

Certifying Low-Level Programs with Hardware Interrupts and Preemptive Threads

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

[†]Yale University, New Haven, CT 06520-8285, U.S.A.

[‡]Toyota Technological Institute at Chicago, Chicago, IL 60637, U.S.A.

[§]Tsinghua University, Beijing 100084, China

^{*}University of Science and Technology of China, Hefei, Anhui 230026, China

feng@tti-c.org shao@cs.yale.edu dongyuan@tsinghua.edu.cn guoyu@mail.ustc.edu.cn

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.

Categories and Subject Descriptors F.3.1 [Logics and Meanings of Programs]: Specifying and Verifying and Reasoning about Programs; D.2.4 [Software Engineering]: Software/Program Verification — Correctness proofs, formal methods, reliability; D.3.1 [Programming Languages]: Formal Definitions and Theory — Semantics; D.4.5 [Operating Systems]: Reliability — Verification

General Terms Languages, Reliability, Security, Verification

Keywords Hardware Interrupts, Preemptive Threads, Certified System Software, Concurrency, Separation Logic

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

PLDI’08, June 7–13, 2008, Tucson, Arizona, USA.

Copyright © 2008 ACM 978-1-59593-860-2/08/06...\$5.00

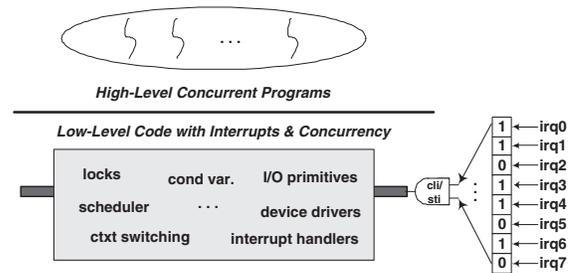


Figure 1. “High-Level” vs. “Low-Level” System Programs

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—including Verisoft/XT (Gargano et al. 2005; Paul et al. 2007), L4.verified (Tuch et al. 2005), and Singularity (Hunt and Larus 2004)—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 statements about system configurations and behaviors (Hunt and Larus 2004).

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 preemptible uniprocessor OS kernel into two layers. At the “higher” abstraction level, we have threads that follow the standard concurrent programming model (Hoare 1972): 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 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 (Hoare 1969), typed assembly language (Morrisset et al.

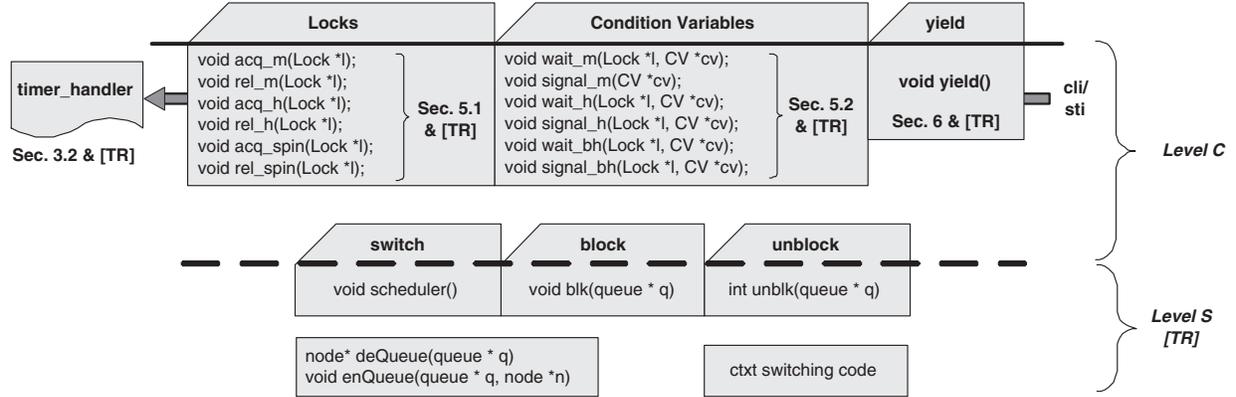


Figure 3. Structure of Our Certified Preemptive Thread Implementation

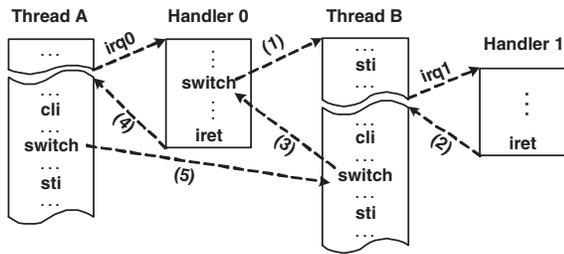


Figure 2. Interaction between Threads and Interrupts

1998), concurrent separation logic (O’Hearn 2004; Brookes 2004), and its assume-guarantee variant (Feng et al. 2007a; Vafeiadis and Parkinson 2007)) 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 (Regehr and Cooperider 2007).
- *Subtle intertwining between interrupts and threads.* In Fig. 2, thread A is interrupted by the interrupt request `irq0`. 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 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 manipulate 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 is linked together using an OCAP-style framework (Feng et al. 2007b).
- 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 Fig. 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 interrupts explicitly and implements interrupt handlers but abstracts away the implementation of threads. Code at Level S is sequential (always executed with interrupts disabled); functions that need to know the concrete representations of thread control blocks (TCBs) and

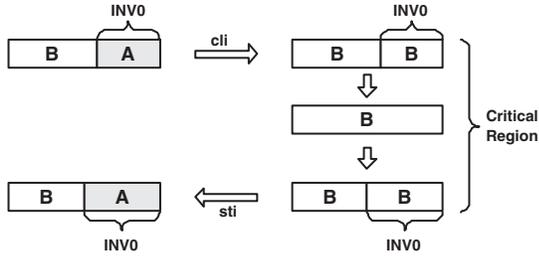


Figure 4. Memory Partition for Handler and Non-Handler

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

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 (O’Hearn 2004; Brookes 2004): *there always 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 (Ishtiaq and O’Hearn 2001; Reynolds 2002), 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 transfers. This semantics completely hides thread queues and thus the complex interleaving between concurrent entities.

We first study the semantics of `cli` and `sti` assuming that the non-handler code is single-threaded. 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

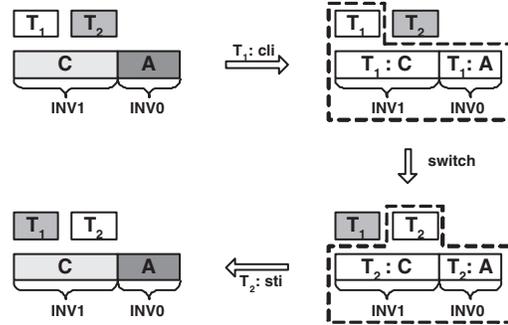


Figure 5. The Memory Model for Multi-Threaded Non-Handler

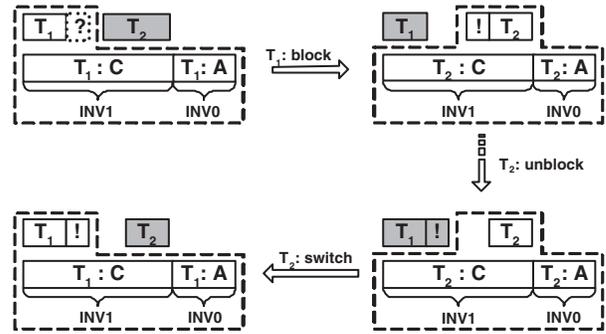


Figure 6. Block and Unblock

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 are only two threads): each thread has its own private memory, and both threads share the block C. When block C is available for sharing, 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 it. 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. The switch operation transfers the ownership of A and C from T_1 to T_2 , knowing that the interrupt remains disabled. Enabling the interrupt (by T_2) releases the ownership.

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 transfers: 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`

(World)	\mathbb{W}	$::= (\mathbb{C}, \mathbb{S}, \mathbb{K}, \text{pc})$
(CodeHeap)	\mathbb{C}	$::= \{\mathbf{f} \rightsquigarrow \mathbf{c}\}^*$
(State)	\mathbb{S}	$::= (\mathbb{H}, \mathbb{R}, \mathbf{ie}, \mathbf{is})$
(Heap)	\mathbb{H}	$::= \{\mathbf{1} \rightsquigarrow \mathbf{w}\}^*$
(RegFile)	\mathbb{R}	$::= \{r_0 \rightsquigarrow w_0, \dots, r_k \rightsquigarrow w_k\}$
(Stack)	\mathbb{K}	$::= \text{nil} \mid \mathbf{f} :: \mathbb{K} \mid (\mathbf{f}, \mathbb{R}) :: \mathbb{K}$
(Bit)	\mathbf{b}	$::= 0 \mid 1$
(Flags)	\mathbf{ie}, \mathbf{is}	$::= \mathbf{b}$
(Labels)	$\mathbf{l}, \mathbf{f}, \text{pc}$	$::= n$ (nat nums)
(Word)	\mathbf{w}	$::= i$ (integers)
(Register)	r	$::= r_0 \mid r_1 \mid \dots$
(Instr)	\mathbf{t}	$::= \text{mov } r_d, r_s \mid \text{movi } r_d, \mathbf{w} \mid \text{add } r_d, r_s \mid \text{sub } r_d, r_s$ $\mid \text{ld } r_d, \mathbf{w}(r_s) \mid \text{st } \mathbf{w}(r_t), r_s \mid \text{beq } r_s, r_t, \mathbf{f} \mid \text{call } \mathbf{f} \mid \text{cli} \mid \text{sti}$
(Commnd)	\mathbf{c}	$::= \mathbf{t} \mid \mathbf{j} \mathbf{f} \mid \text{ret} \mid \text{iret}$
(InstrSeq)	\mathbb{I}	$::= \mathbf{t}; \mathbb{I} \mid \mathbf{j} \mathbf{f} \mid \text{ret} \mid \text{iret}$

