$(C, S, K, pc, tid, T, B) \mapsto W'$ where $S = (H, R, ie, is)$	
if $C(pc) =$	$W' =$
switch	$(C, (H, R', 0, is'), K', pc', tid', T', B)$ if $ie = 0$, $T' = T\{tid \rightsquigarrow (R, K, is, pc+1)\}$, $tid' \in \text{readyQ}(T, B')$, and $T'(tid') = (R', K', is', pc')$
block r_t	$(C, (H, R', ie, is'), K', pc', tid', T', B')$ if $ie = 0$, $w = R(r_t)$, $B(w) = Q$, $B' = B\{w \rightsquigarrow (Q \cup \{tid\})\}$, $tid' \in \text{readyQ}(T, B')$, $T(tid') = (R', K', is', pc')$ and $T' = T\{tid \rightsquigarrow (R, K, is, pc+1)\}$
block r_t	$(C, (H, R, ie, is), K, pc, tid, T, B)$ if $ie = 0$, $w = R(r_t)$, and $\text{readyQ}(T, B') = \{tid\}$
unblock r_t, r_d	$(C, (H, R', ie, is), K, pc+1, tid, T, B)$ if $ie = 0$, $w = R(r_t)$, $B(w) = \emptyset$, and $R' = R\{r_d \rightsquigarrow 0\}$
unblock r_t, r_d	$(C, (H, R', ie, is), K, pc+1, tid, T, B')$ if $ie = 0$, $w = R(r_t)$, $B(w) = Q \cup \{tid'\}$, $B' = B\{w \rightsquigarrow Q\}$, and $R' = R\{r_d \rightsquigarrow tid'\}$
other c	$(C, S', K', pc', tid, T, B)$ if $\text{NextS}_{(c, K)} S S'$, $\text{NextK}_{(pc, c)} K K'$, and $\text{NextPC}_{(c, R, K)} pc pc'$

Figure 12. The Step Relation for AIM-2

$(CdHpSpec)$	Ψ	$::= \{(f_1, s_1), \dots, (f_n, s_n)\}$
$(Spec)$	s	$::= (p, g)$
$(Pred)$	p	$\in Stack \rightarrow State \rightarrow Prop$
$(Guarantee)$	g	$\in State \rightarrow State \rightarrow Prop$
$(MPred)$	$m, INV0, INV1$	$\in Heap \rightarrow Prop$
$(WQSpec)$	Δ	$::= \{w \rightsquigarrow m\}^*$

Figure 14. Specification Constructs

get the preemptive model, we install a timer interrupt handler that executes `switch`.

Figure 13 shows the implementation of a preemptive interrupt handler for the timer. Each time the interrupt comes, the handler test the value of the counter at memory location `CNT`. If the counter reaches 100, the handler switches control to other threads; otherwise it increases the counter by 1 and returns to the interrupted thread. We will explain the meaning of specifications and how the timer handler is certified in Section 4.

4. The Program Logic

4.1 Specification Language

We use the mechanized *meta-logic* implemented in the Coq proof assistant as our specification language. The logic corresponds to higher-order logic with inductive definitions.

As shown in Figure 14, the specification Ψ for the code heap C associates code labels f with specifications s . We allow each f to have more than one s , just as a function may have multiple specified interfaces. The specification s is a pair (p, g) . The assertion p is a predicate over a stack K and a program state S (its meta-type in Coq is a function that takes K and S as arguments and returns a proposition), while g is a predicate over two program states. As we can see, the $\text{NextS}_{(c, K)}$ relation defined in Figure 9 is a special form of g . Following our previous work on reasoning low-level code with stack based control abstractions [7], we use p to specify the precondition over stack and state, and use g to specify the guaranteed behavior from the specified program point to the point when the *current* function returns.

We also use the predicate m to specify data heaps. The invariants $INV0$ and $INV1$ shown in Figures 5 and 6 are both heap predicates. We encode in Figure 15 Separation Logic connectors [13, 21] in our specification language. Assertions in Separation Logic capture

$\text{true} \triangleq \lambda H. \text{True}$	$\text{emp} \triangleq \lambda H. H = \emptyset$
$1 \mapsto w \triangleq \lambda H. H = \{1 \rightsquigarrow w\}$	$1 \mapsto _ \triangleq \lambda H. \exists w. (1 \mapsto w) H$
$H_1 \perp H_2 \triangleq \text{dom}(H_1) \cap \text{dom}(H_2) = \emptyset$	
$H_1 \uplus H_2 \triangleq \begin{cases} H_1 \cup H_2 & \text{if } H_1 \perp H_2 \\ \text{undefined} & \text{otherwise} \end{cases}$	
$m_1 * m_2 \triangleq \lambda H. \exists H_1, H_2. (H_1 \uplus H_2 = H) \wedge m_1 H_1 \wedge m_2 H_2$	
$p * m \triangleq \lambda K, S. \exists H_1, H_2. (H_1 \uplus H_2 = S.H) \wedge p K S _{H_1} \wedge m H_2$	
$m \multimap m' \triangleq \lambda H. \forall H', H''. (H \uplus H' = H'') \wedge m H' \rightarrow m' H''$	
$m \multimap p \triangleq \lambda K, S. \forall H, H'. (H \uplus S.H = H') \wedge m H \rightarrow p K S _{H'}$	
$\text{precise}(m) \triangleq \forall H, H_1, H_2. (H_1 \subseteq H) \wedge (H_2 \subseteq H) \wedge m H_1 \wedge m H_2 \rightarrow (H_1 = H_2)$	

Figure 15. Definitions of Separation Logic Assertions

$\text{enable}(p, g) \triangleq \forall K, S. p K S \rightarrow \exists S', g S S'$
$p \triangleright g \triangleq \lambda K, S. \exists S_0, p K S_0 \wedge g S_0 S$
$g \circ g' \triangleq \lambda S, S'. \exists S''. g S S' \wedge g' S' S''$
$p \circ g \triangleq \lambda S, S'. \exists K. p K S \wedge g S S'$
$p \Rightarrow p' \triangleq \forall K, S. p K S \rightarrow p' K S$
$g \Rightarrow g' \triangleq \forall S, S'. g S S' \rightarrow g' S S'$

Figure 16. Connectors for p and g

ownership of heaps. The assertion “ $1 \mapsto n$ ” holds only if the heap has only one cell at 1 containing value n . It can also be interpreted as the ownership of this memory cell. “ $m * m'$ ” means the heap can be split into two *disjoint* parts, and m and m' hold over one of them respectively. A heap satisfies $m \multimap m'$ if and only if the disjoint union of it with any heap satisfying m would satisfy m' .

The specification Δ maps an identifier w to a heap predicate specifying the well-formedness of the resource that the threads in the block queue $B(w)$ are waiting for.

Specification of the Interrupt Handler. We need to give a specification to the interrupt handler to certify the handler code and ensure the non-interference. We let $(h_entry, (p_i, g_i)) \in \Psi$, where p_i and

$$\begin{aligned}
P ? m : m' &\triangleq \lambda \mathbb{H}. (P \wedge m \mathbb{H}) \vee (\neg P \wedge m' \mathbb{H}) \\
g_{\text{cli}} &\triangleq \lambda (\mathbb{H}, \mathbb{R}, \text{ie}, \text{is}), (\mathbb{H}', \mathbb{R}', \text{ie}', \text{is}'). \\
&\quad (\text{is} = \text{is}') \wedge (\mathbb{R} = \mathbb{R}') \wedge (\text{ie}' = 0) \wedge \\
&\quad \left\{ \begin{array}{l} \text{emp} \\ (\text{ie} = 1 \wedge \text{is} = 0) ? (\text{INV0} * \text{INV1}) : \text{emp} \end{array} \right\} \mathbb{H} \mathbb{H}' \\
g_{\text{sti}} &\triangleq \lambda (\mathbb{H}, \mathbb{R}, \text{ie}, \text{is}), (\mathbb{H}', \mathbb{R}', \text{ie}', \text{is}'). \\
&\quad (\text{is} = \text{is}') \wedge (\mathbb{R} = \mathbb{R}') \wedge (\text{ie}' = 1) \wedge \\
&\quad \left\{ \begin{array}{l} (\text{ie} = 0 \wedge \text{is} = 0) ? (\text{INV0} * \text{INV1}) : \text{emp} \\ \text{emp} \end{array} \right\} \mathbb{H} \mathbb{H}' \\
g_{\text{switch}} &\triangleq \lambda (\mathbb{H}, \mathbb{R}, \text{ie}, \text{is}), (\mathbb{H}', \mathbb{R}', \text{ie}', \text{is}'). \\
&\quad (\text{ie} = 0) \wedge (\text{ie} = \text{ie}') \wedge (\mathbb{R} = \mathbb{R}') \wedge (\text{is} = \text{is}') \wedge \\
&\quad \left\{ \begin{array}{l} \text{INV0} * (\text{is} = 0 ? \text{INV1} : \text{emp}) \\ \text{INV0} * (\text{is} = 0 ? \text{INV1} : \text{emp}) \end{array} \right\} \mathbb{H} \mathbb{H}' \\
g_{\text{block } r_s}^{\Delta} &\triangleq \lambda (\mathbb{H}, \mathbb{R}, \text{ie}, \text{is}), (\mathbb{H}', \mathbb{R}', \text{ie}', \text{is}'). \\
&\quad (\text{ie} = 0) \wedge (\text{ie} = \text{ie}') \wedge (\mathbb{R} = \mathbb{R}') \wedge (\text{is} = \text{is}') \wedge \\
&\quad \exists m. \Delta(\mathbb{R}(r_s)) = m \wedge \\
&\quad \left\{ \begin{array}{l} \text{INV0} * (\text{is} = 0 ? \text{INV1} : \text{emp}) \\ \text{INV0} * (\text{is} = 0 ? \text{INV1} : \text{emp}) * m \end{array} \right\} \mathbb{H} \mathbb{H}' \\
g_{\text{unblock } r_s, r_d}^{\Delta} &\triangleq \lambda (\mathbb{H}, \mathbb{R}, \text{ie}, \text{is}), (\mathbb{H}', \mathbb{R}', \text{ie}', \text{is}'). \\
&\quad (\text{ie} = 0) \wedge (\text{ie} = \text{ie}') \wedge (\text{is} = \text{is}') \wedge \\
&\quad (\forall r \neq r_d. \mathbb{R}(r) = \mathbb{R}'(r)) \wedge \\
&\quad \exists m. \Delta(\mathbb{R}(r_s)) = m \wedge (m * \text{true}) \mathbb{H} \wedge \\
&\quad \left\{ \begin{array}{l} (\mathbb{R}'(r_d) = 0) ? \text{emp} : m \\ \text{emp} \end{array} \right\} \mathbb{H} \mathbb{H}'
\end{aligned}$$

Figure 17. Semantics for cli, sti and Thread Primitives

g_i are defined as follows:

$$p_i \triangleq \lambda \mathbb{K}, \mathbb{S}. ((\text{INV0} * \text{true}) \mathbb{S} \cdot \mathbb{H}) \wedge (\mathbb{S} \cdot \text{is} = 1) \wedge (\mathbb{S} \cdot \text{ie} = 0) \wedge \exists \mathbb{f}, \mathbb{R}, \mathbb{K}'. \mathbb{K} = (\mathbb{f}, \mathbb{R}) :: \mathbb{K}' \quad (1)$$

$$g_i \triangleq \lambda \mathbb{S}, \mathbb{S}'. \left\{ \begin{array}{l} \text{INV0} \\ \text{INV0} \end{array} \right\} \mathbb{S} \cdot \mathbb{H} \mathbb{S}' \cdot \mathbb{H} \wedge (\mathbb{S}' \cdot \text{ie} = \mathbb{S} \cdot \text{ie}) \wedge (\mathbb{S}' \cdot \text{is} = \mathbb{S} \cdot \text{is}) \quad (2)$$

The precondition p_i specifies the stack and state at the entry `h_entry`. It requires that the local heap used by the handler (block A in Figure 5) satisfies `INV0`. `INV0` is a global parameter of our system, whose definition depends on the functionality of the interrupt handler. The guarantee g_i specifies the behavior of the handler. The arguments \mathbb{S} and \mathbb{S}' correspond to program states at the entry and exit points, respectively. It says the `ie` and `is` bits in \mathbb{S}' have the same value as in \mathbb{S} , and the handler's local heap satisfies `INV0` in \mathbb{S} and \mathbb{S}' , while the rest of the heap remains unchanged. The predicate $\left\{ \begin{array}{l} m_1 \\ m_2 \end{array} \right\}$ is defined below.

$$\left\{ \begin{array}{l} m_1 \\ m_2 \end{array} \right\} \triangleq \lambda \mathbb{H}_1, \mathbb{H}_2. \exists \mathbb{H}'_1, \mathbb{H}'_2, \mathbb{H}. (m_1 \mathbb{H}'_1) \wedge (m_2 \mathbb{H}'_2) \wedge (\mathbb{H}'_1 \uplus \mathbb{H} = \mathbb{H}_1) \wedge (\mathbb{H}'_2 \uplus \mathbb{H} = \mathbb{H}_2) \quad (3)$$

It has the following nice monotonicity: for any $\mathbb{H}_1, \mathbb{H}_2$ and \mathbb{H}' , if \mathbb{H}_1 and \mathbb{H}_2 satisfy the predicate, $\mathbb{H}_1 \perp \mathbb{H}'$, and $\mathbb{H}_2 \perp \mathbb{H}'$, then $\mathbb{H}_1 \uplus \mathbb{H}'$ and $\mathbb{H}_2 \uplus \mathbb{H}'$ satisfy the predicate.

