

BARRIERS IN CONCURRENT SEPARATION LOGIC: NOW WITH TOOL SUPPORT!

AQUINAS HOBOR AND CRISTIAN GHERGHINA

e-mail address: {hobor, cristian}@comp.nus.edu.sg

National University of Singapore

ABSTRACT. We develop and prove sound a concurrent separation logic for Pthreads-style barriers. Although Pthreads barriers are widely used in systems, and separation logic is widely used for verification, there has not been any effort to combine the two. Unlike locks and critical sections, Pthreads barriers enable simultaneous resource redistribution between multiple threads and are inherently stateful, leading to significant complications in the design of the logic and its soundness proof. We show how our logic can be applied to a specific example program in a modular way. Our proofs are machine-checked in Coq. We showcase a program verification toolset that automatically applies the logic rules and discharges the associated proof obligations.

1. INTRODUCTION

In a shared-memory concurrent program, threads communicate via a common memory. Programmers use synchronization mechanisms, such as critical sections and locks, to avoid data races. In a data race, threads “step on each others’ toes” by using the shared memory in an unsafe manner. Recently, concurrent separation logic has been used to formally reason about shared-memory programs that use critical sections and (first-class) locks [22, 16, 14, 15]. Programs verified with concurrent separation logic are provably data-race free.

What about shared-memory programs that use other kinds of synchronization mechanisms, such as semaphores? The general assumption is that other mechanisms can be implemented with locks, and that reasonable Hoare rules can be derived by verifying their implementation. Indeed, the first published example of concurrent separation logic was implementing semaphores using critical sections [22]. Unfortunately, not all synchronization mechanisms can be easily reduced to locks in a way that allows for a reasonable Hoare rule to be derived. In this paper we introduce a Hoare rule that natively handles one such synchronization mechanism, the Pthreads-style barrier.

Pthreads (POSIX Threads) is a widely-used API for concurrent programming, and includes various procedures for thread creation/destruction and synchronization [7]. When a thread issues a barrier call it waits until a specified number (typically all) of other threads have also issued a barrier call; at that point, all of the threads continue. Although barriers do not get much attention in theory-oriented literature, they are very common in actual

Aquinas Hobor is supported by a Lee Kuan Yew Fellowship and MOE AcRF Tier 1 #T1 251RES0902.

systems code. PARSEC is the standard benchmarking suite for multicore architectures, and has thirteen workloads selected to provide a realistic cross-section for how concurrency is used in practice today; a total of five (38%) of PARSEC’s workloads use barriers, covering the application domains of financial analysis (blackscholes), computer vision (bodytrack), engineering (canneal), animation (fluidanimate), and data mining (streamcluster) [4]. A common use for barriers is to manage large numbers of threads in a pipeline setting. For example, in a video-processing algorithm, each thread might read from some shared common area containing the most recently completed frame while writing to some private area that will contain some fraction of the next frame. (A thread might need to know what is happening in other areas of the previous frame to properly handle objects entering or exiting its part of the current frame.) In the next iteration, the old private areas become the new shared common area as the algorithm continues.

Our key insight is that a barrier is used to simultaneously redistribute ownership of resources (typically, permission to read/write memory cells) between multiple threads. In the video-processing example, each thread starts out with read-only access to the previous frame and write access to a portion of the current frame. At the barrier call, each thread gives up its write access to its portion of the (just-finished) frame, and receives back read-only access to the entire frame. Separation logic (when combined with fractional permissions [5, 11]) can elegantly model this kind of resource redistribution. Let Pre_i be the preconditions held upon entering the barrier, and $Post_i$ be the postconditions that will hold after being released; then the following equation is *almost* true:

$$\bigstar_i Pre_i = \bigstar_i Post_i \quad (1.1)$$

Pipelined algorithms often operate in stages. Since barriers are used to ensure that one computation has finished before the next can start, the barriers need to have stages as well—a piece of ghost state associated with the barrier. We model this by building a finite automata into the barrier definition. We then need an assertion, written $\mathbf{barrier}(bn, \pi, cs)$, which says that barrier bn , owned with fractional permission π , is currently in state cs . The state of a barrier changes exactly as the threads are released from the barrier. We can correct equation (1.1) by noting that barrier bn is transitioning from state cs (current state) to state ns (next state), and that the other resources (frame F) are not modified:

$$\begin{aligned} \bigstar_i Pre_i &= F * \mathbf{barrier}(bn, \blacksquare, cs) \\ \bigstar_i Post_i &= F * \mathbf{barrier}(bn, \blacksquare, ns) \end{aligned} \quad (1.2)$$

We use the symbol \blacksquare to $\overset{i}{\text{denote}}$ the full ($\sim 100\%$) permission, which we require so that no thread has a “stale” view of the barrier state. Although the on-chip (or *erased*) operational behavior of a barrier is conceptually simple¹, it may be already apparent that the verification can rapidly become quite complicated.

Contributions.

- (1) We give a formal characterization for sound barrier definitions.
- (2) We design a natural Hoare rule in separation logic for verifying barrier calls.
- (3) We give a formal resource-aware *unerased* concurrent operational semantics for barriers and prove our Hoare rules sound with respect to our semantics.

¹Suspend each thread as it arrives; keep a counter of the number of arrived threads; and when all of the threads have arrived, resume the suspended threads.

- (4) Our soundness results are machine-checked in Coq and are available at:

www.comp.nus.edu.sg/~hobor/barrier

- (5) We extended a program verification toolchain to automatically apply our Hoare rules to concurrent programs using barrier synchronization and discharge the resulting proof obligations. Our prototype is available at:

www.comp.nus.edu.sg/~cristian/projects/barriers/tool.html

Relation to Previously Published Work. We previously published on the design of the program logic and its soundness proof [18]; in §8 this presentation additionally presents our work on the modifications to the HIP/SLEEK program verifier we developed to reason about our logic.

2. SYNTAX, SEPARATION ALGEBRAS, SHARES, AND ASSERTIONS

Here we briefly introduce preliminaries: the syntax of our language, separation algebras, share accounting, and the assertions of our separation logic.

2.1. Programming Language Syntax. To let us focus on the barriers, most of our programming language is pure vanilla. We define four kinds of (tagged) values v : **TRUE**, **FALSE**, **ADDR**(\mathbb{N}), and **DATA**(\mathbb{N}). We have two (tagged) expressions e : $\mathbb{C}(v)$ and $\mathbb{V}(x)$, where x are local variable names (just \mathbb{N} in Coq). To make the example more interesting we add the arithmetical operations to e . We write **bn** for a barrier number, with $\text{bn} \in \mathbb{N}$.

We have ten commands c : **skip** (do nothing), $x := e$ (local variable assignment), $x := [e]$ (load from memory), $[e_1] := e_2$ (store to memory), $x := \text{new } e$ (memory allocation), **free** e (memory deallocation), $c_1; c_2$ (instruction sequence), **if** e **then** c_1 **else** c_2 (if-then-else), **while** e $\{c\}$ (loops), and **barrier** **bn** (wait for barrier **bn**). To run commands $c_1 \dots c_n$ in parallel (which, like O’Hearn, we only allow at the top level [22]), we write $c_1 || \dots || c_n$. To avoid clogging the presentation, we elide a setup sequence before the parallel composition.