Figure 7. Definition of AIM-1

$\mathbb{C}[\mathbf{f}]$	$\stackrel{\text{def}}{=} \begin{cases} \mathbf{c} & \mathbf{c} = \mathbb{C}(\mathbf{f}) \text{ and } \mathbf{c} = \mathbf{j} \mathbf{f}', \text{ ret, or iret} \\ \mathbf{t}; \mathbb{I} & \mathbf{t} = \mathbb{C}(\mathbf{f}) \text{ and } \mathbb{I} = \mathbb{C}[\mathbf{f}+1] \end{cases}$
$(F\{a \rightsquigarrow b\})(x)$	$\stackrel{\text{def}}{=} \begin{cases} b & \text{if } x = a \\ F(x) & \text{otherwise.} \end{cases}$
$\mathbb{S}_{ \mathbb{H}'}$	$\stackrel{\text{def}}{=} (\mathbb{H}', \mathbb{S}, \mathbb{R}, \mathbf{ie}, \mathbf{is})$
$\mathbb{S}_{ \mathbb{R}'}$	$\stackrel{\text{def}}{=} (\mathbb{S}, \mathbb{H}, \mathbb{R}', \mathbf{ie}, \mathbf{is})$
$\mathbb{S}_{ \{\mathbf{ie}=\mathbf{b}\}}$	$\stackrel{\text{def}}{=} (\mathbb{S}, \mathbb{H}, \mathbb{S}, \mathbb{R}, \mathbf{b}, \mathbf{ie}, \mathbf{is})$
$\mathbb{S}_{ \{\mathbf{is}=\mathbf{b}\}}$	$\stackrel{\text{def}}{=} (\mathbb{S}, \mathbb{H}, \mathbb{S}, \mathbb{R}, \mathbf{ie}, \mathbf{b})$

Figure 8. Definition of Representations

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 will 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 Fig. 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 \mathbf{c} . Each command \mathbf{c} is either a sequential or branch instruction \mathbf{t} , or jump or return instructions. The state \mathbb{S} contains the data heap \mathbb{H} , the register file \mathbb{R} , and flags \mathbf{ie} and \mathbf{is} . The binary flags \mathbf{ie} and \mathbf{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 \mathbf{f} or a pair (\mathbf{f}, \mathbb{R}) . 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 \mathbf{f} in \mathbb{C} , as defined in Fig. 8. We use the 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 Fig. 8.

NextS _(c,K) S S'	
if c =	S' =
mov r_d, r_s	$(\mathbb{H}, \mathbb{R}\{r_d \rightsquigarrow \mathbb{R}(r_s)\}, \mathbf{ie}, \mathbf{is})$
movi r_d, \mathbf{w}	$(\mathbb{H}, \mathbb{R}\{r_d \rightsquigarrow \mathbf{w}\}, \mathbf{ie}, \mathbf{is})$
add r_d, r_s	$(\mathbb{H}, \mathbb{R}\{r_d \rightsquigarrow (\mathbb{R}(r_s) + \mathbb{R}(r_d))\}, \mathbf{ie}, \mathbf{is})$
sub r_d, r_s	$(\mathbb{H}, \mathbb{R}\{r_d \rightsquigarrow (\mathbb{R}(r_d) - \mathbb{R}(r_s))\}, \mathbf{ie}, \mathbf{is})$
ld $r_d, \mathbf{w}(r_s)$	$(\mathbb{H}, \mathbb{R}\{r_d \rightsquigarrow \mathbb{H}(\mathbb{R}(r_s) + \mathbf{w})\}, \mathbf{ie}, \mathbf{is})$ if $(\mathbb{R}(r_s) + \mathbf{w}) \in \text{dom}(\mathbb{H})$
st $\mathbf{w}(r_t), r_s$	$(\mathbb{H}\{(\mathbb{R}(r_t) + \mathbf{w}) \rightsquigarrow \mathbb{R}(r_s)\}, \mathbb{R}, \mathbf{ie}, \mathbf{is})$ if $(\mathbb{R}(r_t) + \mathbf{w}) \in \text{dom}(\mathbb{H})$
cli	$\mathbb{S}_{ \{\mathbf{ie}=0\}}$
sti	$\mathbb{S}_{ \{\mathbf{ie}=1\}}$
iret	$(\mathbb{H}, \mathbb{R}', \mathbf{1}, 0)$ if $\mathbf{is} = 1, \mathbb{K} = (\mathbf{f}, \mathbb{R}') :: \mathbb{K}'$ for some \mathbf{f} and \mathbb{K}'
other cases	\mathbb{S}

NextK _(pc,c) K K'	
if c =	K' =
call \mathbf{f}	$(\text{pc}+1) :: \mathbb{K}$
ret	\mathbb{K}'' if $\mathbb{K} = \mathbf{f} :: \mathbb{K}''$ for some \mathbf{f}
iret	\mathbb{K}'' if $\mathbb{K} = (\mathbf{f}, \mathbb{R}) :: \mathbb{K}''$ for some \mathbf{f} and \mathbb{R}
other cases	\mathbb{K}

NextPC _(c,R,K) pc pc'	
if c =	pc' =
beq r_s, r_t, \mathbf{f}	\mathbf{f} if $\mathbb{R}(r_s) = \mathbb{R}(r_t)$
beq r_s, r_t, \mathbf{f}	$\text{pc} + 1$ if $\mathbb{R}(r_s) \neq \mathbb{R}(r_t)$
call \mathbf{f}	\mathbf{f}
$\mathbf{j} \mathbf{f}$	\mathbf{f}
ret	\mathbf{f} if $\mathbb{K} = \mathbf{f} :: \mathbb{K}'$ for some \mathbb{K}'
iret	\mathbf{f} if $\mathbb{K} = (\mathbf{f}, \mathbb{R}') :: \mathbb{K}'$ for some \mathbb{K}' and \mathbb{R}'
other cases	$\text{pc} + 1$

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

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

$$\mathbb{W} \Longrightarrow \mathbb{W}' \stackrel{\text{def}}{=} (\mathbb{W} \mapsto \mathbb{W}') \vee (\mathbb{W} \not\leq \mathbb{W}')$$

Figure 9. Operational Semantics of Instructions

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 $\mathbf{h_entry}$. Support of multi-level interrupts is discussed in Sec. 6. An incoming interrupt is processed only if the \mathbf{ie} bit is set, and no interrupt is currently being serviced (*i.e.*, $\mathbf{is} = 0$). The processor pushes (pc, \mathbb{R}) onto the stack \mathbb{K} , clears \mathbf{ie} , sets \mathbf{is} , and sets the new pc to $\mathbf{h_entry}$. The state transition $(\mathbb{W} \not\leq \mathbb{W}')$ is defined in the IRQ rule in Fig. 9.

The operational semantics of each instruction is defined in Fig. 9. We use relations $\text{NextS}_{(c,K)}$, $\text{NextK}_{(pc,c)}$ and $\text{NextPC}_{(c,R,K)}$ to show the change of states, stacks and program counters respectively when \mathbf{c} is executed. Semantics of most instructions are straightforward. The command iret pops the stack frame $(\mathbf{f}, \mathbb{R}')$ on top of \mathbb{K} and resets pc and the register file \mathbb{R} with \mathbf{f} and \mathbb{R}' , respectively. It also restores \mathbf{ie} and \mathbf{is} with the value when the interrupt occurs, which must be 1 and 0 respectively (otherwise the interrupt cannot have been handled). 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

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

Figure 12. The Step Relation for AIM-2

incleft: $\neg\{(p_0, \text{NoG})\}$	h_entry: $\neg\{(p_i, g_i)\}$	h_entry: $\neg\{(p_i, g_i)\}$
movi \$r1, RIGHT	movi \$r1, LEFT	j h_timer
movi \$r2, LEFT	movi \$r2, RIGHT	h_timer: $\neg\{(p_i, g_i)\}$
l_loop: $\neg\{(p_1, \text{NoG})\}$	movi \$r3, 0	movi \$r1, CNT
movi \$r3, 0	ld \$r4, 0(\$r1)	ld \$r2, 0(\$r1) ; \$r2 <- [CNT]
cli	beq \$r3, \$r4, r_win	movi \$r3, 100
$\neg\{(p_2, \text{NoG})\}$	movi \$r3, 1	beq \$r2, \$r3, schd ; if ([CNT]=100)
ld \$r4, 0(\$r1)	sub \$r4, \$r3	movi \$r3, 1 ; goto schd
beq \$r3, \$r4, l_win	st 0(\$r1), \$r4	add \$r2, \$r3
movi \$r3, 1	ld \$r4, 0(\$r2)	st 0(\$r1), \$r2 ; [CNT]++
sub \$r4, \$r3	add \$r4, \$r3	iret
st 0(\$r1), \$r4	st 0(\$r2), \$r4	schd: $\neg\{(p_0, g_0)\}$
ld \$r4, 0(\$r2)	iret	movi \$r2, 0
add \$r4, \$r3	r_win: $\neg\{(p_4, g_4)\}$	st 0(\$r1), \$r2 ; [CNT] := 0
st 0(\$r2), \$r4	iret	switch
sti		iret
$\neg\{(p_1, \text{NoG})\}$		$p_0 \stackrel{\text{def}}{=} \text{enable}_{\text{iret}} \wedge (r_1 = \text{CNT})$
j l_loop		$g_0 \stackrel{\text{def}}{=} \left\{ \begin{array}{l} \text{CNT} \mapsto - \\ \text{INV0} \end{array} \right\} \wedge (\text{ie} = \text{ie}') \wedge (\text{is} = \text{is}')$
l_win: $\neg\{(p_3, \text{NoG})\}$		
sti		
j l_loop		

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) $\iota ::= \dots \mid \text{switch} \mid \text{block } r_t \mid \text{unlock } r_t, r_d \mid \dots$

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

a cleaner presentation. In our implementation (Feng et al. 2007c), 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 $\mathbb{W} \Rightarrow \mathbb{W}'$.