4.2 Inference Rules

Inference rules of the program logic are shown in Figs. 18 and 20. The judgment for well-formed instruction sequences says it is safe to execute the instruction sequence given the imported interface Ψ , the specification of block queues Δ , and a precondition (p, g) .

$$\boxed{\Psi, \Delta \vdash \{s\} f : \mathbb{I}} \quad (\text{Well-Formed Instr. Seq.})$$

$$\frac{\iota \notin \{\text{call} \dots, \text{beq} \dots\} \quad \Psi, \Delta \vdash \{(p', g')\} f+1 : \mathbb{I} \quad \text{enable}(p, g_i) \quad (p \triangleright g_i) \Rightarrow p' \quad (p \circ (g_i \circ g')) \Rightarrow g}{\Psi, \Delta \vdash \{(p, g)\} f : \iota; \mathbb{I}} \quad (\text{SEQ})$$

$$\frac{(\mathbb{f}+1, (p'', g'')) \in \Psi \quad \Psi, \Delta \vdash \{(p'', g'')\} \mathbb{f}+1 : \mathbb{I} \quad (\mathbb{f}', (p', g')) \in \Psi \quad \forall \mathbb{K}, \mathbb{S}, \text{pc}. p \mathbb{K} \mathbb{S} \rightarrow p' (\text{pc} :: \mathbb{K}) \mathbb{S} \quad (p \triangleright g') \Rightarrow p'' \quad (p \circ (g' \circ g'')) \Rightarrow g}{\Psi, \Delta \vdash \{(p, g)\} f : \text{call } \mathbb{f}'; \mathbb{I}} \quad (\text{CALL})$$

$$\frac{(\mathbb{f}', (p', g')) \in \Psi \quad \Psi, \Delta \vdash \{(p'', g'')\} \mathbb{f}+1 : \mathbb{I} \quad (p \triangleright \text{gid}_{r_s=r_t}) \Rightarrow p' \quad (p \circ (\text{gid}_{r_s=r_t} \circ g')) \Rightarrow g \quad (p \triangleright \text{gid}_{r_s \neq r_t}) \Rightarrow p'' \quad (p \circ (\text{gid}_{r_s \neq r_t} \circ g'')) \Rightarrow g}{\Psi, \Delta \vdash \{(p, g)\} f : \text{beq } r_s, r_t, \mathbb{f}'; \mathbb{I}} \quad (\text{BEQ})$$

$$\frac{p \Rightarrow \text{enable}_{\text{ret}} \quad (p \circ \text{gid}) \Rightarrow g}{\Psi, \Delta \vdash \{(p, g)\} f : \text{ret}} \quad (\text{IRET})$$

where $\text{enable}_{\text{ret}} \triangleq \lambda \mathbb{K}, \mathbb{S}. \exists \mathbb{f}, \mathbb{R}, \mathbb{K}'. \mathbb{K} = (\mathbb{f}, \mathbb{R}) :: \mathbb{K}' \wedge \mathbb{S} \cdot \text{is} = 1$

$$\frac{p \Rightarrow \text{enable}_{\text{ret}} \quad (p \circ \text{gid}) \Rightarrow g}{\Psi, \Delta \vdash \{(p, g)\} f : \text{ret}} \quad (\text{RET})$$

where $\text{enable}_{\text{ret}} \triangleq \lambda \mathbb{K}, \mathbb{S}. \exists \mathbb{f}, \mathbb{K}'. \mathbb{K} = \mathbb{f} :: \mathbb{K}'$

$$\frac{(\mathbb{f}', (p', g')) \in \Psi \quad p \Rightarrow p' \quad (p \circ g') \Rightarrow g}{\Psi, \Delta \vdash \{(p, g)\} f : \mathbb{f}'} \quad (\text{J})$$

$$\boxed{\Psi, \Delta \vdash \mathbb{C} : \Psi'} \quad (\text{Well-Formed Code Heap})$$

$$\frac{\text{for all } (\mathbb{f}, s) \in \Psi' : \quad \Psi, \Delta \vdash \{s\} f : \mathbb{C}[\mathbb{f}]}{\Psi, \Delta \vdash \mathbb{C} : \Psi'} \quad (\text{CDHP})$$

Figure 18. Inference Rules

The predicate p specifies the current stack and state, and g specifies the state transition from the current program point to the return point of the current function (or the interrupt handler).

The `SEQ` rule is a schema for instruction sequences starting with instructions except branch and function call instructions. We need to find an intermediate specification (p', g') , with respect to which the remaining instruction sequence is well-formed. It is also used as a post-condition for the first instruction. We use g_i to specify the state transition made by the instruction ι . The premise $\text{enable}(p, g_i)$ is defined in Figure 16. It means that the state transition g_i would not get stuck as long as the starting stack and state satisfy p . The predicate $p \triangleright g_i$, shown in Figure 16, specifies the stack and state resulting from the state transition g_i , knowing the initial state satisfies p . It is the strongest post condition after g_i . The composition of two subsequent transitions g and g' is represented as $g \circ g'$, and $p \circ g$ refines g with the extra knowledge that the initial state satisfies p . We also lift the implication relation between p 's and g 's. The last premise in the `SEQ` rule requires the composition of g_i and g' fulfills g , knowing the current state satisfies p .

If ι is an arithmetic instruction, move instruction or memory operation, we define g_i as `NextS`_{($\iota, _$) (see Figure 9). Since `NextS` does not depend on the stack for these instructions, we use “ $_$ ” to represent arbitrary stacks. Also note that the `NextS` relations for `ld` or `st` require the target address to be in the domain of heap, therefore the premise $\text{enable}(p, g_i)$ requires that p contains the ownership of the target memory cell accessed by `ld` or `st`.}

Interrupts and thread primitive instructions. One of the major technical contributions of this paper is our formulation of g_i for `cli`, `sti` and `switch`, `block` and `unblock`, which, as shown in Figure 17, gives an axiomatic ownership transfer semantics to them.

The transition g_{cli} says that, if cli is executed in the non-handler ($is = 0$) and the interrupt is enabled ($ie = 1$), the current thread gets ownership of the well-formed sub-heap A and C satisfying $INV0 * INV1$, as shown in Figure 5; otherwise there is no ownership transfer because the interrupt has already been disabled before cli . The transition g_{sti} is defined similarly. Note that the premise $enable(p, g_i)$ in the SEQ requires that, before executing sti , the precondition p must contain the ownership ($ie = 0 \wedge is = 0$)? ($INV1 * INV0$): emp .

The transition g_{switch} for $switch$ requires that the sub-heap A and C (in Figure 5) be well-formed before and after $switch$. However, if we execute $switch$ in the interrupt handler ($is = 1$), we know $INV1$ always holds and leave it implicit. Also $enable(p, g_i)$ requires that the precondition p ensures $ie = 0$ and $INV0 * (is = 0 ? INV1 : emp)$ holds over some sub-heap.

The transition $g_{block\ r_s}^\Delta$ for block r_s refers to the specification Δ . It requires $ie = 0$ and r_s contains an identifier of a block queue with specification m in Δ . It is similar to $switch$, except that the thread gets the ownership of m after it is released (see Figure 6). In $g_{unblock\ r_s, r_d}^\Delta$, we require the initial heap must contains a sub-heap satisfying m , because $unblock$ may transfer it to a blocked thread. However, since $unblock$ does not immediately switch controls, we do not need the sub-heap A and C to be well-formed. If the target address r_d contains non-zero value at the end of $unblock$, some thread has been released from the block queue. The current thread transfers m to the released thread and has no access to it any more. Otherwise, no thread is released and there is no ownership transfer.

Function calls and memory polymorphism. In the $CALL$ rule, we treat the state transition g' made by the callee as the transition of the call instruction. We also require that the precondition p implies the precondition p' of the callee, which corresponds to the $enable$ premise in the SEQ rule. Note that the $CALL$ rule supports memory polymorphism in a natural way: if g' of the callee is specified following the pattern $\left\{ \begin{smallmatrix} m \\ m' \end{smallmatrix} \right\}$ defined by formula (3), the callee's specification does not need to mention data required by the caller but not accessed in the callee, thus achieving a similar effects to the frame rule in Separation Logic. We first introduce the following definitions.

$$\begin{aligned} \text{monotonic}(p) &\triangleq \forall \mathbb{K}, \mathbb{S}, \mathbb{H}, \mathbb{H}'. (\mathbb{H}' = \mathbb{H} \uplus \mathbb{S} \cdot \mathbb{H}) \wedge p \mathbb{K} \mathbb{S} \rightarrow p \mathbb{K} \mathbb{S} |_{\mathbb{H}'} \\ \text{monotonic}(g) &\triangleq \forall \mathbb{S}, \mathbb{S}', \mathbb{H}, \mathbb{H}', \mathbb{H}_0. (\mathbb{H}_0 \uplus \mathbb{S} \cdot \mathbb{H} = \mathbb{H}) \wedge (\mathbb{H}_0 \uplus \mathbb{S}' \cdot \mathbb{H} = \mathbb{H}') \\ &\quad \rightarrow (g \mathbb{S} \mathbb{S}' \rightarrow g \mathbb{S} |_{\mathbb{H}} \mathbb{S}' |_{\mathbb{H}'}) \\ \text{frame}(g) &\triangleq \forall \mathbb{S}, \mathbb{S}', \mathbb{H}_0, \mathbb{H}. (\mathbb{H}_0 \uplus \mathbb{H} = \mathbb{S} \cdot \mathbb{H}) \wedge (g \mathbb{S} \mathbb{S}') \wedge (\exists \mathbb{S}'' . g \mathbb{S} |_{\mathbb{H}} \mathbb{S}'') \\ &\quad \rightarrow \exists \mathbb{H}'. (\mathbb{H}_0 \uplus \mathbb{H}' = \mathbb{S}' \cdot \mathbb{H}) \wedge g \mathbb{S} |_{\mathbb{H}} \mathbb{S}' |_{\mathbb{H}'} \\ \text{wff_spec}(p, g) &\triangleq \text{monotonic}(p) \wedge \text{monotonic}(g) \wedge \text{frame}(g) \wedge \text{enable}(p, g) \\ \text{wff_spec}(\Psi) &\triangleq \forall f, s. ((f, s) \in \Psi) \rightarrow \text{wff_spec}(s) \end{aligned}$$

Lemma 4.1 (call-frame)

If $\Psi, \Delta \vdash \{(p, g)\} f : \text{call } f \parallel, \text{wff_spec}(\Psi)$ and $\text{wff_spec}(p, g)$, then we have $\Psi, \Delta \vdash \{(p * m, g)\} f : \text{call } f \parallel$ for any m such that $\text{precise}(m)$.