2.2. Disjoint Multi-unit Separation Algebras. Separation algebras are mathematical structures used to model separation logic. We use a variant described by Dockins *et al.* called a disjoint multi-unit separation algebra (hereafter just “DSA”) [11]. Briefly, a DSA is a set S and an associated three-place partial *join relation* \oplus , written $x \oplus y = z$, such that:

A function:	$x \oplus y = z_1 \Rightarrow x \oplus y = z_2 \Rightarrow z_1 = z_2$
Commutative:	$x \oplus y = y \oplus x$
Associative:	$x \oplus (y \oplus z) = (x \oplus y) \oplus z$
Cancellative:	$x_1 \oplus y = z \Rightarrow x_2 \oplus y = z \Rightarrow x_1 = x_2$
Multiple units:	$\forall x. \exists u_x. x \oplus u_x = x$
Disjointness:	$x \oplus x = y \Rightarrow x = y$

A key concept is the idea of an *identity*: x is an identity if $x \oplus y = z$ implies $y = z$. One fundamental property of identities is that x is an identity if and only if $x \oplus x = x$. Dockins also develops a series of standard constructions (*e.g.*, product, functions, etc.) for building complicated DSAs from simpler DSAs. We make use of this idea to construct a variety of separation algebras as needed, usually with the concept of *share* as the “foundational” DSA.

2.3. Shares. Separation logic is a logic of *resource ownership*. Concurrent algorithms sometimes want to have threads share some common resources. Bornat *et al.* introduced the concept of *fractional share* to handle the necessary accounting [5]. Shares form a DSA; a *full share* (complete ownership of a resource) can be broken into various *partial shares*; these shares can then be rejoined into the full share. The *empty share* is the identity for shares. We often need non-empty (strictly *positive*) shares, denoted by π . A critical invariant is that the sum of each thread’s share of a given object is no more or less than the full share.

The semantic meaning of partial shares varies; here we use them in two distinct ways. We require the full share to modify a memory location; in contrast, we only require a positive share to read from one. There is no danger of a data race even though we do not require the full share to read: if a thread has a positive share of some location, no other thread can have a full share for the same location. We use fractional permissions differently for barriers: each precondition includes some positive share of the barrier itself and we require that the preconditions combine to imply the full share of the barrier (plus a frame F).

In the Coq development we use a share model developed Dockins *et al.* that supports sophisticated fractional ownership schemes [11]. Here we simplify this model into four elements: the full share \blacksquare ; two **distinct** nonempty partial shares, \blacktriangleleft and \blacktriangleright , and the empty share \square . The key point is that $\blacktriangleleft \oplus \blacktriangleright = \blacksquare$.

2.4. Assertion Language. We model the assertions of separation logic following Dockins *et al.* [11]. Our states σ are triples of a stack, heap, and barrier map ($\sigma = (s, h, b)$). Local variables live in stacks s (functions from variable names to values). In contrast, a heap h contains the locations shared between threads; heaps are partial functions from addresses to pairs of positive shares and values. We also equip our heaps with a distinguished location, called the *break*, that tracks the boundary between allocated and unallocated locations. The break lets us provide semantics for the $x := \text{new } e$ instruction in a natural way by setting x equal to the current break and then incrementing the break. Since threads share a common break, there is a backdoor communication channel; however the existence of this channel is a small price to pay for avoiding the necessity of a concurrent garbage collector. We ensure that the threads see the same break by equipping our break with ownership shares just as we equip normal memory locations with shares.

We denote the empty heap (which lacks ownership for both all memory locations and the distinguished break location) by h_0 . Of note, our expressions e are evaluated only in the context of the stack; we write $s \vdash e \Downarrow v$ to mean that e evaluates to v in the context of the stack s . Finally, the barrier map b is a partial function from barrier numbers to pairs of barrier states (represented as natural numbers) and positive shares; we denote the empty barrier map by b_0 .

An *assertion* is a function from states to truth values (**Prop** in Coq). As is common, we define the usual logical connectives via a straightforward embedding into the metalogic; for example, the object-level conjunction $P \wedge Q$ is defined as $\lambda\sigma. (P\sigma) \wedge (Q\sigma)$. We will adopt the convention of using the same symbol for both the object-level operators and the meta-level operators to avoid symbol bloat; it should be clear from the context which operator applies in a given situation. We provide all of the standard connectives ($\top, \perp, \wedge, \vee, \Rightarrow, \neg, \forall, \exists$).

We model the connectives of separation logic in the standard way²:

$$\begin{aligned}
\mathbf{emp} &= \lambda(s, h, b). h = h_0 \wedge b = b_0 \\
P * Q &= \lambda\sigma. \exists\sigma_1, \sigma_2. \sigma_1 \oplus \sigma_2 = \sigma \wedge P(\sigma_1) \wedge Q(\sigma_2) \\
e_1 \overset{\pi}{\mapsto} e_2 &= \lambda(s, h, b). \exists a, v. (s \vdash e_1 \Downarrow \mathbf{ADDR}(a)) \wedge (s \vdash e_2 \Downarrow v) \wedge \\
&\quad b = b_0 \wedge h(a) = (v, \pi) \wedge \mathbf{dom}(h) = \{a\} \wedge \mathbf{break}(h) = \square \\
\mathbf{barrier}(bn, \pi, s) &= \lambda(s, h, b). h = h_0 \wedge b(bn) = (s, \pi) \wedge \mathbf{dom}(b) = \{bn\}
\end{aligned}$$

The fractional maps-to assertion, $e_1 \overset{\pi}{\mapsto} e_2$, means that the expression e_1 is pointing to an address a in memory; a is owned with positive share π , and contains the evaluated value v of e_2 . The fractional maps-to assertion does not include any ownership of the break. The barrier assertion, $\mathbf{barrier}(bn, \pi, s)$, means that the barrier bn , owned with positive share π , is in state s .

We also lift program expressions into the logic: $e \Downarrow v$, which evaluates e with σ 's stack (*i.e.*, $\lambda(s, h, b). h = h_0 \wedge b = b_0 \wedge s \vdash e \Downarrow v$); $[e]$, equivalent to $e \Downarrow \mathbf{TRUE}$; and $x = v$, equivalent to $\mathbf{V}(x) \Downarrow v$. These assertions have a “built-in” \mathbf{emp} .

3. EXAMPLE

We present a detailed example inspired by a video decompression algorithm. The code and a detailed-but-informal description of the barrier definition is given in Figure 1.³ Two threads cooperate to repeatedly compute the elements of two size-two arrays x and y . In each iteration, each thread writes to a single cell of the “current” array, and reads from both cells of the “previous” array.