Fig. 10 shows a sample AIM-1 program. The program specifications in shaded boxes are explained in Sec. 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 interrupts before it accesses LEFT and RIGHT.

3.2 AIM-2

Fig. 11 defines AIM-2 as an extension over AIM-1. We extend the 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 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:

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

We also add three primitive instructions: switch, block and unlock.

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 picks a thread nondeterministically from $\text{readyQ}(\mathbb{T}, \mathbb{B})$ to run. 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 “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 “unlock 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;

$(CdHpSpec)$	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

otherwise r_d contains 0. Here \uplus represents the union of two disjoint sets. 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 $(\mathbb{W} \not\rightarrow \mathbb{W}')$ for AIM-2 is almost the same as the one for AIM-1 defined by the `IRQ` rule. It does not change \mathbb{T} , \mathbb{B} and `tid` either. The definition of $(\mathbb{W} \Longrightarrow \mathbb{W}')$ is unchanged.

A preemptive timer interrupt handler. The design of AIM is very interesting in that it supports both preemptive threads (if the interrupt is enabled and the handler does context switching) and non-preemptive ones (if the interrupt is disabled, or if the interrupt is enabled but the handler does no context switching) for higher-level concurrent programs (see Fig. 1).

Fig. 13 shows the implementation of a preemptive timer interrupt handler. Each time the interrupt comes, the handler tests 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 meanings of specifications and show how the timer handler is certified in Sec. 4.

4. The Program Logic

4.1 Specification Language

We use the mechanized *meta-logic* implemented in the Coq proof assistant (Coq 2006) as our specification language. The logic corresponds to higher-order logic with inductive definitions.

As shown in Fig. 14, the specification Ψ for the code heap \mathbb{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 \mathbb{K} and a program state \mathbb{S} , while g is a predicate over two program states. As we can see, the $\text{NextS}_{(\mathbb{C}, \mathbb{K})}$ relation defined in Fig. 9 is a special form of g . As in SCAP (Feng et al. 2006), 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 where the *current* function returns.

We also use the predicate m to specify data heaps. We encode in Fig. 15 Separation Logic connectors (Ishtiaq and O’Hearn 2001; Reynolds 2002) in our specification language. Assertions in Separation Logic capture ownership of heaps. The assertion “ $1 \mapsto n$ ” holds iff the heap has only one cell at 1 containing 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. “ $m \multimap m'$ ” holds over \mathbb{H} iff, for any disjoint heap \mathbb{H}' satisfying m , $\mathbb{H} \uplus \mathbb{H}'$ satisfies m' .

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

Specifications of the shared resources. Heap predicates `INV0` and `INV1` are part of our program specifications, which specify the well-formedness of the shared sub-heap A and C respectively, as shown in Figs. 5 and 6. The definition of `INV0` depends on the

$$\begin{aligned}
\text{true} &\stackrel{\text{def}}{=} \lambda \mathbb{H}. \text{True} & \text{emp} &\stackrel{\text{def}}{=} \lambda \mathbb{H}. \mathbb{H} = \emptyset \\
1 \mapsto w &\stackrel{\text{def}}{=} \lambda \mathbb{H}. \mathbb{H} = \{1 \rightsquigarrow w\} & 1 \mapsto _ &\stackrel{\text{def}}{=} \lambda \mathbb{H}. \exists w. (1 \mapsto w) \mathbb{H} \\
\mathbb{H}_1 \perp \mathbb{H}_2 &\stackrel{\text{def}}{=} \text{dom}(\mathbb{H}_1) \cap \text{dom}(\mathbb{H}_2) = \emptyset \\
\mathbb{H}_1 \uplus \mathbb{H}_2 &\stackrel{\text{def}}{=} \begin{cases} \mathbb{H}_1 \cup \mathbb{H}_2 & \text{if } \mathbb{H}_1 \perp \mathbb{H}_2 \\ \text{undefined} & \text{otherwise} \end{cases} \\
m_1 * m_2 &\stackrel{\text{def}}{=} \lambda \mathbb{H}. \exists \mathbb{H}_1, \mathbb{H}_2. (\mathbb{H}_1 \uplus \mathbb{H}_2 = \mathbb{H}) \wedge m_1 \mathbb{H}_1 \wedge m_2 \mathbb{H}_2 \\
p * m &\stackrel{\text{def}}{=} \lambda \mathbb{K}, \mathbb{S}. \exists \mathbb{H}_1, \mathbb{H}_2. (\mathbb{H}_1 \uplus \mathbb{H}_2 = \mathbb{S}. \mathbb{H}) \wedge p \mathbb{K} \mathbb{S} |_{\mathbb{H}_1} \wedge m \mathbb{H}_2 \\
m \multimap m' &\stackrel{\text{def}}{=} \lambda \mathbb{H}. \forall \mathbb{H}', \mathbb{H}'' . (\mathbb{H} \uplus \mathbb{H}' = \mathbb{H}'') \wedge m \mathbb{H}' \rightarrow m' \mathbb{H}'' \\
m \multimap p &\stackrel{\text{def}}{=} \lambda \mathbb{K}, \mathbb{S}. \forall \mathbb{H}', \mathbb{H}'' . (\mathbb{H}' \uplus \mathbb{S}. \mathbb{H} = \mathbb{H}'') \wedge m \mathbb{H}' \rightarrow p \mathbb{K} \mathbb{S} |_{\mathbb{H}''} \\
\text{precise}(m) &\stackrel{\text{def}}{=} \forall \mathbb{H}_1, \mathbb{H}_2. \\
&\quad (\mathbb{H}_1 \subseteq \mathbb{H}) \wedge (\mathbb{H}_2 \subseteq \mathbb{H}) \wedge m \mathbb{H}_1 \wedge m \mathbb{H}_2 \rightarrow (\mathbb{H}_1 = \mathbb{H}_2)
\end{aligned}$$

Figure 15. Definitions of Separation Logic Assertions

functionality of the global interrupt handler; and `INV1` depends on the sharing of resources among threads. To simplify the presentation, we treat them as global parameters throughout this paper.¹

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 g_i are defined as follows:

$$\begin{aligned}
p_i &\stackrel{\text{def}}{=} \lambda \mathbb{K}, \mathbb{S}. ((INV0 * \text{true}) \mathbb{S}. \mathbb{H}) \wedge (\mathbb{S}. \text{is} = 1) \wedge (\mathbb{S}. \text{ie} = 0) \\
&\quad \wedge \exists f, \mathbb{R}, \mathbb{K}'. \mathbb{K} = (f, \mathbb{R}) :: \mathbb{K}' \quad (1) \\
g_i &\stackrel{\text{def}}{=} \lambda \mathbb{S}, \mathbb{S}'. \left\{ \begin{array}{l} INV0 \\ INV0 \end{array} \right\} \mathbb{S}. \mathbb{H} \mathbb{S}'. \mathbb{H} \\
&\quad \wedge (\mathbb{S}'. \text{ie} = \mathbb{S}. \text{ie}) \wedge (\mathbb{S}'. \text{is} = \mathbb{S}. \text{is}) \quad (2)
\end{aligned}$$

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 Fig. 5) satisfies `INV0`. 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.

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

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 Fig. 16. The judgment $\Psi, \Delta \vdash \{s\} f : \mathbb{I}$ defines the well-formedness of the instruction sequence \mathbb{I} starting at the code label f , given the imported interfaces in Ψ , the specification Δ of block queues, and a precondition (p, g) .

The `SEQ` rule is a *schema* for instruction sequences starting with an instruction ι (ι cannot be 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

¹ They can also be treated as local parameters threading through judgments in our program logic (as Ψ and Δ in Fig. 16). To avoid the requirement of the global knowledge about shared resources and to have better modularity, frame rules (Reynolds 2002; O’Hearn et al. 2004) can be supported following the same way they are supported in SCAP (Feng and Shao 2008). We do not discuss the details in this paper.