Proof sketch. To prove the lemma, we need to prove the following propositions:

- if $\text{monotonic}(p')$, then $(p \Rightarrow p') \rightarrow (p * m \Rightarrow p')$;
- if $p \Rightarrow p'$, $enable(p', g')$, $frame(g')$ and $\text{monotonic}(p'')$, then $(p \triangleright g' \Rightarrow p'') \rightarrow ((p * m) \triangleright g' \Rightarrow p'')$;

$$p * \text{Inv} \triangleq \lambda \mathbb{K}, \mathbb{S}. (p * \text{Inv}(\mathbb{S} \cdot ie, \mathbb{S} \cdot is)) \mathbb{K} \mathbb{S}$$

$$\text{Inv}(ie, is) \triangleq \begin{cases} INV1 & is = 1 \\ emp & is = 0 \text{ and } ie = 0 \\ INV_s & is = 0 \text{ and } ie = 1 \end{cases}$$

$$\text{where } INV_s \triangleq INV0 * INV1$$

$$\begin{aligned} [g]_{(m, m')} &\triangleq \lambda \mathbb{S}, \mathbb{S}'. \exists \mathbb{H}_1, \mathbb{H}_2, \mathbb{H}'_1, \mathbb{H}'_2. \\ &\quad (\mathbb{H}_1 \uplus \mathbb{H}_2 = \mathbb{S} \cdot \mathbb{H}) \wedge (\mathbb{H}'_1 \uplus \mathbb{H}'_2 = \mathbb{S}' \cdot \mathbb{H}) \\ &\quad \wedge m \mathbb{H}_2 \wedge m' \mathbb{H}'_2 \wedge g \mathbb{S} |_{\mathbb{H}_1} \mathbb{S}' |_{\mathbb{H}'_1} \end{aligned}$$

$$[g] \triangleq \lambda \mathbb{S}, \mathbb{S}'. [g]_{(\text{Inv}(\mathbb{S} \cdot ie, \mathbb{S} \cdot is), \text{Inv}(\mathbb{S}' \cdot ie, \mathbb{S}' \cdot is))}$$

$$\text{WFST}(g, \mathbb{S}, \text{nil}, \Psi) \triangleq \neg \exists \mathbb{S}'. g \mathbb{S} \mathbb{S}'$$

$$\begin{aligned} \text{WFST}(g, \mathbb{S}, f :: \mathbb{K}, \Psi) &\triangleq \\ &\quad \exists p_f, g_f. (f, (p_f, g_f)) \in \Psi \\ &\quad \wedge \forall \mathbb{S}'. g \mathbb{S} \mathbb{S}' \rightarrow (p_f * \text{Inv}) \mathbb{K} \mathbb{S}' \wedge \text{WFST}([g_f], \mathbb{S}', \mathbb{K}, \Psi) \end{aligned}$$

$$\begin{aligned} \text{WFST}(g, \mathbb{S}, (f, \mathbb{R}) :: \mathbb{K}, \Psi) &\triangleq \\ &\quad \exists p_f, g_f. (f, (p_f, g_f)) \in \Psi \\ &\quad \wedge \forall \mathbb{S}'. g \mathbb{S} \mathbb{S}' \rightarrow (p_f * \text{Inv}) \mathbb{K} \mathbb{S}'' \wedge \text{WFST}([g_f], \mathbb{S}'', \mathbb{K}, \Psi) \\ &\quad \text{where } \mathbb{S}'' = (\mathbb{S}' \cdot \mathbb{H}, \mathbb{R}, 1, 0) \end{aligned}$$

Figure 19. Auxiliary Definitions for Program Invariants

- if $p \Rightarrow p'$, $(p \triangleright g' \Rightarrow p'')$, $enable(p', g')$, $frame(g')$, $enable(p'', g'')$, $frame(g'')$, and $\text{monotonic}(g)$, then $((p \circ (g' \circ g'')) \Rightarrow g) \rightarrow (((p * m) \circ (g' \circ g'')) \Rightarrow g)$.
- Each proposition can be proved by above definitions. \square

Also, we can prove $\text{wff_spec}(p, g)$ if:

$$\begin{aligned} p &\triangleq \lambda \mathbb{K}, \mathbb{S}. m * \text{true } \mathbb{S} \cdot \mathbb{H} \\ g &\triangleq \lambda \mathbb{S}, \mathbb{S}'. \left\{ \begin{smallmatrix} m \\ m' \end{smallmatrix} \right\} \mathbb{S} \cdot \mathbb{H} \mathbb{S}' \cdot \mathbb{H}, \end{aligned}$$

for any m, m' such that $\text{precise}(m)$.

Other instructions. In the BEQ rule, we use $\text{gid}_{r_s=r_t}$ to represent an identity transition with the extra knowledge that r_s and r_t contain the same value. $\text{gid}_{r_s \neq r_t}$ is defined similarly. We do not have an $enable$ premise because the branch instruction never gets stuck. $IRET$ and RET rules require that the code has finished its guaranteed transition at this point. So an identity transition gid should satisfy the remaining transition g . The predicates enable_{iret} and enable_{ret} specify the requirements over stacks. The J rule can be viewed as a specialization of the BEQ rule.

Well-formed code heaps. The $CDHP$ rule says the code heap is well-formed if and only if each instruction sequence specified in Ψ' is well-formed.

Program Invariants. The program invariant enforced by our program logic is defined by the $PROG$ rule in Figure 20, which is another major technical contribution of this work. It says that, if there are n threads in the thread queue in addition to the current thread, the heap can be split into $n + 1$ blocks, each for one thread. Each block \mathbb{H}_k ($k > 0$) is for a ready or blocked thread in queues. The block \mathbb{H}_0 is assigned to the current thread, which includes both its private memory and the shared memory A and C shown in Figure 5. The code heap needs to be well-formed, defined by the $CDHP$ rule. The domain of Δ should be the same with the domain of \mathbb{B} , *i.e.*, Δ only specifies block queues in \mathbb{B} .

The $PROG$ rule also requires that the current thread, ready threads and blocked threads are all well-formed. The CTH rule defines the well-formedness of the current thread. It requires that the pc have a specification (p, g) in Ψ' . By the well-formedness of the code

$$\begin{array}{c}
\mathbb{T} \setminus \text{tid} = \{\text{tid}_1 \rightsquigarrow (\mathbb{R}_1, \mathbb{K}_1, \text{is}_1, \text{pc}_1), \dots, \\
\quad \text{tid}_n \rightsquigarrow (\mathbb{R}_n, \mathbb{K}_n, \text{is}_n, \text{pc}_n)\} \\
\mathbb{S}. \mathbb{H} = \mathbb{H}_0 \uplus \dots \uplus \mathbb{H}_n \quad \mathbb{S}_0 = \mathbb{S}|_{\mathbb{H}_0} \\
\Psi, \Delta \vdash \mathbb{C}: \Psi' \quad \Psi \subseteq \Psi' \quad \text{dom}(\Delta) = \text{dom}(\mathbb{B}) \quad \Psi' \vdash^{\mathbb{C}} (\mathbb{S}_0, \mathbb{K}, \text{pc}) \\
\text{for all } 0 < k \leq n \text{ such that } \text{tid}_k \in \text{readyQ}(\mathbb{T}, \mathbb{B}): \\
\quad \Psi' \vdash^{\mathbb{R}} (\mathbb{H}_k, \mathbb{R}_k, \mathbb{K}_k, \text{is}_k, \text{pc}_k) \\
\text{for all } w, \text{tid}_j \text{ such that } \text{tid}_j \in \mathbb{B}(w): \\
\quad \Psi', \Delta, w \vdash^{\mathbb{W}} (\mathbb{H}_j, \mathbb{R}_j, \mathbb{K}_j, \text{is}_j, \text{pc}_j) \\
\hline
\Psi, \Delta \vdash (\mathbb{C}, \mathbb{S}, \mathbb{K}, \text{pc}, \text{tid}, \mathbb{T}, \mathbb{B}) \quad (\text{PROG}) \\
\hline
\frac{(\text{pc}, (\text{p}, \text{g})) \in \Psi' \quad (\text{p} * \text{Inv}) \mathbb{K} \mathbb{S} \quad \text{WFST}(\lfloor \text{g} \rfloor, \mathbb{S}, \mathbb{K}, \Psi')}{\Psi' \vdash^{\mathbb{C}} (\mathbb{S}, \mathbb{K}, \text{pc})} \quad (\text{CTH}) \\
\hline
\frac{\mathbb{S}_k = (\mathbb{H}_k, \mathbb{R}_k, 0, \text{is}_k) \quad (\text{INV}_s * (\Psi' \vdash^{\mathbb{C}} (_, _, \text{pc}_k))) \mathbb{K}_k \mathbb{S}_k}{\Psi' \vdash^{\mathbb{R}} (\mathbb{H}_k, \mathbb{R}_k, \mathbb{K}_k, \text{is}_k, \text{pc}_k)} \quad (\text{RDY}) \\
\hline
\frac{\Delta(w) = \text{m} \quad (\text{m} * (\Psi' \vdash^{\mathbb{R}} (_, \mathbb{R}_j, \mathbb{K}_j, \text{is}_j, \text{pc}_j))) \mathbb{H}_j}{\Psi', \Delta, w \vdash^{\mathbb{W}} (\mathbb{H}_j, \mathbb{R}_j, \mathbb{K}_j, \text{is}_j, \text{pc}_j)} \quad (\text{WAIT}) \\
\hline
\end{array}$$

Figure 20. Inference Rules (cont'd)

heap we know PC points to a well-formed instruction sequence. Also $\text{p} * \text{Inv}$ holds over the stack and state with the sub-heap \mathbb{H}_0 . The definition of $\text{p} * \text{Inv}$ is shown in Figure 19. It specifies the shared memory inaccessible from p . If the current program point is in the interrupt handler ($\text{is} = 1$), p leaves the memory block \mathbb{C} (in Figure 5) unspecified, therefore Inv requires it to satisfy INV1 . Otherwise ($\text{is} = 0$), if $\text{ie} = 0$, memory \mathbb{C} and \mathbb{A} become the current thread's private memory and there is no memory to share. If $\text{ie} = 1$, memory \mathbb{C} and \mathbb{A} are not accessible from p , therefore Inv requires them to be well-formed. The inductively defined predicate WFST is shown in Figure 19. It says there exists a well-formed stack with some depth k . At the end of the current function, when g is fulfilled, there must be a stack frame on top of the stack with a pc pointing to well-formed instruction sequences. Also there is a well-formed stack with depth $k-1$ at the returning state.

Since a judgment is represented as a proposition in our meta-logic, and we reuse the meta-logic as our specification language, we can define $\Psi' \vdash^{\mathbb{C}} (_, _, \text{pc})$ as a special predicate $\text{p}: \lambda \mathbb{K}, \mathbb{S}. \Psi' \vdash^{\mathbb{C}} (\mathbb{S}, \mathbb{K}, \text{pc})$. Then the definition of well-formed ready threads in the RDY rule becomes very straightforward: if the *ready* thread gets the extra ownership of shared memory \mathbb{C} and \mathbb{A} , it becomes a well-formed *current* thread (see Fig. 5). Recall that $\text{m} * \text{p}$ is defined in Fig. 15.

Similarly, we define $\Psi' \vdash^{\mathbb{R}} (_, \mathbb{R}, \mathbb{K}, \text{is}, \text{pc})$ as the memory predicate $\lambda \mathbb{H}. \Psi' \vdash^{\mathbb{R}} (_, \mathbb{R}, \mathbb{K}, \text{is}, \text{pc})$. The WAIT rule says that the *waiting* thread in a block queue with identifier w becomes a well-formed *ready* thread if it gets extra ownership of memory $\Delta(w)$. This is also illustrated in Figure 6. The memory predicate $\text{m} * \text{m}'$ is defined in Figure 15.