In Figure 1 we give a pictorial representation of the state machine associated with the barrier used in the code using the following specialized notation:

$$\begin{array}{c}
\mathbf{x}_1 \quad \mathbf{i} \quad \mathbf{b}\text{-state} \\
\begin{array}{l}
\mathbf{Mv}_1 \quad \begin{array}{|c|c|} \hline \mathbf{A} & \mathbf{T} \\ \hline \end{array} \quad \begin{array}{|c|} \hline \mathbf{1} \\ \hline \end{array} \quad \wedge \mathbf{T} \geq 30 \quad \equiv \quad \exists \mathbf{A}, \mathbf{T}, \mathbf{x}_1 \overset{\blacksquare}{\rightarrow} \mathbf{A} * \mathbf{i} \overset{\blacksquare}{\rightarrow} \mathbf{T} * \mathbf{barrier}(\mathbf{b}, \blacksquare, \mathbf{1}) \quad \wedge \mathbf{T} \geq 30 \\
\mathbf{Mv}_2 \quad \begin{array}{|c|c|} \hline \mathbf{A} & \mathbf{T} \\ \hline \end{array} \quad \begin{array}{|c|} \hline \mathbf{1} \\ \hline \end{array} \quad \wedge \mathbf{T} \geq 30 \quad \equiv \quad \exists \mathbf{A}, \mathbf{T}, \mathbf{x}_1 \rightarrow \mathbf{A} * \mathbf{i} \rightarrow \mathbf{T} * \mathbf{barrier}(\mathbf{b}, \blacksquare, \mathbf{1}) \quad \wedge \mathbf{T} \geq 30 \\
\Downarrow \\
\mathbf{Mv}_1 \quad \begin{array}{|c|c|} \hline \mathbf{A} & \blacksquare \\ \hline \end{array} \quad \begin{array}{|c|} \hline \mathbf{3} \\ \hline \end{array} \quad \equiv \quad \exists \mathbf{A} \quad \mathbf{x}_1 \overset{\blacksquare}{\rightarrow} \mathbf{A} \quad * \mathbf{barrier}(\mathbf{b}, \blacksquare, \mathbf{3}) \\
\mathbf{Mv}_2 \quad \begin{array}{|c|c|} \hline \mathbf{A} & \mathbf{T} \\ \hline \end{array} \quad \begin{array}{|c|} \hline \mathbf{3} \\ \hline \end{array} \quad \wedge \mathbf{T} \geq 30 \quad \equiv \quad \exists \mathbf{A}, \mathbf{T}, \mathbf{x}_1 \rightarrow \mathbf{A} * \mathbf{i} \rightarrow \mathbf{T} * \mathbf{barrier}(\mathbf{b}, \blacksquare, \mathbf{3}) \quad \wedge \mathbf{T} \geq 30
\end{array}
\end{array}$$

This notation is used to express the pre- and postconditions for a given barrier transition. Each row is a pictorial representation (values, barrier states, and shares) of a formula in separation logic as indicated above. The preconditions are on top (one per row) and the postconditions below. Each row is associated with a *move*; move 1 is a pair of the first precondition row and the first postcondition row, etc. A barrier that is waiting for n threads will have n moves; n can be fewer than the total number of threads. We do not require that a given thread always takes the same move each time it reaches a given barrier transition.

²Our Coq definition for \mathbf{emp} is different but equivalent to the definition given here.

³In our Coq development we give the full formal description of the example barrier.

<pre> 0: {x1[■]→0 * x2[■]→0 * y1[■]→0 * y2[■]→0 * i[■]→0 * barrier(bn, ■, 0)} 0': {x1[■]→0 * x2[■]→0 * y1[■]→0 * y2[■]→0 * i[■]→0 * barrier(bn, ■, 0)} ... 1: barrier b; 2: n := 0; 3: while n < 30 { 4: a1 := [x1]; 5: a2 := [x2]; 6: [y1] := (a1+2*a2); 7: barrier b; 8: a1 := [y1]; 9: a2 := [y2]; 10: [x1] := (a1+2*a2); 11: n := (n+1); 12: [i] := n; 13: barrier b; 14: } 15: } 16: barrier b; 17: [i] := 0; ... </pre>	<pre> { x1[■]→0 * x2[■]→0 * y1[■]→0 * y2[■]→0 * i[■]→0 * barrier(bn, ■, 0) } ... barrier b; // b transitions 0→1 m := 0; while m < 30 { a1 := [x1]; a2 := [x2]; [y2] := (a1+3*a2); barrier b; // b transitions 1→2 a1 := [y1]; a2 := [y2]; [x2] := (a1+3*a2); barrier b; // b transitions 2→1 m := [i]; } barrier b; // b transitions 1→3 ... </pre>
----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

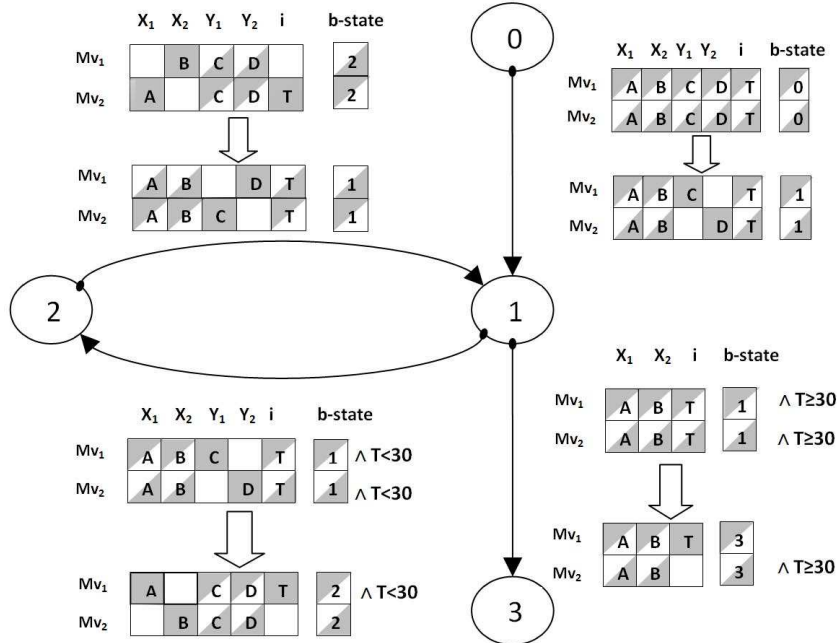


Figure 1: Example: Code and Barrier Diagram

Note that only the permissions on the memory cells change during a transition; the contents (values) do not.⁴ The exception to this is the special column on the right side, which denotes the assertion associated with the barrier itself. As the barrier transitions, this value changes from the previous state to the next; we require that the sum of the preconditions includes the full share of the barrier assertion to guarantee that no thread has an out-of-date view of the barrier's state. Observe that all of the preconditions join together, and, except for the state of the barrier itself, are exactly equal to the join of the postconditions.

The initial state of the machine is given as an assertion in line 0. The machine starts with full ownership of the array cells x_1 , x_2 , y_1 , and y_2 , as well as an additional cell i , used as a condition variable. The barrier b is fully-owned and is in state 0. The initial state is then partitioned into two parts on line 0', with the left thread (A) and right thread (B) getting the shares $\frac{1}{2}$ and $\frac{1}{2}$, respectively.