$\Psi, \Delta \vdash \{s\} f : \mathbb{I}$

 (Well-Formed Instr. Seq.)
$$\frac{\begin{array}{l} \iota \notin \{\text{call}, \dots, \text{beq}, \dots\} \quad \Psi, \Delta \vdash \{(p', g')\} f + 1 : \mathbb{I} \\ \text{enable}(p, g_i) \quad (p \triangleright g_i) \Rightarrow p' \quad (p \circ (g_i \circ g')) \Rightarrow g \\ \text{where } g_i \stackrel{\text{def}}{=} [\iota]_{\Delta} \end{array}}{\Psi, \Delta \vdash \{(p, g)\} f : \iota; \mathbb{I}} \quad (\text{SEQ})$$

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

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

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

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

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

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

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

$\Psi, \Delta \vdash C : \Psi'$

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

$\Psi, \Delta \vdash \mathbb{W}$

 (Well-Formed World)
$$\begin{array}{l} T \setminus \text{tid} = \{\text{tid}_1 \rightsquigarrow (\mathbb{R}_1, \mathbb{K}_1, \text{is}_1, \text{pc}_1), \dots, \\ \quad \text{tid}_n \rightsquigarrow (\mathbb{R}_n, \mathbb{K}_n, \text{is}_n, \text{pc}_n)\} \\ \mathbb{S}. \mathbb{H} = \mathbb{H}_0 \uplus \dots \uplus \mathbb{H}_n \quad \mathbb{S}_i = (\mathbb{H}_i, \mathbb{R}_i, 0, \text{is}_i) \quad (0 < i \leq n) \\ \Psi, \Delta \vdash C : \Psi' \quad \Psi \subseteq \Psi' \quad \text{dom}(\Delta) = \text{dom}(\mathbb{B}) \\ \text{WFCth}(\mathbb{S}_0, \mathbb{K}, \text{pc}, \Psi') \quad \text{where } \mathbb{S}_0 = \mathbb{S}|_{\mathbb{H}_0} \\ \text{for all } 0 < k \leq n \text{ such that } \text{tid}_k \in \text{readyQ}(T, \mathbb{B}) : \\ \quad \text{WFRdy}(\mathbb{S}_k, \mathbb{K}_k, \text{pc}_k, \Psi') \\ \text{for all } w \text{ and } 0 < j \leq n \text{ such that } \text{tid}_j \in \mathbb{B}(w) : \\ \quad \text{WFWait}(\mathbb{S}_j, \mathbb{K}_j, \text{pc}_j, \Psi', \Delta(w)) \end{array}$$

$$\Psi, \Delta \vdash (C, \mathbb{S}, \mathbb{K}, \text{pc}, \text{tid}, T, \mathbb{B}) \quad (\text{WLD})$$

Figure 16. Inference Rules

to represent the state transition $[\iota]_{\Delta}$ made by the instruction ι , which is defined in Fig. 18 and is explained below. The premise $\text{enable}(p, g_i)$ is defined in Fig. 17. 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 Fig. 17, 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 $[\iota]_{\Delta}$ in Fig. 18 as $\text{NextS}_{(\iota, _)}$. Since NextS does not depend on the stack for these instructions (recall its definition

$$\begin{aligned} \text{enable}(p, g) &\stackrel{\text{def}}{=} \forall \mathbb{K}, \mathbb{S}. p \mathbb{K} \mathbb{S} \rightarrow \exists \mathbb{S}', g \mathbb{S} \mathbb{S}' \\ p \triangleright g &\stackrel{\text{def}}{=} \lambda \mathbb{K}, \mathbb{S}. \exists \mathbb{S}_0, p \mathbb{K} \mathbb{S}_0 \wedge g \mathbb{S}_0 \mathbb{S} \\ g \circ g' &\stackrel{\text{def}}{=} \lambda \mathbb{S}, \mathbb{S}'. \exists \mathbb{S}'' . g \mathbb{S} \mathbb{S}' \wedge g' \mathbb{S}' \mathbb{S}'' \quad p \Rightarrow p' \stackrel{\text{def}}{=} \forall \mathbb{K}, \mathbb{S}. p \mathbb{K} \mathbb{S} \rightarrow p' \mathbb{K} \mathbb{S} \\ p \circ g &\stackrel{\text{def}}{=} \lambda \mathbb{S}, \mathbb{S}'. \exists \mathbb{K}. p \mathbb{K} \mathbb{S} \wedge g \mathbb{S} \mathbb{S}' \quad g \Rightarrow g' \stackrel{\text{def}}{=} \forall \mathbb{S}, \mathbb{S}'. g \mathbb{S} \mathbb{S}' \rightarrow g' \mathbb{S} \mathbb{S}' \end{aligned}$$

Figure 17. Connectors for p and g

$$P ? m : m' \stackrel{\text{def}}{=} \lambda \mathbb{H}. (P \wedge m \mathbb{H}) \vee (\neg P \wedge m' \mathbb{H})$$

$$\llbracket \text{cli} \rrbracket_{\Delta} \stackrel{\text{def}}{=} \lambda (\mathbb{H}, \mathbb{R}, \text{ie}, \text{is}), (\mathbb{H}', \mathbb{R}', \text{ie}', \text{is}'). \left. \begin{array}{l} (\text{is} = \text{is}') \wedge (\mathbb{R} = \mathbb{R}') \wedge (\text{ie}' = 0) \wedge \\ \left\{ \begin{array}{l} \text{emp} \\ (\text{ie} = 1 \wedge \text{is} = 0) ? (\text{INV0} * \text{INV1}) : \text{emp} \end{array} \right\} \mathbb{H} \mathbb{H}' \end{array} \right\}$$

$$\llbracket \text{sti} \rrbracket_{\Delta} \stackrel{\text{def}}{=} \lambda (\mathbb{H}, \mathbb{R}, \text{ie}, \text{is}), (\mathbb{H}', \mathbb{R}', \text{ie}', \text{is}'). \left. \begin{array}{l} (\text{is} = \text{is}') \wedge (\mathbb{R} = \mathbb{R}') \wedge (\text{ie}' = 1) \wedge \\ \left\{ \begin{array}{l} (\text{ie} = 0 \wedge \text{is} = 0) ? (\text{INV0} * \text{INV1}) : \text{emp} \\ \text{emp} \end{array} \right\} \mathbb{H} \mathbb{H}' \end{array} \right\}$$

$$\llbracket \text{switch} \rrbracket_{\Delta} \stackrel{\text{def}}{=} \lambda (\mathbb{H}, \mathbb{R}, \text{ie}, \text{is}), (\mathbb{H}', \mathbb{R}', \text{ie}', \text{is}'). \left. \begin{array}{l} (\text{ie} = 0) \wedge (\text{ie} = \text{ie}') \wedge (\mathbb{R} = \mathbb{R}') \wedge (\text{is} = \text{is}') \wedge \\ \left\{ \begin{array}{l} \text{INV0} * (\text{is} = 0 ? \text{INV1} : \text{emp}) \\ \text{INV0} * (\text{is} = 0 ? \text{INV1} : \text{emp}) \end{array} \right\} \mathbb{H} \mathbb{H}' \end{array} \right\}$$

$$\llbracket \text{block } r_s \rrbracket_{\Delta} \stackrel{\text{def}}{=} \lambda (\mathbb{H}, \mathbb{R}, \text{ie}, \text{is}), (\mathbb{H}', \mathbb{R}', \text{ie}', \text{is}'). \left. \begin{array}{l} (\text{ie} = 0) \wedge (\text{ie} = \text{ie}') \wedge (\mathbb{R} = \mathbb{R}') \wedge (\text{is} = \text{is}') \wedge \\ \exists m. \Delta(\mathbb{R}(r_s)) = m \wedge \\ \left\{ \begin{array}{l} \text{INV0} * (\text{is} = 0 ? \text{INV1} : \text{emp}) \\ \text{INV0} * (\text{is} = 0 ? \text{INV1} : \text{emp}) * m \end{array} \right\} \mathbb{H} \mathbb{H}' \end{array} \right\}$$

$$\llbracket \text{unblock } r_s, r_d \rrbracket_{\Delta} \stackrel{\text{def}}{=} \lambda (\mathbb{H}, \mathbb{R}, \text{ie}, \text{is}), (\mathbb{H}', \mathbb{R}', \text{ie}', \text{is}'). \left. \begin{array}{l} (\text{ie} = 0) \wedge (\text{ie} = \text{ie}') \wedge (\text{is} = \text{is}') \wedge (\forall r \neq r_d. \mathbb{R}(r) = \mathbb{R}'(r)) \wedge \\ \exists m. \Delta(\mathbb{R}(r_s)) = m \wedge (m * \text{true}) \mathbb{H} \wedge \\ \left\{ \begin{array}{l} (\mathbb{R}'(r_d) = 0) ? \text{emp} : m \\ \text{emp} \end{array} \right\} \mathbb{H} \mathbb{H}' \end{array} \right\}$$

$$\llbracket \iota \rrbracket_{\Delta} \stackrel{\text{def}}{=} \text{NextS}_{(\iota, _)} \quad (\text{for all other } \iota)$$

Figure 18. Thread-Local State Transitions Made by ι

in Fig. 9), 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.

Interrupts and thread primitive instructions. One of the major technical contributions of this paper is our formulation of $[\iota]_{\Delta}$ for cli , sti , switch , block and unblock , which, as shown in Fig. 18, gives them an axiomatic ownership transfer semantics.