4.3 Soundness

We prove the soundness of the program logic following the syntactic approach. The progress lemma shows that the program invariant defined by the PROG rule ensures the next instruction does not get stuck. The preservation lemma says the new state after executing the next instruction satisfies the invariant. Therefore, the program never gets stuck as long as the initial state satisfies the invariant. *More importantly*, the invariant always holds during execution, from which we can derive rich properties of programs. For instance, we know there is always a partition of heap for all threads, based on which we can derive the well-formedness of the current thread, ready threads and waiting threads. This can be viewed as

a guarantee of non-interference. Here, we only show a soundness theorem formalizing the partial correctness of programs.

Lemma 4.2 (Progress)

If $\Psi, \Delta \vdash \mathbb{W}$, then there exists \mathbb{W}' such that $(\mathbb{W} \mapsto \mathbb{W}')$.

Proof sketch. Suppose $\mathbb{W} = (\mathbb{C}, \mathbb{S}, \mathbb{K}, \text{pc}, \text{tid}, \mathbb{T}, \mathbb{B})$. By the definition of the operational semantics, execution of $\mathbb{C}(\text{pc})$ would always succeed unless $\mathbb{C}(\text{pc})$ is one of ld , st , switch , block , unlock , ret or iret instructions. By $\Psi, \Delta \vdash \mathbb{W}$ we know there exist Ψ' and (p, g) such that $(\text{pc}, (\text{p}, \text{g})) \in \Psi'$, $\text{p} * \text{Inv} \mathbb{K} \mathbb{S}$, and $\Psi, \Delta \vdash \{\text{s}\} \mathbb{C}[\text{pc}]$. Then by the premise $\text{enable}(\text{p}, \text{g}_i)$ in the SEQ rule we know the first 5 instructions would not get stuck, and, by $\text{enable}_{\text{ret}}$ and the RET rule or $\text{enable}_{\text{iret}}$ and the IRET rule, we know the last two instructions would not get stuck either. \square

Lemma 4.3 says the program can always reach the entry point of the interrupt handler as long as the interrupt is enabled and there is no interrupts being serviced.

Lemma 4.3 (IRQ-Progress)

If $\mathbb{W}. \mathbb{S}. \text{ie} = 1$ and $\mathbb{W}. \mathbb{S}. \text{is} = 0$, there always exists \mathbb{W}' such that $\mathbb{W} \not\leq \mathbb{W}'$.

Proof sketch. Obvious by the IRQ rule shown in Section 3.1. \square

The following two lemmas show that the interrupt handler does not interfere with the non-handler code.

Lemma 4.4 (presv-p)

For all $\text{p}, \mathbb{K}, \mathbb{S}$ and \mathbb{S}' , if $\mathbb{S} = (\mathbb{H}, \mathbb{R}, 1, 0)$, $(\text{p} * \text{Inv}) \mathbb{K} \mathbb{S}$, $\text{precise}(\text{INV0})$, $\text{precise}(\text{INV1})$, and $\lfloor \text{g}_i \rfloor (\mathbb{H}, \mathbb{R}, 0, 1) \mathbb{S}'$, then $(\text{p} * \text{Inv}) \mathbb{K} (\mathbb{S}'. \mathbb{H}, \mathbb{R}, 1, 0)$.

Lemma 4.5 (presv-g)

For all g and \mathbb{S} , if $\mathbb{S} = (\mathbb{H}, \mathbb{R}, 1, 0)$ for certain \mathbb{H} and \mathbb{R} , $(\text{INV}_s * \text{true}) \mathbb{H}$, $\text{precise}(\text{INV0})$ and $\text{precise}(\text{INV1})$, then,

$$\forall \mathbb{S}', \mathbb{S}'' . (\lfloor \text{g}_i \rfloor \mathbb{S}|_{\{\text{ie}=0, \text{is}=1\}} \mathbb{S}') \wedge (\lfloor \text{g} \rfloor (\mathbb{S}'. \mathbb{H}, \mathbb{R}, 1, 0) \mathbb{S}'') \rightarrow (\lfloor \text{g} \rfloor \mathbb{S} \mathbb{S}'')$$

The following lemmas illustrate how the program invariant is preserved by cli , sti and thread primitives, based on our definition of their ownership transfer semantics.

Lemma 4.6 (cli)

If $\text{NextS}_{(\text{cli}, \mathbb{K})} \mathbb{S} \mathbb{S}'$, and $(\text{p} * \text{Inv}) \mathbb{K} \mathbb{S}$, then $(\text{sp}(\text{g}_{\text{cli}}, \text{p}) * \text{Inv}) \mathbb{K} \mathbb{S}'$.

Lemma 4.7 (sti)

If $\text{NextS}_{(\text{sti}, \mathbb{K})} \mathbb{S} \mathbb{S}'$, $(\text{p} * \text{Inv}) \mathbb{K} \mathbb{S}$, and $\text{p} \Rightarrow \text{Inv}(1 - \mathbb{S}. \text{ie}, \mathbb{S}. \text{is}) * \text{true}$, then $(\text{sp}(\text{g}_{\text{sti}}, \text{p}) * \text{Inv}) \mathbb{K} \mathbb{S}'$.

Lemma 4.8 (switch)

If $\text{enable}(\text{p}, \text{g}_{\text{switch}})$, $(\text{p} \triangleright \text{g}_{\text{switch}}) \Rightarrow \text{p}'$, and $(\text{p} * \text{Inv}) \mathbb{K} \mathbb{S}$, then

- there exist \mathbb{H}_1 and \mathbb{H}_2 such that $\text{INV}_s \mathbb{H}_1$ and $\mathbb{H}_1 \uplus \mathbb{H}_2 = \mathbb{S}. \mathbb{H}$;
- $\lfloor \text{g}_{\text{sw}} \rfloor \mathbb{S} \mathbb{S}$;
- for all \mathbb{H}'_1 , if $\text{INV}_s \mathbb{H}'_1$, $\mathbb{H}'_1 \perp \mathbb{H}_2$, and $\mathbb{S}' = \mathbb{S}|_{\mathbb{H}'_1 \cup \mathbb{H}_2}$, then $(\text{p}' * \text{Inv}) \mathbb{K} \mathbb{S}'$ and $\lfloor \text{g}_{\text{sw}} \rfloor \mathbb{S} \mathbb{S}'$.

Lemma 4.9 (Ready2Running)

If $\Psi \vdash^{\mathbb{R}} (\mathbb{H}, \mathbb{R}, \mathbb{K}, \text{is}, \text{pc})$, then

- pc is a valid code pointer specified in Ψ with some specification (p, g) ;
- for all \mathbb{H}' , if $\text{INV}_s \mathbb{H}'$, $\mathbb{H}' \perp \mathbb{H}$, and $\mathbb{S} = (\mathbb{H} \cup \mathbb{H}', \mathbb{R}, 0, \text{is})$, then $(\text{p} * \text{Inv}) \mathbb{K} \mathbb{S}$.

Lemma 4.10 (block)

If $\text{enable}(p, g_{\text{block } r_s}^\Delta)$, $(p \triangleright g_{\text{block } r_s}^\Delta) \Rightarrow p'$, and $(p * \text{Inv}) \mathbb{K} \mathbb{S}$, then

- there exist \mathbb{H}_1 and \mathbb{H}_2 such that $\text{INV}_s \mathbb{H}_1$ and $\mathbb{H}_1 \uplus \mathbb{H}_2 = \mathbb{S}.\mathbb{H}$;
- there exist m such that $\Delta(\mathbb{S}.\mathbb{R}(r_s)) = m$;
- for all $\mathbb{H}'_1, \mathbb{H}'_3$ and \mathbb{H} , if $\text{INV}_s \mathbb{H}'_1, m' \mathbb{H}'_3, \mathbb{H} = \mathbb{H}'_1 \uplus \mathbb{H}_2 \uplus \mathbb{H}'_3$, and $\mathbb{S}' = \mathbb{S}|_{\mathbb{H}}$, then $(p' * \text{Inv}) \mathbb{K} \mathbb{S}'$ and $\lfloor g_{\text{block } r_s}^\Delta \rfloor \mathbb{S} \mathbb{S}'$.

Lemma 4.11 (unblock)

If $\text{enable}(p, g_{\text{unblock } r_s, r_d}^\Delta)$, $(p \triangleright g_{\text{block } r_s}^\Delta) \Rightarrow p'$, and $(p * \text{Inv}) \mathbb{K} \mathbb{S}$ (where $\mathbb{S} = (\mathbb{H}, \mathbb{R}, \text{ie}, \text{is})$), then

- there exists m such that $\Delta(\mathbb{R}(r_s)) = m$;
- there exist \mathbb{H}_1 and \mathbb{H}_2 such that $\mathbb{H}_1 \uplus \mathbb{H}_2 = \mathbb{H}$ and $m \mathbb{H}_1$;
- let $\mathbb{S}' = (\mathbb{H}, \mathbb{R}\{r_d \rightsquigarrow 0\}, \text{ie}, \text{is})$, then $\lfloor g_{\text{block } r_s}^\Delta \rfloor \mathbb{S} \mathbb{S}'$ and $(p' * \text{Inv}) \mathbb{K} \mathbb{S}'$;
- for all $n > 0$, let $\mathbb{S}' = (\mathbb{H}_2, \mathbb{R}\{r_d \rightsquigarrow n\}, \text{ie}, \text{is})$, then $\lfloor g_{\text{block } r_s}^\Delta \rfloor \mathbb{S} \mathbb{S}'$ and $(p' * \text{Inv}) \mathbb{K} \mathbb{S}'$.

Lemma 4.12 (Waiting2Ready)

If $\Psi, \Delta, w \vdash^w (\mathbb{H}, \mathbb{R}, \mathbb{K}, \text{is}, \text{pc})$, then

- there exists m such that $\Delta(w) = m$;
- for any \mathbb{H}' , if $\mathbb{H} \perp \mathbb{H}'$ and $m \mathbb{H}'$, then $\Psi \vdash^R (\mathbb{H} \uplus \mathbb{H}', \mathbb{R}, \mathbb{K}, \text{is}, \text{pc})$.

The following lemmas are also used to prove the preservation lemma.

Lemma 4.13 (Stack-Strengthen)

For all g, g', \mathbb{S} and \mathbb{S}' , if $\forall \mathbb{S}'' . g \mathbb{S} \mathbb{S}'' \rightarrow g' \mathbb{S}' \mathbb{S}''$, and $\text{WFST}(g', \mathbb{S}', \mathbb{K}, \Psi)$, then $\text{WFST}(g, \mathbb{S}, \mathbb{K}, \Psi)$.

Lemma 4.14 (Code Heap)

If $\Psi, \Delta \vdash \mathbb{C} : \Psi'$, then, for all f and s , if $(fs) \in \Psi'$, we know $f \in \text{dom}(\mathbb{C})$.

Lemma 4.15 (Spec. Extension)

If $\Psi, \Delta \vdash \mathbb{C} : \Psi'$ and $\Psi, \Delta \vdash \{s\} f : \mathbb{C}[f]$, then we have $\Psi, \Delta \vdash \mathbb{C} : \Psi''$, where $\Psi'' = \Psi' \cup \{(f, s)\}$.

The preservation lemma is formalized below.

Lemma 4.16 (Preservation)

If INV0 and INV1 are precise (preciseness is defined in Figure 15), $\Psi, \Delta \vdash \mathbb{W}$ and $(\mathbb{W} \Longrightarrow \mathbb{W}')$, we have $\Psi, \Delta \vdash \mathbb{W}'$.

Below we show the soundness theorem. Recall that preciseness is defined in Figure 15, and the specification (p_i, g_i) for the interrupt handler is defined by Formulae (1) and (2). The soundness theorem captures the partial correctness of programs in the sense that, when we reach the target address of jump or branch instructions, we know the assertion specified in Ψ holds. These assertions corresponds to loop-invariants and pre- and post-conditions for functions at high-level.