Not shown (between lines 0' and 1) is thread-specific initialization code; perhaps both threads read both arrays and perform consistency checks. The real action starts with the barrier call on line 1, which ensures that this initialization code has completed. Thread A takes move 1 and thread B takes move 2. Afterwards, thread A has full ownership over y_1 and thread B has full ownership over y_2 ; the ownership of x_1 , x_2 , and i remains split between A and B. While the ownership of the barrier is unchanged, it is now in state 1.

We then enter the main loop on line 3. On lines 4–5, both threads read from the shared cells x_1 and x_2 , and on line 6 both threads update their fully-owned cell. The barrier call on line 7 ensures that these updates have been completed before the threads continue. Since the value T at memory location i is less than 30, only the 1–2 transition is possible; the 1–3 transition requires $T \geq 30$. Thread A takes move 1 and thread B takes move 2⁵; afterwards, both threads have partial shares of y_1 and y_2 , thread A has the full share of x_1 and the condition cell i , and thread B has the full share of x_2 ; the barrier is in state 2.

Lines 8–10 are mirrors of lines 4–6. On lines 11–12, thread A updates the condition cell i . The barrier on line 13 ensures that the updates on lines 10 and 12 have completed before the threads continue; thread A takes move 2 while thread B takes move 1. Afterwards, the threads have the same permissions they had on entering the loop: A has full ownership of y_1 , B has full ownership of y_2 , and they share ownership of x_1 , x_2 , and i ; the barrier is again in state 1.

On line 14, thread B reads from the condition variable i , and then the program loops back to line 3. After 30 iterations, the loop exits and control moves to the barrier on line 16. Observe that since the (shared) value T at memory location i is greater than or equal to 30, only the 1–3 transition is possible; the 1–2 transition requires $T < 30$. Thread A takes move 1 while thread B takes move 2; afterwards, both threads are sharing ownership of x_1 , x_2 , y_1 , and y_2 (since the transition from 1 to 3 does not mention y_1 and y_2 they are unchanged). Thread A has full permission over the condition variable i ; the barrier is in state 3. Finally, on line 17, thread A updates i ; the barrier on line 16 ensures that thread B's read of i on line 14 has already occurred.

⁴We use the same quantified variable names before and after the transition because an outside observer can tell that the values are the same. A local verification can use ghost state to prove the equality; alternatively we could add the ability to move the quantifier to other parts of the diagram, *e.g.*, over an entire pre-post pair.

⁵In this example a given thread always takes the same move for a given transition; however, this is not forced by the rules of our logic.

BarDef (barrier definition)	≡	{ bd_bn : Nat bd_limit : Nat bd_states : list BarStateDef }	barrier id # of threads state list
BarStateDef (barrier state)	≡	{ bsd_bn : Nat bsd_cs : Nat bsd_directions : list BarMoveList bsd_limit : Nat }	barrier id state id transition list # of threads
BarMoveList (transition)	≡	{ bml_dest : Nat bml_bn : Nat bml_cs : Nat bml_limit : Nat bml_moves : list (assert × assert) }	next state barrier id current state # of threads pre/post pairs

Figure 2: Barrier Definitions

4. BARRIER DEFINITIONS AND CONSISTENCY REQUIREMENTS

We present the type of a barrier definition in Figure 2 in the form of a data structure. The definitions include numerous consistency requirements; in Coq these are maintained with dependent types. From the top down, a barrier definition (`BarDef`) consists of a barrier identifier (*i.e.*, barrier number), the number of threads the barrier is synchronizing, and a list of barrier state definitions. For programs that have more than one barrier, the individual barrier definitions will be collected into a list and barrier number j will be in list slot j .

A barrier state definition (`BarStateDef`) consists of a barrier number, the number of threads synchronized, a state id, and a transition list; such that:

- (1) the barrier number matches the barrier number in the containing `BarDef`
- (2) the limit matches the limit of the containing `BarDef`⁶
- (3) the state identifier j indicates that this `BarStateDef` is the j element of the containing `BarDef`'s list of state definitions
- (4) the directions are *mutually exclusive*

The first three are unexciting; we will discuss mutual exclusion shortly.

A transition (`BarMoveList`) contains a barrier number (`bn`), number of threads synchronized, current state identifier (`cs`), next state identifier (`ns`), and list of precondition/postcondition pairs (the *move list*). We require that:

- (1) `bn` matches the barrier number in the containing `BarStateDef`
- (2) the limit matches the limit in the containing `BarStateDef`
- (3) `cs` matches the state identifier in the containing `BarStateDef`
- (4) the length of list of moves (`bml_moves`) is equal to the limit (`bml_limit`)
- (5) all of the pre/postconditions in the movelist ignore the stack, focusing only on the memory and barrier map. Since stacks are private to each thread (on a processor these would be registers), it does not make sense for them to be mentioned in the “public” pre/post conditions.

⁶A command to dynamically alter the number of threads a barrier managed might allow different states/transitions to wait for different numbers of threads.

- (6) all of the preconditions in the movelist are *precise*. Precision is a technical property involving the identifiability of states satisfying an assertion.⁷
- (7) each precondition P includes some positive share of the barrier assertion with \mathbf{bn} and \mathbf{cs} , *i.e.*, $\exists\pi. P \Rightarrow \top * \mathbf{barrier}(\mathbf{bn}, \pi, \mathbf{cs})$.
- (8) the sum of the preconditions must equal the sum of the postconditions, except for the state of the barrier; moreover, the sum of the preconditions must include the full share of the barrier (equation (1.2), repeated here):

$$\begin{aligned} * \underset{i}{Pre}_i &= F * \mathbf{barrier}(\mathbf{bn}, \blacksquare, \mathbf{cs}) \\ * \underset{i}{Post}_i &= F * \mathbf{barrier}(\mathbf{bn}, \blacksquare, \mathbf{ns}) \end{aligned}$$

Items 1–4 are simple bookkeeping; items 5–7 are similar to technical requirements required in other variants of concurrent separation logic [22, 15, 14]. As previously mentioned, the fundamental insight of this approach is property (8).

The function `lookup_move` simplifies the lookup of a move in a `BarDef`:

$$\mathbf{lookup_move}(bd, cs, dir, mv) = bd.\mathbf{bd_states}[cs].\mathbf{bsd_directions}[dir].\mathbf{bml_moves}[mv]$$

Using this notation, we can express the important requirement that all directions in the barrier state cs of the barrier definition \mathbf{bd} are mutually exclusive:

$$\begin{aligned} \forall dir_1, dir_2, mv_1, mv_2, pre_1, pre_2. dir_1 \neq dir_2 \Rightarrow \\ \mathbf{lookup_move}(bd, cs, dir_1, mv_1) = (pre_1, _) \Rightarrow \\ \mathbf{lookup_move}(bd, cs, dir_2, mv_2) = (pre_2, _) \Rightarrow \\ (\top * pre_1) \wedge (\top * pre_2) \equiv \perp \end{aligned}$$

In other words, it is *impossible* for any of the preconditions of more than one transition (of a given state) to be true at a time. The simplest way to understand this is to consider the 1–2 and 1–3 transitions in the example program. The 1–2 transition requires that the value in memory cell i be strictly less than 30; in contrast, the 1–3 transition requires that *the same cell* contains a value greater than or equal to 30. Plainly these are incompatible; but in fact the above property is stronger: *both* of the moves on the 1–2 transition, and *both* of the moves on the 1–3 transition include the incompatibility. Thus, if thread A takes transition 1–2, it knows for certain that thread B *cannot* take transition 1–3. This way we ensure that both threads always agree on the barrier’s current state.