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

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

$\llbracket \text{block } r_s \rrbracket_{\Delta}$ requires $\text{ie} = 0$ and r_s contains an identifier of a block queue with specification m in Δ . It is similar to $\llbracket \text{switch} \rrbracket_{\Delta}$, except that the thread gets the ownership of m after it is released (see Fig. 6). In $\llbracket \text{unblock } r_s, r_d \rrbracket_{\Delta}$, we require the initial heap must

$$\begin{aligned}
\text{Inv}(\text{ie}, \text{is}) &\stackrel{\text{def}}{=} \begin{cases} \text{INV1} & \text{is} = 1 \\ \text{emp} & \text{is} = 0 \text{ and } \text{ie} = 0 \\ \text{INV}_s & \text{is} = 0 \text{ and } \text{ie} = 1 \end{cases} \\
\text{where } \text{INV}_s &\stackrel{\text{def}}{=} \text{INV0} * \text{INV1} \\
\text{p} * \text{Inv} &\stackrel{\text{def}}{=} \lambda \mathbb{K}, \mathbb{S}. (\text{p} * \text{Inv}(\mathbb{S}.\text{ie}, \mathbb{S}.\text{is})) \mathbb{K} \mathbb{S} \\
\lfloor \mathbf{g} \rfloor &\stackrel{\text{def}}{=} \lambda (\mathbb{H}, \mathbb{R}, \text{ie}, \text{is}), (\mathbb{H}', \mathbb{R}', \text{ie}', \text{is}'). \exists \mathbb{H}_1, \mathbb{H}_2, \mathbb{H}'_1, \mathbb{H}'_2. \\
&\quad (\mathbb{H}_1 \uplus \mathbb{H}_2 = \mathbb{H}) \wedge (\mathbb{H}'_1 \uplus \mathbb{H}'_2 = \mathbb{H}') \\
&\quad \wedge \mathbf{g}(\mathbb{H}_1, \mathbb{R}, \text{ie}, \text{is}) (\mathbb{H}'_1, \mathbb{R}', \text{ie}', \text{is}') \\
&\quad \wedge \text{Inv}(\text{ie}, \text{is}) \mathbb{H}_2 \wedge \text{Inv}(\text{ie}', \text{is}') \mathbb{H}'_2 \\
\text{WFST}(\mathbf{g}, \mathbb{S}, \text{nil}, \Psi) &\stackrel{\text{def}}{=} \neg \exists \mathbb{S}'. \mathbf{g} \mathbb{S} \mathbb{S}' \\
\text{WFST}(\mathbf{g}, \mathbb{S}, \mathbf{f} :: \mathbb{K}, \Psi) &\stackrel{\text{def}}{=} \\
&\quad \exists \mathbf{p}_{\mathbf{f}}, \mathbf{g}_{\mathbf{f}}. (\mathbf{f}, (\mathbf{p}_{\mathbf{f}}, \mathbf{g}_{\mathbf{f}})) \in \Psi \\
&\quad \wedge \forall \mathbb{S}'. \mathbf{g} \mathbb{S} \mathbb{S}' \rightarrow (\mathbf{p}_{\mathbf{f}} * \text{Inv}) \mathbb{K} \mathbb{S}' \wedge \text{WFST}(\lfloor \mathbf{g}_{\mathbf{f}} \rfloor, \mathbb{S}', \mathbb{K}, \Psi) \\
\text{WFST}(\mathbf{g}, \mathbb{S}, (\mathbf{f}, \mathbb{R}) :: \mathbb{K}, \Psi) &\stackrel{\text{def}}{=} \\
&\quad \exists \mathbf{p}_{\mathbf{f}}, \mathbf{g}_{\mathbf{f}}. (\mathbf{f}, (\mathbf{p}_{\mathbf{f}}, \mathbf{g}_{\mathbf{f}})) \in \Psi \\
&\quad \wedge \forall \mathbb{S}'. \mathbf{g} \mathbb{S} \mathbb{S}' \rightarrow (\mathbf{p}_{\mathbf{f}} * \text{Inv}) \mathbb{K} \mathbb{S}' \wedge \text{WFST}(\lfloor \mathbf{g}_{\mathbf{f}} \rfloor, \mathbb{S}', \mathbb{K}, \Psi) \\
&\quad \text{where } \mathbb{S}' = (\mathbb{S}'.\mathbb{H}, \mathbb{R}, 1, 0) \\
\text{WFCth}(\mathbb{S}, \mathbb{K}, \text{pc}, \Psi) &\stackrel{\text{def}}{=} \exists \mathbf{p}, \mathbf{g}. (\text{pc}, (\mathbf{p}, \mathbf{g})) \in \Psi \\
&\quad \wedge (\mathbf{p} * \text{Inv}) \mathbb{K} \mathbb{S} \wedge \text{WFST}(\lfloor \mathbf{g} \rfloor, \mathbb{S}, \mathbb{K}, \Psi) \\
\text{WFCth}(\text{pc}, \Psi) &\stackrel{\text{def}}{=} \lambda \mathbb{K}, \mathbb{S}. \text{WFCth}(\mathbb{S}, \mathbb{K}, \text{pc}, \Psi) \\
\text{WFRdy}(\mathbb{S}, \mathbb{K}, \text{pc}, \Psi) &\stackrel{\text{def}}{=} ((\text{INV0} * \text{INV1}) \multimap \text{WFCth}(\text{pc}, \Psi)) \mathbb{K} \mathbb{S} \\
\text{WFRdy}(\text{pc}, \Psi) &\stackrel{\text{def}}{=} \lambda \mathbb{K}, \mathbb{S}. \text{WFRdy}(\mathbb{S}, \mathbb{K}, \text{pc}, \Psi) \\
\text{WFWait}(\mathbb{S}, \mathbb{K}, \text{pc}, \Psi, \mathbf{m}) &\stackrel{\text{def}}{=} (\mathbf{m} \multimap \text{WFRdy}(\text{pc}, \Psi)) \mathbb{K} \mathbb{S}
\end{aligned}$$

Figure 19. Well-Formed Current, Ready and Waiting Threads

contain a sub-heap satisfying \mathbf{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 r_d contains non-zero value at the end of unblock, some thread has been released from the block queue. The current thread transfers \mathbf{m} to the released thread and has no access to it any more. Otherwise, no thread is released and there is no ownership transfer.

Other instructions. In the `CALL` rule in Fig. 16, we treat the state transition \mathbf{g}' made by the callee as the transition of the call instruction. We also require that the precondition \mathbf{p} implies the precondition \mathbf{p}' of the callee, which corresponds to the `enable` premise in the `SEQ` rule. `IRET` and `RET` rules require that the function has finished its guaranteed transition at this point. So an identity transition \mathbf{g} should satisfy the remaining transition \mathbf{g} . The predicates `enableiret` and `enableret` specify the requirements over stacks. In the `BEQ` rule, we use `gidrs=rt` and `gidrs≠rt` to represent identity transitions with extra knowledge about r_s and r_t :

$$\begin{aligned}
\text{gid} &\stackrel{\text{def}}{=} \lambda \mathbb{S}, \mathbb{S}'. \mathbb{S} = \mathbb{S}' \\
\text{gid}_{r_s=r_t} &\stackrel{\text{def}}{=} \lambda \mathbb{S}, \mathbb{S}'. (\text{gid } \mathbb{S} \mathbb{S}') \wedge (\mathbb{S}.\mathbb{R}(r_s) = \mathbb{S}.\mathbb{R}(r_t)) \\
\text{gid}_{r_s \neq r_t} &\stackrel{\text{def}}{=} \lambda \mathbb{S}, \mathbb{S}'. (\text{gid } \mathbb{S} \mathbb{S}') \wedge (\mathbb{S}.\mathbb{R}(r_s) \neq \mathbb{S}.\mathbb{R}(r_t))
\end{aligned}$$

We do not have an `enable` premise because `beq` never gets stuck. The `J` rule can be viewed as a specialization of `BEQ`.

Well-formed code heap. The `CDHP` rule says the code heap is well-formed if and only if each instruction sequence specified in Ψ' is well-formed. Ψ and Ψ' can be viewed as the imported and exported interfaces of \mathbb{C} respectively.

Program invariants. The `WLD` rule formulates the program invariant enforced by our program logic. If there are n threads in \mathbb{T} in addition to the current thread, the heap can be split into $n + 1$ blocks. Each block \mathbb{H}_k ($k > 0$) is for a ready or blocked thread in

$$\begin{aligned}
\mathbf{p} &\stackrel{\text{def}}{=} (\text{ie} = 1) \wedge (\text{is} = 0) & \mathbf{p}' &\stackrel{\text{def}}{=} (\text{ie} = 0) \wedge (\text{is} = 0) & \mathbf{p}_0 &\stackrel{\text{def}}{=} \mathbf{p} \\
\mathbf{p}_1 &\stackrel{\text{def}}{=} \mathbf{p} \wedge (r_1 = \text{RIGHT}) \wedge (r_2 = \text{LEFT}) \\
\mathbf{p}_2 &\stackrel{\text{def}}{=} \mathbf{p}' \wedge (r_1 = \text{RIGHT}) \wedge (r_2 = \text{LEFT}) \wedge (r_3 = 0) \wedge (\text{INV0} * \text{true}) \\
\mathbf{p}_3 &\stackrel{\text{def}}{=} \mathbf{p}' \wedge (r_1 = \text{RIGHT}) \wedge (r_2 = \text{LEFT}) \wedge (\text{INV0} * \text{true}) \\
\mathbf{p}_4 &\stackrel{\text{def}}{=} \text{enable}_{\text{iret}} & \text{NoG} &\stackrel{\text{def}}{=} \lambda \mathbb{S}, \mathbb{S}'. \text{False}
\end{aligned}$$

Figure 20. Specifications of the Teeter-Totter Example

queues. The block \mathbb{H}_0 is assigned to the current thread, which includes both its private heap and the shared part (blocks A and C, as shown in Fig. 5). The code heap \mathbb{C} needs to be well-formed, as defined by the `CDHP` rule. We require the imported interface Ψ is a subset of the exported interface Ψ' , therefore \mathbb{C} is self-contained and each imported specification has been certified. The domain of Δ should be the same with the domain of \mathbb{B} , *i.e.*, Δ specifies and only specifies block queues in \mathbb{B} . The `WLD` rule also requires that the local heaps and execution contexts of the current thread, ready threads and blocked threads are all well-formed (see Fig. 19).