Theorem 4.17 (Soundness)

If INV0 and INV1 are precise, $\Psi, \Delta \vdash \mathbb{W}$, and $(h.\text{entry}, (p_i, g_i)) \in \Psi$, then, for any n , there exists \mathbb{W}' such that $\mathbb{W} \Longrightarrow^n \mathbb{W}'$; and, if $\mathbb{W}' = (\mathbb{C}, \mathbb{S}, \mathbb{K}, \text{pc}, \text{tid}, \mathbb{T}, \mathbb{B})$, then

1. if $\mathbb{C}(\text{pc}) = j f$, then there exists (p, g) such that $(f, (p, g)) \in \Psi$ and $p \mathbb{K} \mathbb{S}$ holds;
2. if $\mathbb{C}(\text{pc}) = \text{beq } r_s, r_t, f$ and $\mathbb{S}.\mathbb{R}(r_s) = \mathbb{S}.\mathbb{R}(r_t)$, then there exists (p, g) such that $(f, (p, g)) \in \Psi$ and $p \mathbb{K} \mathbb{S}$ holds;
3. if $\mathbb{C}(\text{pc}) = \text{call } f$, then there exists (p, g) such that $(f, (p, g)) \in \Psi$ and $p(\text{pc} :: \mathbb{K}) \mathbb{S}$ holds;

$$p \triangleq (\text{ie} = 1) \wedge (\text{is} = 0)$$

$$p' \triangleq (\text{ie} = 0) \wedge (\text{is} = 0)$$

$$p_0 \triangleq p$$

$$p_1 \triangleq p \wedge (r_1 = \text{RIGHT}) \wedge (r_2 = \text{LEFT})$$

$$p_2 \triangleq p' \wedge (r_1 = \text{RIGHT}) \wedge (r_2 = \text{LEFT}) \wedge (r_3 = 0) \wedge (\text{INV0} * \text{true})$$

$$p_3 \triangleq p' \wedge (r_1 = \text{RIGHT}) \wedge (r_2 = \text{LEFT}) \wedge (\text{INV0} * \text{true})$$

$$p_4 \triangleq \text{enable}_{\text{iret}}$$

$$\text{NoG} \triangleq \lambda \mathbb{S}, \mathbb{S}'. \text{False}$$

Figure 21. Specifications for the Teeter-Totter Example

4. if $\mathbb{C}(\text{pc}) = \text{ret}$, then there exist pc' , \mathbb{K}' , and (p, g) such that $\mathbb{K} = \text{pc}' :: \mathbb{K}'$, $(\text{pc}', (p, g)) \in \Psi$, and $p \mathbb{K}' \mathbb{S}$ holds.

Proof sketch. The theorem describes both non-stuckness and partial correctness. For non-stuckness, we can simply do induction over n and apply Progress (4.2 and 4.3) and Preservation (4.16) lemmas. We actually prove a stronger version than non-stuckness: $(\Psi, \Delta \vdash _)$ holds over every intermediate world configuration. Therefore, we know $\Psi, \Delta \vdash \mathbb{W}'$ holds. Then the next four bullets about partial correctness can be easily derived by an inversion of the `prog` rule. \square

4.4 The Teeter-Totter Example

With our program logic, we can now certify the Teeter-Totter example shown in Fig. 10. We first define INV0 , the interrupt handler's specification for its local memory.

$$\text{INV0} \triangleq \exists w_l, w_r. ((\text{LEFT} \mapsto w_l) * (\text{RIGHT} \mapsto w_r)) \wedge (w_l + w_r = 100)$$

Then we can get the concrete specification of the interrupt handler, following Formulae (1) and (2). We let INV1 be `emp`, since the non-handler code is sequential.

Specifications are shown in Figure 21. Recall $\text{enable}_{\text{iret}}$ is defined in Figure 18. To simplify our presentation, we present the predicate p in the form of a proposition with free variables referring to components of the state \mathbb{S} . Also, we use m as a shorthand for the proposition $m \mathbb{H}$ when there is no confusion.

If we compare p_1 and p_2 , we will see that the non-handler code cannot access memory at addresses `LEFT` and `RIGHT` without first disabling the interrupt because, from p_1 , we cannot prove that the target addresses `LEFT` and `RIGHT` are in the domain of memory, as required in the instantiation of the `seq` rule for `ld` and `st`. Since the non-handler never returns, we do not really care about the guarantee g for the state transition from the specified point to the return point of the function. Here we simply use `NoG` (see Figure 21) as guarantees.

Certifying the code with respect to the specifications is left to the reader and not shown here.

4.5 The Timer Handler

Here we also briefly explain the specification for the preemptive timer handler shown in Figure 13. The only memory the handler uses is the memory cell at location `CNT`. We define INV0 below.

$$\text{INV0} \triangleq \exists w. (\text{CNT} \mapsto w) \wedge (w \leq 100)$$

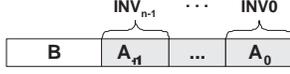
Then we get the specification of the handler (p_i, g_i) by Equations (1) and (2). In g_0 , we use primed variable (e.g., ie' and is') to refer to components in the second state.

4.6 Further Extensions

In AIM, we only support one interrupt in the system, which cannot be interrupted again. It is actually easy to extend the machine to

support multi-level interrupts: we change the `is` bit into a vector of bits `ivec` corresponding to interrupts in service. An interrupt can only be interrupted by other interrupts with higher priorities, which can also be disabled by clearing the `ie` bit. At the end of each interrupt handler, the corresponding in-service bit will be cleared so that interrupts at the same or lower level can be served.

Extension of the program logic to support multi-level interrupts is also straightforward, following the same idea of memory partition. Suppose there are n interrupts in the system, the memory will be partitioned into $n+1$ blocks, as shown below:



where block A_k will be used by the interrupt handler k . To take care of the preemption relations with multiple handlers, we need to change our definition of $\text{Inv}(\text{ie}, \text{is})$ into $\text{Inv}(\text{ie}, \text{ivec})$, which models the switch of memory ownership at the points of `cli`, `sti` and boundaries of interrupt handlers.

Another simplification in our work is the assumption of a global interrupt handler entry. It is easy to extend our machine and program logic to support runtime installation of interrupt handlers. In our machine, we can add a special register and an “install” to update this register. When interrupt comes, we look up the entry point from this register. This extension has almost no effects over our program logic, thanks to our support of modular reasoning. We only need to add a command rule for the “install” instruction to enforce that the new handler’s interface is compatible to the specification (p_i, g_i) .

Also, we do not consider dynamic thread creation in this paper. In our previous work [6], we have shown how to support dynamic thread creation following a similar technique to support dynamic memory allocation in type systems. The technique is fairly orthogonal and can be easily incorporated into this work. Gotsman *et al.* [9] recently showed an extension of concurrent separation logic with dynamic threads and locks. Their technique might be applied here as well to support dynamic creation of block queues.

We will not show the details of supporting multi-level interrupts, dynamic installation of handlers, and dynamic creation of threads and block queues, which are extensions orthogonal to the focus of this paper, *i.e.*, interaction between threads and interrupts.

5. More Examples and Implementations

In this section, we show how to use AIM and the program logic to implement and certify common synchronization primitives.

5.1 Implementations of Locks.

Threads use locks to achieve exclusive access to shared heap. We use Γ , a partial mapping from lock ids to heap predicates, to specify invariants of memory blocks protected by locks.

$$(LockID) \quad l ::= 1$$

$$(LockSpec) \quad \Gamma ::= \{l \rightsquigarrow m\}^*$$

In our implementations, we use memory pointers (label `1`) as lock ids l . Each l points to a memory cell containing a binary flag that records whether the lock has been acquired (flag is 0) or not. The heap used to implement locks and the heap protected by locks are shared by threads in the non-handler code. The invariant $\text{INV}(\Gamma)$ over this part of heap is defined below. We require $\text{INV}_s \Rightarrow \text{INV}(\Gamma) * \text{true}$ (recall that INV_s is a shorthand for $\text{INV}0 * \text{INV}1$).

$$\text{INV}(l, m) \triangleq \exists w. (l \mapsto w) * ((w = 0) \wedge \text{emp} \vee (w = 1) \wedge m) \quad (4)$$

$$\text{INV}(\Gamma) \triangleq \forall_* l \in \text{dom}(\Gamma). \text{INV}(l, \Gamma(l)) \quad (5)$$

```

ACQ_H:  -{(p01, g01)}
        cli
        -{(p02, g02)}
        call ACQ_H_a
        -{(p03, g03)}
        sti
        -{(p04, gid)}
        ret

ACQ_H_a: -{(p11, g11)}
        ld  $r2, 0($r1)      ;; $r2 <- [l]
        movi $r3, 0
        beq  $r2, $r3, gowait ;; ([l] == 0)?
        st  0($r1), $r3     ;; [l] <> 0:
        ret                 ;; [l] <- 0

gowait: -{(p12, g11)}      ;; [l] == 0:
        block $r1           ;; block
        -{(p13, gid)}
        ret

REL_H:  -{(p21, g21)}
        cli
        call REL_H_a
        sti
        ret

REL_H_a: -{(p31, g31)}
        unblock $r1, $r2
        -{(p32, g32)}
        movi $r3, 0
        beq  $r2, $r3, rel_lock
        ret

rel_lock: -{(p33, g33)}
        movi $r2, 1
        st  0($r1), $r2
        -{(p34, gid)}
        ret

```

Figure 22. Hoare-Style Implementation of Lock

where \forall_* is an indexed, finitely iterated separating conjunction, which is defined as:

$$\forall_* x \in S. P(x) \triangleq \begin{cases} \text{emp} & \text{if } S = \emptyset \\ P(x_i) * \forall_* x \in S'. P(x) & \text{if } S = S' \uplus \{x_i\} \end{cases}$$

We first show two block-based implementations, in which we use the lock id as the identifier of the corresponding block queue in \mathbb{B} . Then we show an implementation of spin locks.

The Hoare-style implementation. In Hoare style, the thread gets the lock (and the resource protected by the lock) immediately after it is released from the block queue. The implementation and specifications are shown in Figs. 22 and 23. The precondition for `ACQ_H` is (p_{01}, g_{01}) . The assertion p_{01} requires that r_1 contains a lock id and $\Delta(r_1) = \Gamma(r_1)$. The guarantee g_{01} shows that the function obtains the ownership of $\Gamma(r_1)$ when it returns. Here we use primed variables (*e.g.*, ie' and is') to refer to components in the return state, and use $\text{trash}(\{r_2, r_3\})$ to mean that values of all registers other than r_2 and r_3 are preserved.

Although the meaning of p_{01} and g_{01} is obvious, they can be further simplified when exported to the high-level concurrent programs shown in Figure 3. The high-level specification does not need to refer to `ie` and `is` since they are always 1 and 0 respectively. We do not need to specify stack and trashing of registers,