5. HOARE LOGIC

Our Hoare judgment has the form $\Gamma \vdash \{P\} c \{Q\}$, where Γ is a list of barrier definitions as given in §4, P and Q are assertions in separation logic, and c is a command. Our Hoare rules come in three groups: standard Hoare logic (Skip, If, Sequence, While, Assignment, Consequence); standard separation logic (Frame, Store, Load, New, Free); and the barrier rule. We give groups two and three in Figure 3; group one is standard and elided. We note four points for group two.

First, as explained in §2.4, the assertions $e \Downarrow v$, $[e]$ and $x = v$ are bundled with an assertion that the heap and barrier map are empty (*i.e.*, $e \Downarrow v \Rightarrow \mathbf{emp}$); thus, we use the separating conjunction when employing them. Second, the rules are in “side-condition-free

⁷Precision may not be required; another property (tentatively christened “token”) that might serve would be if, for any precondition P , $P * P \equiv \perp$. Note that precision in conjunction with item (6) implies P is a token.

$$\begin{array}{c}
\frac{}{\Gamma \vdash \{P\} \text{ skip } \{P\}} \text{Skip} \quad \frac{\Gamma \vdash \{P * [e]\} c_t \{Q\} \quad \Gamma \vdash \{P * \neg[e]\} c_f \{Q\}}{\Gamma \vdash \{P\} \text{ if } e \text{ then } c_t \text{ else } c_f \{Q\}} \text{If} \\
\\
\frac{\Gamma \vdash \{P\} c_1 \{Q\} \quad \Gamma \vdash \{Q\} c_2 \{R\}}{\Gamma \vdash \{P\} c_1 ; c_2 \{R\}} \text{Seq.} \quad \frac{\Gamma \vdash \{I * [e]\} c \{I\}}{\Gamma \vdash \{I\} \text{ while } e \{c\} \{I * \neg[e]\}} \text{While} \\
\\
\frac{}{\Gamma \vdash \{e \Downarrow v\} x := e \{x = v\}} \text{Assign} \quad \frac{P' \vdash P \quad \Gamma \vdash \{P\} c \{Q\} \quad Q \vdash Q'}{\Gamma \vdash \{P'\} c \{Q'\}} \text{Conseq.} \\
\\
\frac{\Gamma \vdash \{P\} c \{Q\} \text{ closed}(F, c)}{\Gamma \vdash \{F * P\} c \{F * Q\}} \text{Frame} \quad \frac{}{\Gamma \vdash \{e_1 \overset{\blacksquare}{\mapsto} _ \} [e_1] := e_2 \{e_1 \overset{\blacksquare}{\mapsto} e_2\}} \text{Store} \\
\\
\frac{}{\Gamma \vdash \{e_1 \overset{\pi}{\mapsto} e_2 * e_1 \Downarrow v_1 * e_2 \Downarrow v_2\} x := [e_1] \{C(v_1) \overset{\pi}{\mapsto} C(v_2) * x = v_2\}} \text{Load} \\
\\
\frac{}{\Gamma \vdash \{e \Downarrow v\} x := \text{new } e \{V(x) \overset{\blacksquare}{\mapsto} C(v)\}} \text{New} \quad \frac{}{\Gamma \vdash \{e_1 \overset{\blacksquare}{\mapsto} e_2\} \text{ free } e_1 \{\text{emp}\}} \text{Free} \\
\\
\boxed{\frac{\Gamma[bn] = bd \quad \text{lookup_move}(bd, cs, dir, mv) = (P, Q)}{\Gamma \vdash \{P\} \text{ barrier bn } \{Q\}} \text{Barrier}}
\end{array}$$

Figure 3: Hoare rules

form". Thus, instead of presenting the load rule as $\Gamma \vdash \{e_1 \overset{\pi}{\mapsto} e_2\} x := [e_1] \{x = e_2 * e_1 \overset{\pi}{\mapsto} e_2\}$, which is aesthetically attractive but untrue in the pesky case when e_2 depends on x (e.g., $x := [x]$), we use a form that is less visually pleasing but does not require side conditions.⁸ It is straightforward to restore rules with side conditions via the Consequence rule. Third, our Store and Free rules require the full share of location e_1 ; in contrast, our Load rule only requires some positive share; this is consistent with our use of fractional permissions as explained in §2.3. Fourth, memory allocation and deallocation are more complicated in concurrent settings than in sequential settings, and so the New and Free rules cause nontrivial complications in the semantic model.

The Hoare rule for barriers is so simple that at first glance it may be hard to understand. The variables for the current state cs , direction dir , and move mv appear to be free in the `lookup_move`! However, things are not quite as unconstrained as they initially appear. Recall from §4 that one of the consistency requirements for the precondition P is that P implies an assertion about the barrier itself: $P \Rightarrow Q * \text{barrier}(bn, \pi, cs)$; thus at a given program point we can only use directions and moves from the current state. Similarly, recall from §4 that since the directions are mutually exclusive, dir is uniquely determined.

This leaves the question of the uniqueness of mv . If a thread only satisfies a single precondition, then the move mv is uniquely determined. Unfortunately, it is simple to construct programs in which a thread enters a barrier while satisfying the preconditions of multiple moves. What saves us is that we are developing a logic of partial correctness. Since

⁸Recall from §2: $V(x)$ and $C(v)$ are expression constructors for locals and constants. In addition, $\text{closed}(F, c)$ means that F does not depend on locals modified by c .

preconditions to moves must be precise and nonempty (*i.e.*, token), only one thread is able to satisfy a given precondition at a time. The pigeonhole principle guarantees that if a thread holds multiple preconditions then some other thread will not be able to enter the barrier; in this case, the barrier call will never return and we can guarantee any postcondition.