`WFCth` defines the well-formedness of the current thread. It requires that the `pc` has a specification (\mathbf{p}, \mathbf{g}) in Ψ , thus $\mathbb{C}[\text{pc}]$ is well-formed with respect to (\mathbf{p}, \mathbf{g}) . The current thread's stack and its local state (containing the sub-heap \mathbb{H}_0) need to satisfy $\mathbf{p} * \text{Inv}$. Here \mathbf{p} specifies the state accessible by the current thread, while `Inv`, defined in Fig. 19, specifies the inaccessible part of the shared heap. If the current program point is in the interrupt handler ($\text{is} = 1$), \mathbf{p} leaves the memory block C (in Fig. 5) unspecified, therefore `Inv` is defined as `INV1` and specifies the well-formedness of C. Otherwise ($\text{is} = 0$), if $\text{ie} = 0$, blocks A and C become the current thread's private memory and the inaccessible part is empty. If $\text{ie} = 1$, A and C are inaccessible; `Inv` specifies their well-formedness in this case. The predicate `WFST`, defined in Fig. 19, says there exists a well-formed stack with some depth k . The definition is similar to the one in SCAP (Feng et al. 2006) and is not explained here.

The definition of well-formed ready threads `WFRdy` is very straightforward: if the *ready* thread gets the extra ownership of shared memory A and C, it becomes a well-formed *current* thread (see Fig. 5). Recall that $\mathbf{m} \multimap \mathbf{p}$ is defined in Fig. 15. Similarly, `WFWait` says that the *waiting* thread in a block queue waiting for the resource \mathbf{m} becomes a well-formed *ready* thread if it gets \mathbf{m} (see Fig. 6). The definitions of `WFRdy` and `WFWait` concisely formulate the relationship between current, ready and waiting threads.

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

$$\text{INV0} \stackrel{\text{def}}{=} \exists \mathbf{w}_l, \mathbf{w}_r. ((\text{LEFT} \mapsto \mathbf{w}_l) * (\text{RIGHT} \mapsto \mathbf{w}_r)) \wedge (\mathbf{w}_l + \mathbf{w}_r = n),$$

where n is an auxiliary logical variable. Then we can get the concrete specification of the interrupt handler, following Formulae (1) and (2) in Sec. 4.1. We let `INV1` be `emp`, since the non-handler code is sequential.

Specifications are shown in Fig. 20. Recall `enableiret` is defined in Fig. 16. To simplify our presentation, we present the predicate \mathbf{p} in the form of a proposition with free variables referring to components of the state \mathbb{S} . Also, we use \mathbf{m} as a shorthand for the proposition $\mathbf{m} \mathbb{H}$ when there is no confusion.

If we compare \mathbf{p}_1 and \mathbf{p}_2 , we will see that the non-handler code cannot access memory at addresses `LEFT` and `RIGHT` without first disabling the interrupt because \mathbf{p}_1 does not contain the ownership of `LEFT` and `RIGHT`. Since the non-handler never returns, we simply use `NoG` (see Fig. 20) as the guarantee for the state transition from the specified point to the return point.

The timer handler. We also briefly explain the specification for the preemptive timer handler shown in Fig. 13. The handler only

accesses the memory cell at the location CNT. We instantiate INV0 below:

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

Then we get the specification of the handler (p_i, g_i) by Formulae (1) and (2). In g_0 (shown in Fig. 13), we use primed variable (e.g., ie' and is') to refer to components in the second state.

Soundness. We prove the soundness of the program logic following the syntactic approach. Based on the progress and preservation lemmas, we know the program never gets stuck as long as the initial state satisfies the program invariant defined by the WLD rule. *More importantly*, we know the invariant always holds during execution, from which we can derive rich properties of programs. Here, we only show the soundness theorem formalizing the partial correctness of programs. See the TR (Feng et al. 2007c) for proof details.

Theorem 4.1 (Soundness)

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

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

Recall that preciseness is defined in Fig. 15, and the specification (p_i, g_i) is defined by Formulae (1) and (2).

5. More Examples

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 to specify invariants of memory blocks protected by locks.

$$\begin{aligned} (\text{LockID}) \quad l &::= 1 \\ (\text{LockSpec}) \quad \Gamma &::= \{l \rightsquigarrow m\}^* \end{aligned}$$

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 $\text{INV0} * \text{INV1}$).

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

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

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

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

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 spinlocks. More detailed explanations are given in the TR (Feng et al. 2007c).

The Hoare-style implementations. 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. 21 and 22. 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

```

ACQ_H:  -{(p01, g01)}
        cli
        call ACQ_H_a
        sti
        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 21. Hoare-Style Implementation of Locks

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

Figure 22. Specifications of Hoare-Style Locks

$$p_0 \stackrel{\text{def}}{=} (is = 0) \wedge \text{enable}_{\text{ret}} \wedge (r_1 \in \text{dom}(\Gamma)) \wedge (\Delta(r_1) = \text{emp})$$

$$p_{11} \stackrel{\text{def}}{=} p_0 \wedge (ie = 1) \quad p_{12} \stackrel{\text{def}}{=} p_{11} \wedge (r_3 = 0)$$

$$g_{11} \stackrel{\text{def}}{=} \left\{ \begin{array}{l} \text{emp} \\ \Gamma(r_1) \end{array} \right\} \wedge (ie = ie') \wedge (is = is') \wedge \text{trash}(\{r_2, r_3\})$$

$$p_{13} \stackrel{\text{def}}{=} p_0 \wedge (ie = 0) \wedge (\text{INV}_s * \text{true})$$

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

$$p_{21} \stackrel{\text{def}}{=} p_0 \wedge (ie = 1) \wedge (\Gamma(r_1) * \text{true})$$

$$g_{21} \stackrel{\text{def}}{=} \left\{ \begin{array}{l} \Gamma(r_1) \\ \text{emp} \end{array} \right\} \wedge (ie = ie') \wedge (is = is') \wedge \text{trash}(\{r_2\})$$

```

ACQ_M:     $\neg\{(p_{11}, g_{11})\}$ 
          movi    $r3, 0
acq_loop:  $\neg\{(p_{12}, g_{11})\}$ 
          cli
          ld     $r2, 0($r1)          ;; $r2 <- [l]
          beq   $r2, $r3, gowait     ;; ([l] == 0)?
          st    0($r1), $r3          ;; No: [l] <- 0
          sti
          ret
gowait:    $\neg\{(p_{13}, g_{13})\}$ 
          block $r1
          sti
          j     acq_loop
REL_M:     $\neg\{(p_{21}, g_{21})\}$ 
          cli
          unblock $r1, $r2
          movi   $r2, 1
          st    0($r1), $r2
          sti
          ret

```

Figure 23. Mesa-Style Locks

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.

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 and REL_{H_a} are specified by (p_{21}, g_{21}) and (p_{31}, g_{31}) , respectively. 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. Fig. 23 shows the Mesa-style implementation of locks. 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 . More intermediate assertions are given in the technical report.

Spinlocks. An implementation of spinlocks for uniprocessor systems and its specifications are shown in Fig. 24. 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 af-

$$p \stackrel{\text{def}}{=} (is = 0) \wedge \text{enable}_{\text{ret}} \wedge (r_1 \in \text{dom}(\Gamma))$$

$$p_{11} \stackrel{\text{def}}{=} p \wedge (ie = 1) \quad p_{12} \stackrel{\text{def}}{=} p_{11} \wedge (r_2 = 1)$$

$$g_{11} \stackrel{\text{def}}{=} \left\{ \begin{array}{l} \text{emp} \\ \Gamma(r_1) \end{array} \right\} \wedge (ie = ie') \wedge (is = is') \wedge \text{trash}(\{r_2, r_3\})$$

$$p_{13} \stackrel{\text{def}}{=} p \wedge (ie = 0) \wedge ([r_1] = 1) \wedge (\text{INV}_s * \text{true})$$

$$g_{13} \stackrel{\text{def}}{=} \left\{ \begin{array}{l} \text{INV}_s \\ \Gamma(r_1) \end{array} \right\} \wedge (ie = 1 - ie') \wedge (is = is') \wedge \text{trash}(\{r_2\})$$

$$p_{21} \stackrel{\text{def}}{=} p \wedge (ie = 1) \wedge (\Gamma(r_1) * \text{true})$$

$$g_{21} \stackrel{\text{def}}{=} \left\{ \begin{array}{l} \Gamma(r_1) \\ \text{emp} \end{array} \right\} \wedge (ie = ie') \wedge (is = is') \wedge \text{trash}(\{r_2\})$$

```

;; acquire(l): $r1 contains l
spin_acq:  $\neg\{(p_{11}, g_{11})\}$ 
          movi   $r2, 1
spin_loop:  $\neg\{(p_{12}, g_{11})\}$ 
          cli
          ld     $r3, 0($r1)
          beq   $r2, $r3, spin_set
          sti
          j     spin_loop
spin_set:   $\neg\{(p_{13}, g_{13})\}$ 
          movi   $r2, 0
          st    0($r1), $r2
          sti
          ret
;; release(l): $r1 contains l
spin_rel:   $\neg\{(p_{21}, g_{21})\}$ 
          movi   $r2, 1
          cli
          st    0($r1), $r2
          sti
          ret

```

Figure 24. A Spinlock