$$\begin{aligned}
p_0 &\triangleq (is = 0) \wedge \text{enable}_{\text{ret}} \wedge (r_1 \in \text{dom}(\Gamma)) \wedge (\Delta(r_1) = \Gamma(r_1)) \\
p_{01} &\triangleq p_0 \wedge (ie = 1) \\
g_{01} &\triangleq \left\{ \begin{array}{l} \text{emp} \\ \Gamma(r_1) \end{array} \right\} \wedge (ie = ie') \wedge (is = is') \wedge \text{trash}(\{r_2, r_3\}) \\
p_{02} &\triangleq p_0 \wedge (ie = 0) \wedge (\text{INV}_s * \text{true}) \\
g_{02} &\triangleq \left\{ \begin{array}{l} \text{INV}_s \\ \Gamma(r_1) \end{array} \right\} \wedge (ie = 1 - ie') \wedge (is = is') \wedge \text{trash}(\{r_2, r_3\}) \\
p_{03} &\triangleq p_0 \wedge (ie = 0) \wedge (\Gamma(r_1) * \text{INV}_s * \text{true}) \\
g_{03} &\triangleq \left\{ \begin{array}{l} \text{INV}_s \\ \text{emp} \end{array} \right\} \wedge (ie = 1 - ie') \wedge (is = is') \wedge \text{trash}(\{r_2, r_3\}) \\
p_{04} &\triangleq p_0 \wedge (ie = 1) \wedge (\Gamma(r_1) * \text{true}) \\
p_{11} &\triangleq p_0 \wedge (ie = 0) \wedge (\text{INV}_s * \text{true}) \\
g_{11} &\triangleq \left\{ \begin{array}{l} \text{INV}_s \\ \text{INV}_s * \Gamma(r_1) \end{array} \right\} \wedge (ie = ie') \wedge (is = is') \wedge \text{trash}(\{r_2, r_3\}) \\
p_{12} &\triangleq p_0 \wedge (ie = 0) \wedge ([r_1] = 0) \wedge (\text{INV}_s * \text{true}) \\
p_{13} &\triangleq p_0 \wedge (ie = 0) \wedge (\text{INV}_s * \text{true} * \Gamma(r_1)) \\
p_{21} &\triangleq p_0 \wedge (ie = 1) \wedge (\Gamma(r_1) * \text{true}) \\
g_{21} &\triangleq \left\{ \begin{array}{l} \Gamma(r_1) \\ \text{emp} \end{array} \right\} \wedge (ie = ie') \wedge (is = is') \wedge \text{trash}(\{r_2, r_3\}) \\
p_{31} &\triangleq p_0 \wedge (ie = 0) \wedge (\Gamma(r_1) * \text{INV}_s * \text{true}) \\
g_a &\triangleq \left\{ \begin{array}{l} \Gamma(r_1) \\ \text{emp} \end{array} \right\} \quad g_b \triangleq \left\{ \begin{array}{l} r_1 \mapsto _ \\ r_1 \mapsto _ \end{array} \right\} \quad \text{hid} \triangleq \left\{ \begin{array}{l} \text{emp} \\ \text{emp} \end{array} \right\} \\
g_{31} &\triangleq (g_a \vee g_b) \wedge (ie = ie') \wedge (is = is') \wedge \text{trash}(\{r_2, r_3\}) \\
p_{32} &\triangleq p_0 \wedge (ie = 0) \wedge (((r_2 = 0) \wedge (\Gamma(r_1) * \text{INV}_s)) \vee (r_2 \neq 0) \wedge \text{INV}_s) * \text{true}) \\
g_{32} &\triangleq ((r_2 = 0 \wedge g_b) \vee (r_2 \neq 0 \wedge \text{hid})) \\
&\quad \wedge (ie = ie') \wedge (is = is') \wedge \text{trash}(\{r_2, r_3\}) \\
p_{33} &\triangleq p_0 \wedge (ie = 0) \wedge (\Gamma(r_1) * \text{INV}_s * \text{true}) \\
g_{33} &\triangleq g_b \wedge (ie = ie') \wedge (is = is') \wedge \text{trash}(\{r_2, r_3\}) \\
p_{34} &\triangleq p_0 \wedge (ie = 0) \wedge (\text{INV}_s * \text{true})
\end{aligned}$$

Figure 23. Specifications of Hoare-Style Lock

```

ACQ_M:  -{(p11, g11)}
        movi    $r3, 0
acq_loop: -{(p12, g12)}
        cli
        ld     $r2, 0($r1)      ;; $r2 <- [I]
        beq   $r2, $r3, gowait  ;; ([I] == 0)?
        st   0($r1), $r3       ;; [I] <> 0:
        j    acq_done          ;; [I] <- 0
gowait:  -{(p13, g13)}
        block $r1
        -{(p13, g13)}
        sti
        j     acq_loop
acq_done: -{(p14, g14)}
        sti
        ret

REL_M:  -{(p21, g21)}
        cli
        -{(p22, g22)}
        unblock $r1, $r2
        -{(p23, g22)}
        -{(p22, g22)}
        movi    $r2, 1
        st   0($r1), $r2
        sti
        ret

```

Figure 24. Mesa-Style Implementation of Locks

which can be inferred from the calling convention. We can also hide Δ , since `block` and `unblock` are not visible from the high level. So the specification exported to the high level would be:

$$\begin{aligned}
p_{01} &\triangleq r_1 \in \text{dom}(\Gamma) \\
g_{01} &\triangleq \left\{ \begin{array}{l} \text{emp} \\ \Gamma(r_1) \end{array} \right\}
\end{aligned}$$

We also show some intermediate specifications used during verification. Comparing (p_{01}, g_{01}) and (p_{11}, g_{11}) , we can see that (p_{01}, g_{01}) hides INV_s and the implementation details of the lock from the client code. Readers can also compare p_{12} and p_{13} and see how the `BLK` rule is applied.

Functions `REL_h_a` and `REL_h` releases the lock with the interrupt disabled and enabled, respectively. They are specified by (p_{21}, g_{21}) and (p_{31}, g_{31}) . Depending on whether there are threads waiting for the lock, the current thread may either transfer the ownership of $\Gamma(r_1)$ to a waiting thread or simply set the lock to be available, as specified in g_{31} , but these details are hidden in g_{21} .

The Mesa-style implementation. Figure 24 shows the Mesa-style implementation of locks. The specifications are shown in Figure 25. In the `ACQ_M` function, the thread needs to start another round of loop to test the availability of the lock after `block`. The `REL_M` function always sets the lock to be available, even if it releases a waiting thread. Specifications are the same with Hoare style except that the assertion p_0 requires $\Delta(r_1) = \text{emp}$, which implies the Mesa-style semantics of `block` and `unblock`.

Spin Locks An implementation of spin locks and its specifications are shown in Figure 26. The specifications (p_{11}, g_{11}) and (p_{21}, g_{21}) describes the interface of lock acquire/release. They look very similar to specifications for block-based implementations: “acquire” gets the ownership of the extra resource $\Gamma(r_1)$ protected by the lock in r_1 , while “release” loses the ownership so that the client can no longer use the resource after calling “release”.

$$\begin{aligned}
p_0 &\triangleq (\text{is} = 0) \wedge \text{enable}_{\text{ret}} \wedge (r_1 \in \text{dom}(\Gamma)) \wedge (\Delta(r_1) = \text{emp}) \\
p_{11} &\triangleq p_0 \wedge (\text{ie} = 1) \\
g_{11} &\triangleq \left\{ \begin{array}{l} \text{emp} \\ \Gamma(r_1) \end{array} \right\} \wedge (\text{ie} = \text{ie}') \wedge (\text{is} = \text{is}') \wedge \text{trash}(\{r_2, r_3\}) \\
p_{12} &\triangleq p_{11} \wedge (r_3 = 0) \\
p_{13} &\triangleq p_0 \wedge (\text{ie} = 0) \wedge (\text{INV}_s * \text{true}) \\
g_{13} &\triangleq \left\{ \begin{array}{l} \text{INV}_s \\ \Gamma(r_1) \end{array} \right\} \wedge (\text{ie} = 1 - \text{ie}') \wedge (\text{is} = \text{is}') \wedge \text{trash}(\{r_2, r_3\}) \\
p_{14} &\triangleq p_0 \wedge (\text{ie} = 0) \wedge (\Gamma(r_1) * \text{INV}_s * \text{true}) \\
g_{14} &\triangleq \left\{ \begin{array}{l} \text{INV}_s \\ \text{emp} \end{array} \right\} \wedge (\text{ie} = 1 - \text{ie}') \wedge (\text{is} = \text{is}') \wedge \text{trash}(\{r_2, r_3\}) \\
p_{21} &\triangleq p_0 \wedge (\text{ie} = 1) \wedge (\Gamma(r_1) * \text{true}) \\
g_{21} &\triangleq \left\{ \begin{array}{l} \Gamma(r_1) \\ \text{emp} \end{array} \right\} \wedge (\text{ie} = \text{ie}') \wedge (\text{is} = \text{is}') \wedge \text{trash}(\{r_2\}) \\
p_{22} &\triangleq p_0 \wedge (\text{ie} = 0) \wedge (\Gamma(r_1) * \text{INV}_s * \text{true}) \\
g_{22} &\triangleq \left\{ \begin{array}{l} \Gamma(r_1) * \text{INV}_s \\ \text{emp} \end{array} \right\} \wedge (\text{ie} = 1 - \text{ie}') \wedge (\text{is} = \text{is}') \wedge \text{trash}(\{r_2\}) \\
p_{23} &\triangleq p_0 \wedge (\text{ie} = 0) \wedge \\
&\quad ((([r_1] = 0) \wedge (\Gamma(r_1) * \text{INV}_s) \vee ([r_1] \neq 0) \wedge (\Gamma(r_1) * \text{INV}_s)) * \text{true})
\end{aligned}$$

Figure 25. Specification of Mesa-Style Implementation of Locks

These specifications also hide the implementation details (*e.g.*, the lock name l is a pointer pointing to a binary value) from the client code.

5.2 Implementations of Condition Variables

Now we show implementations of Mesa style [14], Hoare style [11] and Brinch Hansen style [2] condition variables. Below we use Υ , a partial mapping from condition variables cv to heap predicates m , to specify the conditions associated with condition variables.

$$\begin{aligned}
(\text{CondVar}) \quad cv &::= n \quad (\text{nat nums}) \\
(\text{CVSpec}) \quad \Upsilon &::= \{cv \rightsquigarrow m\}^*
\end{aligned}$$

In our implementation, we let cv be an identifier pointing to a block queue in \mathbb{B} . A lock l needs to be associated with cv to guarantee exclusive access of the resource specified by $\Gamma(l)$. The difference between $\Gamma(l)$ and $\Upsilon(cv)$ is that $\Gamma(l)$ specifies the basic well-formedness of the resource (*e.g.*, a well-formed queue), while $\Upsilon(cv)$ specifies an extra condition (*e.g.*, the queue is not empty).

Hoare style and Brinch Hansen style. The implementation and specifications of Hoare-style and Brinch Hansen style are shown in Figs. 27, 28 and 29. The precondition for `WAIT_H` is (p_{11}, g_{11}) . As p_{11} shows, r_1 contains a Hoare-style lock in the sense that $\Delta(r_1) = \Gamma(r_1)$. The register r_2 contains the condition variable with specification $\Upsilon(r_2)$. For Hoare-style, we require $\Delta(r_2) = \Gamma(r_1) \wedge (\Upsilon(r_2) * \text{true})$. Therefore, when the blocked thread is released, it gets the resource protected by the lock with the extra knowledge that the condition associated with the condition variable holds. Here the condition $\Upsilon(r_2)$ does not have to specify the whole resource protected by the lock, therefore we use $\Upsilon(r_2) * \text{true}$. Before calling `WAIT_H`, p_{11} requires that the lock must have been acquired, thus we have the ownership $\Gamma(r_1)$. The condition $\Upsilon(r_2)$ needs to be false. This is not an essential requirement, but we use

```

;; acquire(l): $r1 contains l
spin_acq:  -{(p11, g11)}
           movi $r2, 1
spin_loop: -{(p12, g11)}
           cli
           ld  $r3, 0(r1)
           beq $r2, $r3, spin_set
           sti
           j   spin_loop
spin_set:  -{(p13, g13)}
           movi $r2, 0
           st  0($r1), $r2
           sti
           ret
;; release(l): $r1 contains l
spin_rel:  -{(p21, g21)}
           movi $r2, 1
           cli
           st  0($r1), $r2
           sti
           ret

```

$$\begin{aligned}
p &\triangleq (\text{is} = 0) \wedge \text{enable}_{\text{ret}} \wedge (r_1 \in \text{dom}(\Gamma)) \\
p_{11} &\triangleq p \wedge (\text{ie} = 1) \\
g_{11} &\triangleq \left\{ \begin{array}{l} \text{emp} \\ \Gamma(r_1) \end{array} \right\} \wedge (\text{ie} = \text{ie}') \wedge (\text{is} = \text{is}') \wedge \text{trash}(\{r_2, r_3\}) \\
p_{12} &\triangleq p_{11} \wedge (r_2 = 1) \\
p_{13} &\triangleq p \wedge (\text{ie} = 0) \wedge (\text{INV}_s * \text{true}) \\
g_{13} &\triangleq \left\{ \begin{array}{l} \text{INV}_s \\ \Gamma(r_1) \end{array} \right\} \wedge (\text{ie} = 1 - \text{ie}') \wedge (\text{is} = \text{is}') \wedge \text{trash}(\{r_2\}) \\
p_{21} &\triangleq p \wedge (\text{ie} = 1) \wedge (\Gamma(r_1) * \text{true}) \\
g_{21} &\triangleq \left\{ \begin{array}{l} \Gamma(r_1) \\ \text{emp} \end{array} \right\} \wedge (\text{ie} = \text{ie}') \wedge (\text{is} = \text{is}') \wedge \text{trash}(\{r_2\})
\end{aligned}$$

Figure 26. A Spin Lock

it to prevent waiting without testing the condition. The guarantee g_{11} says that, when `WAIT_H` returns, the current thread still owns the lock (and $\Gamma(r_1)$) and it also knows the condition specified in Υ holds. The precondition for `SIGNAL_H` is (p_{21}, g_{21}) . `SIGNAL_H` requires the thread owns the lock and the condition $\Upsilon(r_2)$ holds at the beginning. When it returns, the thread still owns the lock, but the condition may no longer hold. Figure 28 also shows important intermediate specifications we use during verification.