We now apply the Barrier rule to the barrier calls in line 13 from our example program; the `lookup_moves` are direct from the barrier state diagram:

$$\begin{array}{l}
 \text{Thread A} \left\{ \begin{array}{l}
 \text{lookup_move}(b, 2, 1, 2) = (P, Q) \\
 P = y_1 \overset{\blacksquare}{\mapsto} v_{y_1} * y_2 \overset{\blacksquare}{\mapsto} v_{y_2} * x_1 \overset{\blacksquare}{\mapsto} v_{x_1} * i \overset{\blacksquare}{\mapsto} v_i * \text{barrier}(bn, \blacksquare, 2) \\
 Q = y_1 \overset{\blacksquare}{\mapsto} v_{y_1} * x_1 \overset{\blacksquare}{\mapsto} v_{x_1} * x_2 \overset{\blacksquare}{\mapsto} v_{x_2} * i \overset{\blacksquare}{\mapsto} v_i * \text{barrier}(bn, \blacksquare, 1)
 \end{array} \right. \\
 \hline
 \Gamma \vdash \{P\} \text{ barrier } \mathbf{b} \{Q\} \\
 \text{Thread B} \left\{ \begin{array}{l}
 \text{lookup_move}(b, 2, 1, 1) = (P, Q) \\
 P = y_1 \overset{\blacksquare}{\mapsto} v_{y_1} * y_2 \overset{\blacksquare}{\mapsto} v_{y_2} * x_2 \overset{\blacksquare}{\mapsto} v_{x_2} * \text{barrier}(bn, \blacksquare, 2) \\
 Q = y_2 \overset{\blacksquare}{\mapsto} v_{y_2} * x_1 \overset{\blacksquare}{\mapsto} v_{x_1} * x_2 \overset{\blacksquare}{\mapsto} v_{x_2} * i \overset{\blacksquare}{\mapsto} v_i * \text{barrier}(bn, \blacksquare, 1)
 \end{array} \right. \\
 \hline
 \Gamma \vdash \{P\} \text{ barrier } \mathbf{b} \{Q\}
 \end{array}$$

Note that in this line of the example program, the frame is `emp` in both threads.

Not shown in Figure 3 is a parallel composition rule. As in [15], each thread is verified independently using the Hoare rules given; a top-level safety theorem proves that the entire concurrent machine behaves as expected.

6. SEMANTIC MODELS

Our operational semantics is divided into three parts: purely sequential, which executes all of the instructions except for `barrier` in a thread-local manner; concurrent, which manages thread scheduling and handles the barrier instruction; and oracular, which provides a pseudo-sequential view of the concurrent machine to enable simple proofs of the sequential Hoare rules. Our setup follows Hobor *et al.* very closely and we refer readers there for more detail [16, 15].

Purely sequential semantics. The purely sequential semantics executes the instructions `skip`, `x := e`, `x := [e]`, `[e1] := e2`, `x := new e`, `free e`, `c1; c2`, `if e then c1 else c2`, and `while e {c}`. The form of the sequential step judgment is $(\sigma, c) \mapsto (\sigma', c')$. Here σ is a state (triple of stack, heap, barrier map), just as in §2.4 and c is a command of our language. The semantics of the sequential instructions is standard; the only “tricky” part is that the machine gets stuck if one tries to write to a location for which one does not have full permission or read from a location for which one has no permission; *e.g.*, here is the store rule:

$$\frac{s \vdash e_1 \Downarrow \mathbf{C}(\text{ADDR}(n)) \quad s \vdash e_2 \Downarrow v \quad n < \text{break}(h) \quad h(n) = (\blacksquare, v') \quad h' = [n \mapsto (\blacksquare, v)]h}{((s, h, b), [e_1] := e_2; c) \mapsto ((s, h', b), c)} \text{ sstep - store}$$

The test that $n < \text{break}(h)$ ensures that the address for the store is “in bounds”—that is, less than the current value of the break between allocated and unallocated memory; since we are updating the memory we require that the permission associated with the location

n full (■). We say that this step relation is *unerased* since these bounds and permission checks are virtual rather than on-chip.

We define the other cases of the step relation in a similar way. Observe that if we were in a sequential setting the proof of the Hoare store rule would be straightforward; this is likewise the case for the other cases of the sequential step relation and their associated Hoare rules. If the sequential step relation reaches a barrier call **barrier** bn then it simply gets stuck.

Concurrent semantics. We define the notion of a *concurrent state* in Figure 4. A concurrent state contains a scheduler Ω (modeled as a list of natural numbers), a distinguished heap called the *allocation pool*, a list of *threads*, and a *barrier pool*⁹. The allocation pool is the owner of all of the unallocated memory cells (plus the ownership of the break between allocated and unallocated cells); before we run a thread we transfer the allocation pool into the local heap owned by the thread so that **new** can transfer a cell from this pool into the local heap of a thread when required. When we suspend the thread we remove (what is left of) the allocation pool from its heap so that we can transfer it to the next thread.

A thread contains a (sequential) state (stack, heap, and barrier map) and a *concurrent control*, which is either **Running**(c), meaning the thread is available to run command c , or **Waiting**(bn, dir, mv, c), meaning that the thread is currently waiting on barrier bn to make move mv in direction dir ; after the barrier call completes the thread will resume running with command c .

The barrier pool (**Barpool**) contains a list of *dynamic barrier statuses* (DBSes) as well as a state which is the join of all of the states inside the DBSes. Each DBS consists of a barrier number (which must be its index into the array of its containing **Barpool**), a barrier definition (from §4), and a *waitpool* (WP). A waitpool consists of a direction option (**None** before the first barrier call in a given state; thereafter the unique direction for the next state), a limit (the number of threads synchronized by the barrier, and comes from the barrier definition in the enclosing DBS), a *slot* list, and a state (which is the join of all of the states in the slot list). A slot is a heap and barrier map (the stack is unneeded since barrier pre/postconditions ignore it) as well as a thread id (whence the heap and barrier map came as a precondition, and to which the postcondition will return).

The concurrent step relation has the form $(\Omega, ap, thds, bp) \rightsquigarrow (\Omega', ap', thds', bp')$, where $\Omega, ap, thds$, and bp are the scheduler, allocation pool, thread list, and barrier pool respectively. The concurrent step relation has only four cases; the following case CStep-Seq is used to run all of the sequential commands:

$$\frac{\begin{array}{l} thds[i] = (s, h, b, \text{Running}(c)) \quad h \oplus ap = h' \quad ((s, h', b), c) \mapsto ((s', h'', b), c') \\ h''' \oplus ap' = h'' \quad \text{isAllocPool}(ap') \quad thds' = [i \mapsto (s', h''', b, \text{Running}(c'))] thds \end{array}}{(i :: \Omega, thds, ap, bp) \rightsquigarrow (i :: \Omega, thds', ap', bp)} \quad \text{CStep-Seq}$$

That is, we look up the thread whose thread id is at the head of the scheduler, join in the allocation pool, and run the sequential step relation. If the command c is a barrier call then the sequential relation will not be able to run and so the CStep-Seq relation will not hold; otherwise the sequential step relation will be able to handle any command. After

⁹There is also a series of consistency requirements such as the fact that all of the heaps in the threads and barrier pool join together with the allocation pool into one consistent heap; in the mechanization this is carried around via a dependent type as a fifth component of the concurrent state. We elide this proof from the presentation.