```

WAIT_H:    $\neg\{(p_{11}, g_{11})\}$           ;; wait(l, cv)
          cli
          mov   $r4, $r2
          call REL_H_a
          block $r4
          sti
          ret
SIGNAL_H:  $\neg\{(p_{21}, g_{21})\}$           ;; signal(l, cv)
          cli
          unblock $r2, $r3
          movi   $r4, 0
          beq   $r3, $r4, sig_done
          block $r1
sig_done:  $\neg\{(p_{22}, g_{22})\}$ 
          sti
          ret
SIGNAL_BH:  $\neg\{(p_{31}, g_{31})\}$           ;; signal(l, cv)
          cli
          unblock $r2, $r3          ;; $r2 contains cv
          movi   $r4, 0
          beq   $r3, $r4, sig_cont
          sti
          ret
sig_cont:  $\neg\{(p_{32}, g_{32})\}$ 
          call   REL_H_a          ;; $r1 contains l
          sti
          ret

```

Figure 25. Impl. of CV - Hoare Style and Brinch Hansen Style

$$\begin{aligned}
\text{Cond}(r, r') &\stackrel{\text{def}}{=} \Gamma(r) \wedge (\Upsilon(r') * \text{true}) & \overline{\text{Cond}}(r, r') &\stackrel{\text{def}}{=} \Gamma(r) \wedge \neg(\Upsilon(r') * \text{true}) \\
p(r, r') &\stackrel{\text{def}}{=} (\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) \\
&\quad \wedge (\Upsilon(cv) = m') \wedge (\Delta(cv) = \text{Cond}(r, r')) \\
p_{11} &\stackrel{\text{def}}{=} p(r_1, r_2) \wedge (\text{ie} = 1) \wedge (\overline{\text{Cond}}(r_1, r_2) * \text{true}) \\
g_{11} &\stackrel{\text{def}}{=} \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_{21} &\stackrel{\text{def}}{=} p(r_1, r_2) \wedge (\text{ie} = 1) \wedge (\text{Cond}(r_1, r_2) * \text{true}) \\
g_{21} &\stackrel{\text{def}}{=} \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} &\stackrel{\text{def}}{=} p(r_1, r_2) \wedge (\text{ie} = 0) \wedge (\text{INV}_s * \Gamma(r_1) * \text{true}) \\
g_{22} &\stackrel{\text{def}}{=} \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_{31} &\stackrel{\text{def}}{=} p_{21} & p_{32} &\stackrel{\text{def}}{=} p_{22} \\
g_{31} &\stackrel{\text{def}}{=} \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\}) \\
g_{32} &\stackrel{\text{def}}{=} \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 26. Spec. of CV - Hoare Style and Brinch Hansen Style

terwards. 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 Hoare style (Hoare 1974), Brinch Hansen style (Brinch Hansen 1975), and Mesa style (Lampson and Redell 1980) condition variables. Below we use Υ to specify the conditions associated with condition variables.

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

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 implementations and specifications are shown in Figs. 25 and 26. 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 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.

$$\begin{aligned}
p(r, r') &\stackrel{\text{def}}{=} (\text{is} = 0) \wedge \text{enable}_{\text{ret}} \wedge \\
&\quad \exists l, cv, m, m'. (r = l) \wedge (r' = cv) \wedge (\Gamma(l) = m) \wedge (\Delta(l) = m) \\
&\quad \wedge (\Upsilon(cv) = m') \wedge (\Delta(cv) = \text{emp}) \\
p_{11} &\stackrel{\text{def}}{=} p(r_1, r_2) \wedge (\text{ie} = 1) \wedge (\overline{\text{Cond}}(r_1, r_2) * \text{true}) \\
g_{11} &\stackrel{\text{def}}{=} \left\{ \begin{array}{l} \overline{\text{Cond}}(r_1, r_2) \\ \Gamma(r_1) \end{array} \right\} \wedge (\text{ie} = \text{ie}') \wedge (\text{is} = \text{is}') \wedge \text{trash}(\{r_2, r_3, r_4\}) \\
p'(r) &\stackrel{\text{def}}{=} (\text{is} = 0) \wedge \exists cv, m. (r = cv) \wedge (\Upsilon(cv) = m) \wedge (\Delta(cv) = \text{emp}) \\
p_{21} &\stackrel{\text{def}}{=} p'(r_1) \wedge (\text{ie} = 1) \\
g_{21} &\stackrel{\text{def}}{=} \text{hid} \wedge (\text{ie} = \text{ie}') \wedge (\text{is} = \text{is}') \wedge \text{trash}(\{r_2\}) \\
\text{WAIT_M:} &\quad \text{-(}\{p_{11}, g_{11}\}\text{)} && ; ; \text{wait}(l, cv) \\
&\quad \text{cli} \\
&\quad \text{mov} \quad \$r4, \$r2 \\
&\quad \text{call} \quad \text{REL_H_a} \\
&\quad \text{block} \quad \$r4 \\
&\quad \text{sti} \\
&\quad \text{call} \quad \text{ACQ_H} \\
&\quad \text{ret} \\
\text{SIGNAL_M:} &\quad \text{-(}\{p_{21}, g_{21}\}\text{)} && ; ; \text{signal}(cv) \\
&\quad \text{cli} \\
&\quad \text{unblock} \quad \$r1, \$r2 \\
&\quad \text{sti} \\
&\quad \text{ret}
\end{aligned}$$

Figure 27. Impl. and Spec. of CV - Mesa Style

Brinch Hansen style condition variables are similar to Hoare-style. The wait function is the same as `WAIT_H`. The signal function `SIGNAL_BH` is specified by (p_{31}, g_{31}) defined in Fig. 26. Here p_{31} is the same as p_{21} . The definition of g_{31} shows the difference: the lock is released when signal returns. Therefore, calling the signal function must be the last command in the critical region.

Mesa-style. Fig. 27 shows Mesa-style condition variables. `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.

`SIGNAL_M` is specified by (p_{21}, g_{21}) . The assertion `hid` is defined in Fig. 22, 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. Intermediate assertions for these examples are given in the TR.

6. Implementations and Further Extensions

The program logic presented in this paper has been adapted for the 16-bit, real-mode x86 architecture. We have formalized a subset of the x86 assembly language, its operational semantics, and the program logic in the Coq proof assistant (Coq 2006). In our implementation, we assume that all interrupts except the timer have been masked. Soundness of the program logic is proved in an OCAP-like (Feng et al. 2007b) framework: inference rules are proved as lemmas in the foundational framework; the soundness of the framework itself is then proved following the syntactic approach. The proof is also formalized in Coq and is machine-checkable.

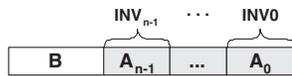
Our preemptive thread libraries (shown in Fig. 3) are also implemented in the x86 assembly code and works in real-mode. Synchronization primitives at Level C are certified using the AIM program logic. The timer handler calls the scheduler implemented at the low-level, which corresponds to the `switch` instruction in AIM. The yield function simply wraps the scheduler by disabling the interrupt at the beginning and enabling it at the end. They are also certified using this logic. Thread primitives at Level S in Fig. 3 are

certified as sequential code. Linking of the certified low-level code with the middle-level libraries and the timer handler is done in the OCAP-like framework. Linking of the thread library code at the middle level with the high-level concurrent programs (see Fig. 1) can be done in a similar way and is left as future work.

The Coq implementation has taken many man-months, out of which a significant amount of efforts has been put on the implementation of basic facilities, including lemmas and tactics for partial mappings, heaps, queues, and Separation Logic assertions. These common facilities are independent of the task of certifying thread libraries and can be reused in future projects. The size of the proof scripts, in terms of the number of lines of Coq tactics, is huge compared with the size of the x86 code. For instance, the proof for the Mesa-style condition variables (26 lines of x86 code) is around 5400 lines, including comments and white spaces. However, as observed by McCreight et al. (2007), the length of proof is probably a poor metric of complexity. There is a lot of redundancy in the proof: when an instruction is seen a second time in the code, we simply copy and paste the previous proof, and do some minor changes. We hope the length of the proof can be greatly reduced given better abstractions and tactics. Also, the 5400 line of proof was finished only in two days by one of the authors, who is an experienced Coq user. We believe this is a very reasonable price to pay for fully certified subroutines with machine checkable proofs. The whole Coq implementation is released as part of the TR (Feng et al. 2007c).

Extensions and future work. 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 partitions. 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 $\text{Inv}(\text{ie}, \text{is})$ in Fig. 19 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 run-time 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 creation of threads and block queues in this paper. In our previous work (Feng and Shao 2005), 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. (2007) and Hobor et al. (2008) extended concurrent separation logic with dynamic creation

of locks. Their techniques might be applied here as well to support dynamic block queues.

It is also interesting to extend our logic to support multi-processor machines in the future. The general idea of memory partitions and ownership transfers used here would still apply in a multi-processor setting, except that we need to know which interrupt interrupts which processor. The implementation of kernel-level threads at the Level S in Fig. 3 becomes more complicated because it is no longer sequential, but it still prohibits interrupts at this level and can be certified based on existing work on concurrency verification. Disabling interrupts plays a less important role to bootstrap the implementation of synchronization primitives. To implement spinlocks, we need to use atomic instructions provided by the hardware, *e.g.*, the compare and swap instruction (`cas`). Also, we would like to see how relaxed memory models affect the reasoning about concurrent programs.