Brinch Hansen style condition variables is similar to Hoare-style. The wait function and its specifications are the same as `WAIT_H`. The signal function `SIGNAL_BH`, which is omitted here, has specification (p_{31}, g_{31}) defined in Figure 28. Here p_{31} is the same as p_{21} for `SIGNAL_H`. The definition of g_{31} shows the difference between Hoare style and Brinch Hansen style: the lock is released when `SIGNAL_BH` returns. Therefore, calling the `SIGNAL_BH` function must be the last command in the critical region.

Mesa-style. Figure 30 shows Mesa-style condition variables, with specifications shown in Figure 31. `WAIT_M` is specified by (p_{11}, g_{11}) . The assertion p_{11} is similar to the precondition for Hoare-style, except that we require $\Delta(r_2) = \text{emp}$. Therefore, as g_{11} shows, the current thread has no idea about the validity of the condition when it returns.

```

WAIT_H:  --{(p11, g11)}           ;; wait(l, cv)
cli
mov     $r4, $r2
--{(p12, g12)}
call   REL_H_a
--{(p13, g13)}
block  $r4
--{(p14, g14)}
sti
ret

SIGNAL_H: --{(p21, g21)}           ;; signal(l, cv)
cli
--{(p22, g22)}
unblock $r2, $r3
--{(p23, g23)}
movi   $r4, 0
beq    $r3, $r4, sig_done
--{(p24, g24)}
block  $r1
sig_done: --{(p25, g25)}
sti
ret

SIGNAL_BH: --{(p31, g31)}           ;; signal(l, cv)
cli
--{(p32, g32)}
unblock $r2, $r3           ;; $r2 contains cv
--{(p33, g33)}
movi   $r4, 0
beq    $r3, $r4, sig_cont
--{(p34, g34)}
j      sig_done
sig_cont: --{(p35, g35)}
--{(p36, g36)}
call   REL_0               ;; $r1 contains l
sig_done: --{(p34, g34)}
sti
ret

```

Figure 27. Implementation of CV - Hoare Style

SIGNAL_M is specified by (p_{21}, g_{21}) . The assertion hid is defined in Figure 23, which means the function has no effects over data heap. From g_{21} we can see that, if we hide the details of releasing a blocked thread, the signal function in Mesa style is just like a skip command. We do not require the current thread to own the lock l before it calls SIGNAL_M, since it has no effects over data heap.

5.3 Other Implementation Details

As shown in Figure 3, we have certified the preemptive thread implementations and libraries extracted from our simplified OS kernel, which is implemented in 16-bit x86 assembly code and works in real-mode. At the lowest level, we have concrete implementations of the scheduler and block/unblock primitives. Thread queues are implemented as a doubly linked list containing thread control blocks. The synchronization primitives in the middle level also have been implemented in x86, which call the underlying primitives. The timer handler simply saves the context and calls the scheduler. The yield function just wraps the scheduler by disabling the interrupt at the beginning and enabling it at the end.

Since the code are at different abstraction levels, we certify them using different program logics, following the technique proposed

$$\text{Cond}(r, r') \triangleq \Gamma(r) \wedge (\Upsilon(r') * \text{true})$$

$$\overline{\text{Cond}}(r, r') \triangleq \Gamma(r) \wedge \neg(\Upsilon(r') * \text{true})$$

$$p(r, r') \triangleq (\text{is} = 0) \wedge \text{enable}_{\text{ret}} \wedge \exists l, cv, m, m'. (r = l) \wedge (r' = cv) \wedge (\Gamma(l) = m) \wedge (\Delta(l) = m) \wedge (\Upsilon(cv) = m') \wedge (\Delta(cv) = \text{Cond}(r, r'))$$

$$p_{11} \triangleq p(r_1, r_2) \wedge (\text{ie} = 1) \wedge (\overline{\text{Cond}}(r_1, r_2) * \text{true})$$

$$g_{11} \triangleq \left\{ \begin{array}{l} \overline{\text{Cond}}(r_1, r_2) \\ \text{Cond}(r_1, r_2) \end{array} \right\} \wedge (\text{ie} = \text{ie}') \wedge (\text{is} = \text{is}') \wedge \text{trash}(\{r_2, r_3, r_4\})$$

$$p_{12} \triangleq p(r_1, r_4) \wedge (\text{ie} = 0) \wedge (\overline{\text{Cond}}(r_1, r_4) * \text{INV}_s * \text{true})$$

$$g_{12} \triangleq \left\{ \begin{array}{l} \overline{\text{Cond}}(r_1, r_4) * \text{INV}_s \\ \text{Cond}(r_1, r_4) \end{array} \right\} \wedge (\text{ie} = 1 - \text{ie}') \wedge (\text{is} = \text{is}') \wedge \text{trash}(\{r_2, r_3, r_4\})$$

$$p_{13} \triangleq p(r_1, r_4) \wedge (\text{ie} = 0) \wedge (\text{INV}_s * \text{true})$$

$$g_{13} \triangleq \left\{ \begin{array}{l} \text{INV}_s \\ \text{Cond}(r_1, r_4) \end{array} \right\} \wedge (\text{ie} = 1 - \text{ie}') \wedge (\text{is} = \text{is}') \wedge \text{trash}(\{r_2, r_3, r_4\})$$

$$p_{14} \triangleq p(r_1, r_4) \wedge (\text{ie} = 0) \wedge (\text{Cond}(r_1, r_4) * \text{INV}_s * \text{true})$$

$$g_{14} \triangleq \left\{ \begin{array}{l} \text{INV}_s \\ \text{emp} \end{array} \right\} \wedge (\text{ie} = 1 - \text{ie}') \wedge (\text{is} = \text{is}') \wedge \text{trash}(\{r_2, r_3, r_4\})$$

$$p_{21} \triangleq p(r_1, r_2) \wedge (\text{ie} = 1) \wedge (\text{Cond}(r_1, r_2) * \text{true})$$

$$g_{21} \triangleq \left\{ \begin{array}{l} \text{Cond}(r_1, r_2) \\ \Gamma(r_1) \end{array} \right\} \wedge (\text{ie} = \text{ie}') \wedge (\text{is} = \text{is}') \wedge \text{trash}(\{r_2, r_3, r_4\})$$

$$p_{22} \triangleq p(r_1, r_2) \wedge (\text{ie} = 0) \wedge ((\Gamma(r_1) \wedge \text{Cond}) * \text{INV}_s)$$

$$g_a \triangleq \left\{ \begin{array}{l} \text{Cond}(r_1, r_2) * \text{INV}_s \\ \Gamma(r_1) \end{array} \right\} \quad g_b \triangleq \left\{ \begin{array}{l} \text{INV}_s \\ \Gamma(r_1) \end{array} \right\}$$

$$g_{22} \triangleq g_a \wedge (\text{ie} = 1 - \text{ie}') \wedge (\text{is} = \text{is}') \wedge \text{trash}(\{r_2, r_3, r_4\})$$

$$p_{23} \triangleq p(r_1, r_2) \wedge (\text{ie} = 0) \wedge ((r_3 = 0) \wedge (\text{Cond}(r_1, r_2) * \text{INV}_s * \text{true}) \vee (r_3 \neq 0) \wedge (\text{INV}_s * \text{true}))$$

$$g_{23} \triangleq (r_3 = 0 \wedge g_a \vee r_3 \neq 0 \wedge g_b) \wedge (\text{ie} = 1 - \text{ie}') \wedge (\text{is} = \text{is}') \wedge \text{trash}(\{r_2, r_3, r_4\})$$

$$p_{24} \triangleq p(r_1, r_2) \wedge (\text{ie} = 0) \wedge (\text{INV}_s * \text{true})$$

$$g_{24} \triangleq g_b \wedge (\text{ie} = 1 - \text{ie}') \wedge (\text{is} = \text{is}') \wedge \text{trash}(\{r_2, r_3, r_4\})$$

$$p_{25} \triangleq p(r_1, r_2) \wedge (\text{ie} = 0) \wedge (\text{INV}_s * \Gamma(r_1) * \text{true})$$

$$g_{25} \triangleq \left\{ \begin{array}{l} \text{INV}_s \\ \text{emp} \end{array} \right\} \wedge (\text{ie} = 1 - \text{ie}') \wedge (\text{is} = \text{is}') \wedge \text{trash}(\{r_2, r_3, r_4\})$$

Figure 28. Spec. of CV - Hoare Style

$$\begin{aligned}
P_{31} &\triangleq p(r_1, r_2) \wedge (\text{ie} = 1) \wedge \Gamma(r_1) \wedge (\text{Cond}(r_1, r_2) * \text{true}) \\
g_{31} &\triangleq \left\{ \begin{array}{l} \text{Cond}(r_1, r_2) \\ \text{emp} \end{array} \right\} \wedge (\text{ie} = \text{ie}') \wedge (\text{is} = \text{is}') \wedge \text{trash}(\{r_2, r_3, r_4\}) \\
P_{32} &\triangleq p(r_1, r_2) \wedge (\text{ie} = 0) \wedge (\text{Cond}(r_1, r_2) * \text{INV}_s * \text{true}) \\
g_a &\triangleq \left\{ \begin{array}{l} \text{Cond}(r_1, r_2) * \text{INV}_s \\ \text{emp} \end{array} \right\} & g_b &\triangleq \left\{ \begin{array}{l} \text{INV}_s \\ \text{emp} \end{array} \right\} \\
g_{32} &\triangleq g_a \wedge (\text{ie} = 1 - \text{ie}') \wedge (\text{is} = \text{is}') \wedge \text{trash}(\{r_2, r_3, r_4\}) \\
P_{33} &\triangleq p(r_1, r_2) \wedge (\text{ie} = 0) \wedge \\
&\quad ((r_3 = 0) \wedge (\text{Cond}(r_1, r_2) * \text{INV}_s) \vee (r_3 \neq 0) \wedge \text{INV}_s) * \text{true} \\
g_{33} &\triangleq (r_3 = 0 \wedge g_a \vee r_3 \neq 0 \wedge g_b) \\
&\quad \wedge (\text{ie} = 1 - \text{ie}') \wedge (\text{is} = \text{is}') \wedge \text{trash}(\{r_2, r_3, r_4\}) \\
p'(r) &\triangleq (\text{is} = 0) \wedge \exists l, m. (r = l) \wedge (\Gamma(l) = m) \wedge (\Delta(l) = m) \\
P_{34} &\triangleq p'(r_1) \wedge (\text{ie} = 0) \wedge (\text{INV}_s * \text{true}) \\
g_{34} &\triangleq g_b \wedge (\text{ie} = 1 - \text{ie}') \wedge (\text{is} = \text{is}') \wedge \text{trash}(\{r_2, r_3, r_4\}) \\
P_{35} &\triangleq p(r_1, r_2) \wedge (\text{ie} = 0) \wedge (\text{Cond}(r_1, r_2) * \text{INV}_s * \text{true}) \\
g_{35} &\triangleq g_a \wedge (\text{ie} = 1 - \text{ie}') \wedge (\text{is} = \text{is}') \wedge \text{trash}(\{r_2, r_3, r_4\}) \\
P_{36} &\triangleq p'(r_1) \wedge (\text{ie} = 0) \wedge (\Gamma(r_1) * \text{INV}_s * \text{true}) \\
g_{36} &\triangleq \left\{ \begin{array}{l} \Gamma(r_1) * \text{INV}_s \\ \text{emp} \end{array} \right\} \\
&\quad \wedge (\text{ie} = 1 - \text{ie}') \wedge (\text{is} = \text{is}') \wedge \text{trash}(\{r_2, r_3, r_4\})
\end{aligned}$$