Cstate	≡	{ cs_sched : list \mathbb{N} cs_allocpool : heap cs_thds : list Thread cs_barpool : Barpool }	schedule alloc pool thread pool barrier pool
Thread	≡	{ th_stk : stack th_hp : heap th_bs : BarrierMap th_ctl : conc_ctl }	local view of barrier states running or waiting
conc_ctl	≡	Running(<i>c</i>) Waiting(<i>bn, dir, mv, c</i>)	executing code <i>c</i> waiting on <i>bn</i>
Barpool	≡	{ bp_bars : list DyBarStatus bp_st : stack \times heap \times BarrierMap }	dynamic barrier status current state
DyBarStatus	≡	{ dbs_bn : \mathbb{N} dbs_wp : Waitpool dbs_bd : BarDef }	barrier id waiting thread pool
Waitpool	≡	{ wp_dir : \mathbb{N} <i>option</i> wp_slots : list slot <i>option</i> wp_limit : \mathbb{N} wp_st : stack \times heap \times BarrierMap }	direction id taken slots current state
slot	≡	(thread_id \times heap \times BarrierMap)	waiting slot

Figure 4: Concurrent state

we have taken a sequential step, we subtract out the (possibly diminished) allocation pool, and reinsert the modified sequential state into the thread list. Since we quantify over all schedulers and our language does not have input/output, it is sufficient to utilize a non-preemptive scheduler; for further justification on the use of such schedulers see [15].

The second case of the concurrent step relation handles the case when a thread has reached the last instruction, which must be a `skip`:

$$\frac{thds[i] = \text{Running}(\text{skip})}{(i :: \Omega, thds, ap, bp) \rightsquigarrow (\Omega, thds, ap, bp)} \text{CStep-Exit}$$

When we reach the end of a thread we simply context switch to the next thread.

The interesting cases occur when the instruction for the running thread is a barrier call; here the CStep-Seq rule does not apply. The concurrent semantics handles the barrier call directly via the last two cases of the step relation; before presenting these cases we will first give a technical definition called `fill_barrier_slot`:

$$\frac{\begin{array}{l} thds[i] = \text{Thread}(stk, hp, bs, (\text{Running}(\text{barrier } bn ; c))) \\ \text{lookup_move}(bp.bp_bars[bn], dir, mv) = (pre, post) \\ hp' \oplus hp'' = hp \quad bs' \oplus bs'' = bs \quad pre(stk, hp', bs') \\ \text{bp_inc_waitpool}(bp, bn, dir, mv, (i, (hp', bs'))) = bp' \\ thds' = [i \rightarrow (\text{Thread}(stk, hp'', bs'', (\text{Waiting}(bn, dir, mv, c)))] thds \end{array}}{\text{fill_barrier_slot}(thds, bp, bn, i) = (thds', bp')}$$

The predicate `fill_barrier_slot` gives the details of removing the (sub)state satisfying the precondition of the barrier from the thread’s state, inserting it into the barrier pool, and suspending the calling thread. The predicate `bp_inc_waitpool` does the insertion into the barrier pool; the details of manipulating the data structure are straightforward but lengthy to formalize¹⁰.

We are now ready to give the first case for the barrier, used when a thread executes a barrier but is not the last thread to do so:

$$\frac{\text{fill_barrier_slot } (thds, bp, bn, i) = (thds', bp') \quad \neg \text{bp_ready } (bp', bn)}{((i :: \Omega), ap, thds, bp) \rightsquigarrow (\Omega, ap, thds', bp')} \text{ CStep-Suspend}$$

After using `fill_barrier_slot`, `CStep-Suspend` checks to see if the barrier is full by counting the number of slots that have been filled in the appropriate wait pool by using the `bp_ready` predicate, and then context switches.

If the barrier is ready then instead of using the `CStep-Suspend` case of the concurrent step relation, we must use the `CStep-Release` case:

$$\frac{\text{fill_barrier_slot } (thds, bp, bn, i) = (thds', bp') \quad \text{bp_ready } (bp', bn) \quad \text{bp_transition } (bp', bn, out) = bp'' \quad \text{transition_threads } (out, thds') = thds''}{((i :: \Omega), ap, thds, bp) \rightsquigarrow (\Omega, ap, thds'', bp'')} \text{ CStep-Release}$$

The first requirement of `CStep-Release` is exactly the same as `CStep-Suspend`: we suspend the thread and transfer the appropriate resources to the barrier pool. However, now all of the threads have arrived at the barrier and so it is ready. We use the `bp_transition` predicate to go through the barrier’s slots in the `waitpool`, combine the associated heaps and barrier maps, redivide these resources according to the barrier postconditions, and remove the associated resources from the barrier pool into a list of slots called `out`. Finally, the states in `out` are combined with the suspended threads, which are simultaneously resumed by the `transition_threads` predicate. The formal definitions of the `bp_transition` and `transition_threads` predicates are extremely complex and very tedious and we refer interested readers to the mechanization.

Oracle semantics. Following Hobor *et al.* [16, 15], we define a third *oracular semantics*: $(\sigma, o, c) \mapsto (\sigma', o', c')$. Here the sequential state σ and command c are exactly the same as in the purely sequential step. The new parameter o is an oracle, a kind of box containing “the rest” of the concurrent machine—that is, o contains a scheduler, a list of other threads, and a barrier pool.

The oracle semantics behaves exactly the same way as the purely sequential semantics on all of the instructions except for the barrier call, with the oracle o being passed through unchanged. That is to say:

$$\frac{(\sigma, c) \mapsto (\sigma', c')}{(\sigma, o, c) \mapsto (\sigma', o, c')} \text{ os-seq}$$

¹⁰In Coq things are trickier since we track some technical side conditions via dependent types so this relation also ensures that these side conditions remain satisfied.

When the oracle semantics reaches a barrier instruction, it consults the oracle o to determine the state of the machine after the barrier:

$$\frac{\text{consult}(h, b, o) = (h', b', o')}{((s, h, b), o, \mathbf{barrier\ bn}; c) \mapsto ((s, h', b'), o', c)} \text{os-consult}$$

The formal definition of the `consult` relation is detailed in [16, 15] but the idea is simple. To consult the oracle, one unpacks the concurrent machine and runs (classically) all of the other threads until control returns to the original thread; `consult` then returns the current h' and b' (that resulted from the barrier call) and repackages the concurrent machine into the new oracle o' . The final case of the oracle semantics occurs when the concurrent machine never returns control (because it got stuck or due to sheer perversion of the scheduler):

$$\frac{\nexists r. \text{consult}(h, b, o) = r \quad (\text{i.e., consult diverges})}{((s, h, b), o, \mathbf{barrier\ bn}; c) \mapsto ((s, h, b), o, \mathbf{barrier\ bn}; c)} \text{os-diverge}$$

When control will never return, it does not matter what this thread does as long as it does not get stuck; accordingly we enter an (infinite) loop.

Soundness proof outline. Our soundness argument falls into several parts. We define our Hoare tuple in terms of our oracle semantics using a definition by Appel and Blazy [2]; this definition was designed for a sequential language and we believe that other standard sequential definitions for Hoare tuples would work as well¹¹. We then prove (in Coq) all of the Hoare rules for the sequential instructions; since the `os-seq` case of the oracle semantics provides a straight lift into the purely sequential semantics this is straightforward¹².

Next, we prove (in Coq) the soundness for the barrier rule. This turns out to be much more complicated than a proof of the soundness of (non-first-class) locks and took the bulk of the effort. There are two points of particular difficulty: first, the excruciatingly painful accounting associated with tracking resources during the barrier call as they move from a source thread (as a precondition), into the barrier pool, and redistribution to the target thread(s) as postcondition(s). The second difficulty is proving that a thread that enters a barrier while holding more than one precondition will never wake up; the analogy is a door with n keys distributed among n owners; if an owner has a second key in his pocket when he enters then one of the remaining owners will not be able to get in.