There are other possible extensions of the program logic to increase its expressiveness. Bornat et al. (2005) showed refinements of Separation Logic assertions to distinguish read-only access and read/write access of memory cells. The refinements can be applied to our program logic to support verification of reader/writer locks. Also, we can change the current invariant-based specifications for the well-formedness of shared memory into rely-guarantee style specifications, where assertions specify transitions of states and are more expressive than invariants (Feng et al. 2007a; Vafeiadis and Parkinson 2007).

7. Related Work and Conclusion

Regehr and Cooper (2007) 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 (2007) 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 (2002) 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 (2001) 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 (1989) certified Kit, an OS kernel implemented in machine code. Gargano et al. (2005) showed a framework for a certified OS kernel in the Verisoft project. Ni et al. (2007) 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 present a new Hoare-style framework for certifying low-level programs involving both interrupts and concurrency. Following Separation Logic, we formalize the interaction among threads and interrupt handlers in terms of memory ownership transfers. Instead of using the operational semantics of `cli`, `sti` and thread primitives, our program logic formulates their local ef-

fects over the current thread, as shown in Fig. 18, 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).

Practitioners doing informal proofs can also benefit from our logic by learning how to do informal reasoning in a systematic way for general concurrent programs, whose correctness is usually not obvious. Although the primitives shown in this paper are similar to standard routines in many OS textbooks, we are not aware of any (even informal) proofs for code that involves both hardware interrupts and preemptive concurrency. Saying that the code should work is one thing (it often still requires leap-of-faith in our experience) — knowing why it works (which this paper does) is another thing. The idea of memory partitions and ownership transfers shown in this paper (and inspired by Separation Logic) gives general guidelines even for informal proofs.

Acknowledgments

We thank anonymous referees for suggestions and comments on an earlier version of this paper. Wei Wang, Haibo Wang, and Xi Wang helped prove some of the lemmas in our Coq implementation. Xinyu Feng and Zhong Shao are supported in part by gift from Microsoft and NSF grant CCR-0524545. Yuan Dong is supported in part by National Natural Science Foundation of China (under grant No. 60573017), China Scholarship Council, and Basic Research Foundation of Tsinghua National Laboratory for Information Science and Technology (TNList). Yu Guo is supported in part by grants from National Natural Science Foundation of China (under grants No. 60673126 and No. 90718026) and Intel China Research Center. Any opinions, findings, and contributions contained in this document are those of the authors and do not reflect the views of these agencies.

References

- William R. Bevier. Kit: A study in operating system verification. *IEEE Trans. Softw. Eng.*, 15(11):1382–1396, 1989.
- Richard Bornat, Cristiano Calcagno, Peter O’Hearn, and Matthew Parkinson. Permission accounting in separation logic. In *Proc. 32nd ACM Symp. on Principles of Prog. Lang.*, pages 259–270, January 2005.
- Per Brinch Hansen. The programming language concurrent pascal. *IEEE Trans. Software Eng.*, 1(2):199–207, 1975.
- Stephen Brookes. A semantics for concurrent separation logic. In *Proc. 15th Int’l Conf. on Concurrency Theory (CONCUR’04)*, volume 3170 of *LNCS*, pages 16–34. Springer, September 2004.
- Coq Development Team. The Coq proof assistant reference manual. The Coq release v8.1, 2006.
- Rober DeLine and Manual Fähndrich. Enforcing high-level protocols in low-level software. In *Proc. 2001 ACM Conf. on Prog. Lang. Design and Impl.*, pages 59–69. ACM Press, June 2001.
- Xinyu Feng and Zhong Shao. Modular verification of concurrent assembly code with dynamic thread creation and termination. In *Proc. 2005 ACM Int’l Conf. on Functional Prog.*, pages 254–267, September 2005.
- Xinyu Feng and Zhong Shao. Local reasoning and information hiding in SCAP. Technical Report YALEU/DCS/TR-1398, Dept. of Computer Science, Yale University, New Haven, CT, February 2008. <http://flint.cs.yale.edu/publications/SCAPFrame.html>.
- Xinyu Feng, Zhong Shao, Alexander Vaynberg, Sen Xiang, and Zhaozhong Ni. Modular verification of assembly code with stack-based control abstractions. In *Proc. 2006 ACM Conf. on Prog. Lang. Design and Impl.*, pages 401–414. ACM Press, June 2006.
- Xinyu Feng, Rodrigo Ferreira, and Zhong Shao. On the relationship between concurrent separation logic and assume-guarantee reasoning. In *Proc. 16th European Symp. on Prog. (ESOP’07)*, volume 4421 of *LNCS*, pages 173–188. Springer, March 2007a.
- Xinyu Feng, Zhaozhong Ni, Zhong Shao, and Yu Guo. An open framework for foundational proof-carrying code. In *Proc. 2007 ACM Workshop on Types in Lang. Design and Impl.*, pages 67–78, January 2007b.
- Xinyu Feng, Zhong Shao, Yuan Dong, and Yu Guo. Certifying low-level programs with hardware interrupts and preemptive threads. Technical Report YALEU/DCS/TR-1396 and Coq Implementations, Dept. of Computer Science, Yale University, New Haven, CT, November 2007c. <http://flint.cs.yale.edu/publications/aim.html>.
- Mauro Gargano, Mark A. Hillebrand, Dirk Leinenbach, and Wolfgang J. Paul. On the correctness of operating system kernels. In *Proc. 18th Int’l Conf. on Theorem Proving in Higher Order Logics*, volume 3603 of *LNCS*, pages 1–16, August 2005.
- Alexey Gotsman, Josh Berdine, Byron Cook, Noam Rinetzkyy, and Mooly Sagiv. Local reasoning for storable locks and threads. In *Proc. 5th ASIAN Symp. on Prog. Lang. and Sys. (APLAS’07)*, pages 19–37, 2007.
- C. A. R. Hoare. An axiomatic basis for computer programming. *Communications of the ACM*, 26(1):53–56, October 1969.
- C. A. R. Hoare. Monitors: An operating system structuring concept. *Communications of the ACM*, 17(10):549–557, October 1974.
- C. A. R. Hoare. Towards a theory of parallel programming. In *Operating Systems Techniques*, pages 61–71. Academic Press, 1972.
- Aquinas Hobor, Andrew W. Appel, and Francesco Zappa Nardelli. Oracle semantics for concurrent separation logic. In *Proc. 17th European Symp. on Prog. (ESOP’08)*, page to appear, 2008.
- Galen C. Hunt and James R. Larus. Singularity design motivation. Technical Report MSR-TR-2004-105, Microsoft Corporation, December 2004.
- Samin S. Ishtiaq and Peter W. O’Hearn. BI as an assertion language for mutable data structures. In *Proc. 28th ACM Symp. on Principles of Prog. Lang.*, pages 14–26. ACM Press, January 2001.
- Butler W. Lampson and David D. Redell. Experience with processes and monitors in Mesa. *Commun. ACM*, 23(2):105–117, 1980.
- Andrew McCreight, Zhong Shao, Chunxiao Lin, and Long Li. A general framework for certifying garbage collectors and their mutators. In *Proc. 2007 ACM Conf. on Prog. Lang. Design and Impl.*, pages 468–479. ACM Press, June 2007.
- Greg Morrisett, David Walker, Karl Cray, and Neal Glew. From System F to typed assembly language. In *Proc. 25th ACM Symp. on Principles of Prog. Lang.*, pages 85–97. ACM Press, January 1998.
- Zhaozhong Ni, Dachuan Yu, and Zhong Shao. Using XCAP to certify realistic systems code: Machine context management. In *Proc. 20th Int’l Conf. on Theorem Proving in Higher Order Logics*, volume 4421 of *LNCS*, pages 189–206. Springer, September 2007.
- Peter W. O’Hearn. Resources, concurrency and local reasoning. In *Proc. 15th Int’l Conf. on Concurrency Theory (CONCUR’04)*, volume 3170 of *LNCS*, pages 49–67. Springer, September 2004.
- Peter W. O’Hearn, Hongseok Yang, and John C. Reynolds. Separation and information hiding. In *Proc. 31th ACM Symp. on Principles of Prog. Lang.*, pages 268–280. ACM Press, January 2004.
- Jens Palsberg and Di Ma. A typed interrupt calculus. In *Proc. 7th Int’l Symp. on Formal Tech. in Real-Time and Fault-Tolerant Sys. (FTRTFT’02)*, volume 2469 of *LNCS*, pages 291–310, September 2002.
- Wolfgang Paul, Manfred Broy, and Thomas In der Rieden. The Verisoft XT project. URL: <http://www.verisoft.de>, 2007.
- John Regehr and Nathan Coopridge. Interrupt verification via thread verification. *Electron. Notes Theor. Comput. Sci.*, 174(9), 2007.
- John C. Reynolds. Separation logic: A logic for shared mutable data structures. In *Proc. 17th Annual IEEE Symp. on Logic in Comp. Sci. (LICS’02)*, pages 55–74. IEEE Computer Society, July 2002.
- Kohei Suenaga and Naoki Kobayashi. Type based analysis of deadlock for a concurrent calculus with interrupts. In *Proc. 16th European Symp. on Prog. (ESOP’07)*, volume 4421 of *LNCS*, pages 490–504, March 2007.
- Harvey Tuch, Gerwin Klein, and Gernot Heiser. OS verification — now! In *Proc. 10th Workshop on Hot Topics in Operating Systems*, June 2005.
- Viktor Vafeiadis and Matthew Parkinson. A marriage of rely/guarantee and separation logic. In *Proc. 18th Int’l Conf. on Concurrency Theory (CONCUR’07)*, volume 4703 of *LNCS*, pages 256–271, September 2007.