Figure 29. Spec. of CV - Brinch Hansen Style

```

WAIT_M:  -{p11, g11}           ;; wait(l, cv)
cli
mov     $r4, $r2
-{p12, g12}
call   REL_H_a
-{p13, g13}
block  $r4
-{p14, g13}
sti
-{p15, g15}
call   ACQ_H
-{p16, gid}
ret

SIGNAL_M: -{p21, g21}       ;; signal(cv)
cli
-{p22, g22}
unblock $r1, $r2
-{p22, g22}
sti
ret

```

Figure 30. Implementation of Condition Variable - Mesa Style

$$\begin{aligned}
p(r, r') &\triangleq (\text{is} = 0) \wedge \text{enable}_{\text{ret}} \wedge \\
&\quad \exists l, cv, m, m'. (r = l) \wedge (r' = cv) \wedge (\Gamma(l) = m) \wedge (\Delta(l) = (m, 1)) \\
&\quad \wedge (\Upsilon(cv) = m') \wedge (\Delta(cv) = \text{emp}) \\
P_{11} &\triangleq p(r_1, r_2) \wedge (\text{ie} = 1) \wedge (\overline{\text{Cond}}(r_1, r_2) * \text{true}) \\
g_{11} &\triangleq \left\{ \begin{array}{l} \overline{\text{Cond}}(r_1, r_2) \\ \Gamma(r_1) \end{array} \right\} \wedge (\text{ie} = \text{ie}') \wedge (\text{is} = \text{is}') \wedge \text{trash}(\{r_2, r_3, r_4\}) \\
P_{12} &\triangleq p(r_1, r_4) \wedge (\text{ie} = 0) \wedge (\overline{\text{Cond}}(r_1, r_4) * \text{INV}_s * \text{true}) \\
g_{12} &\triangleq \left\{ \begin{array}{l} \overline{\text{Cond}}(r_1, r_4) * \text{INV}_s \\ \Gamma(r_1) \end{array} \right\} \\
&\quad \wedge (\text{ie} = 1 - \text{ie}') \wedge (\text{is} = \text{is}') \wedge \text{trash}(\{r_2, r_3, r_4\}) \\
P_{13} &\triangleq p(r_1, r_4) \wedge (\text{ie} = 0) \wedge (\text{INV}_s * \text{true}) \\
g_{13} &\triangleq \left\{ \begin{array}{l} \text{INV}_s \\ \Gamma(r_1) \end{array} \right\} \\
&\quad \wedge (\text{ie} = 1 - \text{ie}') \wedge (\text{is} = \text{is}') \wedge \text{trash}(\{r_2, r_3, r_4\}) \\
P_{14} &\triangleq p(r_1, r_4) \wedge (\text{ie} = 0) \wedge (\text{INV}_s * \text{true}) \\
P_{15} &\triangleq p(r_1, r_4) \wedge (\text{ie} = 1) \\
g_{15} &\triangleq \left\{ \begin{array}{l} \text{emp} \\ \Gamma(r_1) \end{array} \right\} \wedge (\text{ie} = \text{ie}') \wedge (\text{is} = \text{is}') \wedge \text{trash}(\{r_2, r_3, r_4\}) \\
P_{16} &\triangleq p(r_1, r_4) \wedge (\text{ie} = 1) \wedge \Gamma(r_1) \\
p'(r) &\triangleq (\text{is} = 0) \wedge \exists cv, m. (r = cv) \wedge (\Delta(cv) = \text{emp}) \\
P_{21} &\triangleq p'(r_1) \wedge (\text{ie} = 1) \\
g_{21} &\triangleq \text{hid} \wedge (\text{ie} = \text{ie}') \wedge (\text{is} = \text{is}') \wedge \text{trash}(\{r_2\}) \\
P_{22} &\triangleq p'(r_1) \wedge (\text{ie} = 0) \wedge (\text{INV}_s * \text{true}) \\
g_{22} &\triangleq \left\{ \begin{array}{l} \text{INV}_s \\ \text{emp} \end{array} \right\} \wedge (\text{ie} = 1 - \text{ie}') \wedge (\text{is} = \text{is}') \wedge \text{trash}(\{r_2\})
\end{aligned}$$

Figure 31. Specifications of Condition Variable - Mesa Style

by Feng *et al.* [5]. We use SCAP [7] to certify the low-level implementations of thread primitives, including de-queue/en-queue functions and context-switching code, which are all treated as sequential code. SCAP can be viewed as a specialization of our logic for AIM, assuming *ie* is always 0 and prohibiting the execution of *cli* and *sti*. The operational semantics for *switch*, *block* and *unblock* instructions in AIM are used as specifications for the concrete implementations at the low level, which also refers to a concrete specification about the data structure of thread control blocks and thread queues.

To certify the library code in the middle level, we adapted our program logic to x86, and proved its soundness in Coq. Instead of implementing an abstract machine like AIM and then compiling AIM code to real x86 code, we simply replace these primitives with function calls to the low-level implementations. However, we still need to define a mapping from the abstract thread queues to their concrete representation in memory. Since code at this level does not touch thread queues, this mapping is always preserved

We link the certified code at different levels in an OCAP-like framework [5]. The basic idea is based on the observation that the low-level code only manipulates TCBS and queues and does not touch data used at high-level, while the high-level code does not

access queues. Therefore, the invariant at each level is preserved by the other side, and the safety property certified at each level still holds when all the code are linked together.

Linking of the thread library code at the middle level with the high-level concurrent programs can be done in a similar way. If we assume $ie = 1$ and $is = 0$, and prohibit the code from executing `cli` and `sti`, we can derive a specialized logic from the logic for AIM. Such a logic can be applied to certify high-level concurrent code. We can also expose our specifications of synchronization libraries to the high level code, which would treat them as primitive instructions. We will leave this as future work.

6. Related Work and Conclusions

Regehr and Cooper [20] showed how to translate interrupt-driven programs to thread-based programs. However, their technique cannot be directly applied for our goal to build certified OS kernel. First, proof of the correctness of the translation is non-trivial and has not been formalized. As Regehr and Cooper pointed out, the proof requires a formal semantics of interrupts. Our work actually provides such formal semantics. Second, their translation requires higher-level language constructs such as locks, while we certify the implementation of locks based on our AIM.

Suenaga and Kobayashi [22] presented a type system to guarantee deadlock-freedom in a concurrent calculus with interrupts. Their calculus is an ML-style language with built-in support of threads, locks and interrupts. Our AIM is at a lower abstraction level than theirs with no built-in locks. Also, their type system is designed mainly for preventing deadlocks with automatic type inference, while our program logic supports verification of general safety properties, including partial correctness.

Palsberg and Ma [18] proposed a calculus of interrupt driven systems, which has multi-level interrupts but no threads. Instead of a general program logic like ours, they proposed a type system to guarantee an upper bound of stack space. DeLine and Fähndrich [4] showed how to enforce protocols with regard to interrupts levels as an application of Vault's type system. However, it is not clear how to use the type system for general properties of interrupts.

Bevier [1] showed how to formally certify Kit, an OS kernel implemented in machine code. Gargano *et al.* [8] showed a framework for a certified OS kernel in the Verisoft project. Ni *et al.* [16] certified a non-preemptive thread implementation. In all these cases, implementations of kernels or thread libraries are all sequential. They cannot be interrupted and there is no preemptive concurrency.

In this paper we have presented a new Hoare-style framework for certifying low-level programs involving both interrupts and concurrency. Following Separation Logic, we formalized the interaction among threads and interrupt handlers in terms of memory ownership transfer. Instead of using the operational semantics of `cli`, `sti` and thread primitives, our program logic formulates their local effects over the current thread, as shown in Figure 17, which is the key for our logic to achieve modular verification. We have also certified various lock and condition-variable primitives; our specifications are both abstract (hiding implementation details) and precise (capturing the semantic difference among these variations).

References

- [1] W. R. Bevier. Kit: A study in operating system verification. *IEEE Trans. Softw. Eng.*, 15(11):1382–1396, 1989.
- [2] P. Brinch Hansen. The programming language concurrent pascal. *IEEE Trans. Software Eng.*, 1(2):199–207, 1975.
- [3] S. Brookes. A semantics for concurrent separation logic. In *CONCUR'04*, volume 3170 of *LNCS*, pages 16–34, 2004.
- [4] R. DeLine and M. Fähndrich. Enforcing high-level protocols in low-level software. In *PLDI'01*, pages 59–69, 2001.
- [5] X. Feng, Z. Ni, Z. Shao, and Y. Guo. An open framework for foundational proof-carrying code. In *TLDI'07*, pages 67–78, 2007.
- [6] X. Feng and Z. Shao. Modular verification of concurrent assembly code with dynamic thread creation and termination. In *Proc. ICFP'05*, pages 254–267, 2005.
- [7] X. Feng, Z. Shao, A. Vaynberg, S. Xiang, and Z. Ni. Modular verification of assembly code with stack-based control abstractions. In *PLDI'06*, pages 401–414, June 2006.
- [8] M. Gargano, M. A. Hillebrand, D. Leinenbach, and W. J. Paul. On the correctness of operating system kernels. In *TPHOLS'05*, 2005.
- [9] A. Gotsman, J. Berdine, B. Cook, N. Rinetzky, and M. Sagiv. Local reasoning for storable locks and threads. In *Proceedings of The Fifth ASIAN Symposium on Programming Languages and Systems (APLAS'07)*, page to appear, Nov. 2007.
- [10] C. A. R. Hoare. Towards a theory of parallel programming. In *Operating Systems Techniques*, pages 61–71. Academic Press, 1972.
- [11] C. A. R. Hoare. Monitors: An operating system structuring concept. *Communications of the ACM*, 17(10):549–557, Oct. 1974.
- [12] G. C. Hunt and J. R. Larus. Singularity design motivation. Technical Report MSR-TR-2004-105, Microsoft Corporation, December 2004.
- [13] S. S. Ishtiaq and P. W. O'Hearn. BI as an assertion language for mutable data structures. In *POPL'01*, pages 14–26, 2001.
- [14] B. W. Lampson and D. D. Redell. Experience with processes and monitors in Mesa. *Commun. ACM*, 23(2):105–117, 1980.
- [15] G. Morrisett, D. Walker, K. Crary, and N. Glew. From System F to typed assembly language. In *POPL'98*, pages 85–97, 1998.
- [16] Z. Ni, D. Yu, and Z. Shao. Using XCAP to certify realistic systems code: Machine context management. In *TPHOLS'07*, 2007.
- [17] P. W. O'Hearn. Resources, concurrency and local reasoning. In *CONCUR'04*, volume 3170 of *LNCS*, pages 49–67, 2004.
- [18] J. Palsberg and D. Ma. A typed interrupt calculus. In *FTRTFT'02*, pages 291–310, London, UK, 2002. Springer-Verlag.
- [19] W. Paul, M. Broy, and T. In der Rieden. The verisoft xt project. URL: <http://www.verisoft.de>, 2007.
- [20] J. Regehr and N. Cooper. Interrupt verification via thread verification. *Electron. Notes Theor. Comput. Sci.*, 174(9), 2007.
- [21] J. C. Reynolds. Separation logic: A logic for shared mutable data structures. In *Proc. LICS'02*, pages 55–74, July 2002.
- [22] K. Suenaga and N. Kobayashi. Type based analysis of deadlock for a concurrent calculus with interrupts. In *ESOP'07*, March 2007.
- [23] H. Tuch, G. Klein, and G. Heiser. OS verification — now! In *Proc. HotOS-X*, June 2005.
- [24] V. Vafeiadis and M. Parkinson. A marriage of rely/guarantee and separation logic. In *CONCUR'07*, page to appear, 2007.