After proving the Hoare rules from Figure 3 sound with respect to the oracle semantics, the remaining task is to connect the oracle semantics to the concurrent semantics—that is, *oracle soundness*. Oracle soundness says that if each of the threads on a machine are safe with respect to the oracle semantics, then the entire concurrent machine combining the threads together is safe. The (very rough) analogy to this result in Brookes' semantics is the parallel decomposition lemma. Here we use a progress/preservation style proof closely following that given in [15, pp.242–255]; the proof was straightforward and quite short to mechanize. A technical advance over previous work is that the progress/preservation proofs do not require that the concurrent semantics be deterministic. In fact, allowing the semantics to be nondeterministic simplified the proofs significantly.

¹¹We change Appel and Blazy's definition so that our Hoare tuple guarantees that the allocation pool is available for verifying the Hoare rule for $x := \mathbf{new\ } e$.

¹²The Hoare rule for loops (While) is only proved on paper. The loop rule is known to be painful to mechanize and so the mechanization was skipped due to time constraints. It has been proved in Coq for similar (indeed, more complicated from a sequential control-flow perspective) settings in previous work [2, 16].

File	LOC	Time	Description
SLB_Base	1,182	2s	Utility lemmas (largely list facts)
SLB_Lang	1,240	11s	States, program syntax, assertion model
SLB_BarDefs	265	2s	Barrier definitions
SLB_CLang	3,230	1m7s	Dynamic concurrent state
SLB_SSem	415	17s	Sequential semantics
SLB_Sem	784	33s	Concurrent semantics
SLB_ESSem	230	5s	Erased semantics
SLB_ESEquiv	3,352	30s	Erasure proofs
SLB_OSem	1,942	2m10s	Oracular semantics
SLB_HRules	170	2s	Definition of Hoare tuples
SLB_OSound	426	30s	Soundness of oracle semantics
SLB_HRulesSound	1,664	1m14s	Soundness proofs for Hoare rules
SLB_Ex	2,700	48s	Example of a barrier definition
Total	16,598	7m34s	

Figure 5: Proof structure, size and compilation times (2.66GHz, 8GB)

A direct consequence of oracle soundness is that if each thread is verified with the Hoare rules, and is loaded onto a single concurrent machine, then if the machine does not get stuck and if it halts then all of the postconditions hold.

Erasure. One can justly observe that our concurrent semantics is not especially realistic; *e.g.*, we: explicitly track resource ownership permissions (*i.e.*, our semantics is *unerased*); have an unrealistic memory allocator/deallocator and scheduler; ignore issues of byte-addressable memory; do not store code in the heap; and so forth. We believe that we could connect our semantics to a more realistic semantics that could handle each of these issues, but most of them are orthogonal to barriers. For brevity we will comment only on erasing the resource accounting since it forms the heart of our soundness result.

We have defined, in Coq, an *erased* sequential and concurrent semantics. An erased memory is simply a pair of a break address and a total function from addresses to values. The run-time state of an erased barrier is simply a pair of naturals: the first tracking the number of threads currently waiting on the barrier, and the second giving the final number of threads the barrier is waiting for. We define a series of `erase` functions that take an unerased type (memory/barrier status/thread/etc.) to an erased one by “forgetting” all permission information. The sequential erased semantics is quite similar to the unerased one, with the exception that we do not check if we have read/write permission before executing a load/store. The concurrent erased semantics is much simpler than the complicated accounting-enabled semantics explained above since all that is needed to handle the barrier is incrementing/resetting a counter, plus some modest management of the thread list to suspend/resume threads. Critically, our erased semantics is a computable function, enabling program evaluation. Finally, we have proved that our unerased semantics is a conservative approximation to our erased one: that is, if our unerased concurrent machine can take a step from some state Σ to Σ' , then our erased machine takes a step from `erase`(Σ) to `erase`(Σ').

7. COQ DEVELOPMENT

We detail our Coq development in Figure 5. We use the Mechanized Semantic Library [1] for the definitions of share models, separation algebras, and various utility lemmas/tactics. In addition to the standard Coq axioms, we use dependent and propositional extensionality and the law of excluded middle.

Over 7,000 lines of the development is devoted to proving the soundness of the Hoare rule for barriers, largely in the files `SLB_BarDefs.v`, `SLB_CLang.v`, `SLB_Sem.v`, `SLB_OSem.v`, `SLB_HRules.v`, and a small part of `SLB_HRulesSound.v`. The rest of the concurrent semantics, the oracle semantics, and the soundness of the oracle semantics (\sim the parallel decomposition lemma) require approximately 1,000 lines, largely in the files `SLB_Sem.v`, `SLB_HRules.v`, and `SLB_OSound`. The erased semantics requires 230 lines in `SLB_ESSem.v`, while the associated equivalence proofs require 3352 lines in the file `SLB_ESEquiv.v`.

The sequential semantics and proofs for the associated Hoare rules require approximately 2,000 lines drawn from the files `SLB_Lang.v`, `SLB_SSem.v`, `SLB_HRules.v`, and `SLB_HRulesSound.v`. We estimate that the proof of the loop rule would require a further 2,000-3,000 lines. The model of our assertions and the program syntax are both in `SLB_Lang.v`. Utility lemmas/tactics (`SLB_Base.v`) and the example barrier (`SLB_Ex.v`) complete the development.

8. TOOL SUPPORT

We have integrated our program logic for barriers into the HIP/SLEEK program verification toolset [20, 12]. SLEEK is an entailment checker for separation logic and HIP applies Hoare rules to programs and uses SLEEK to discharge the associated proof obligations. We proceeded as follows:

- (1) We developed an equational solver over the sophisticated fractional share model of Dockins *et al.* [11]. Permissions can be existentially or universally quantified and arbitrarily related to permission constants.
- (2) We integrated our equational solver over shares into SLEEK to handle fractional permissions on separation logic assertions (*e.g.*, maps-to, etc.). We believe that SLEEK is the first automatic entailment checker for separation logic that can handle a sophisticated share model.
- (3) We developed an encoding of barrier definitions (diagrams) in SLEEK, which now automatically verifies the side conditions from §4.
- (4) We modified HIP to recognize barrier definitions (whose side conditions are then verified in SLEEK) and barrier calls using the Hoare rule from Figure 3.

Next we describe our equational solver for the Dockins *et al.* share model before giving a more technical background to the HIP/SLEEK system and describing our modifications to it in detail. Most of the technical work occurred in developing the equational solver and its integration with the rest of the separation logic entailment procedures in SLEEK. Once SLEEK understood fractional permissions, checking the validity side conditions on barriers was quite simple.

8.1. Decision Procedure for Shares. SLEEK discharges the heap related proof obligations but relies on external decision procedures for the pure logical fragments it extracts from separation logic formulae. For example, SLEEK utilizes Omega for Presburger arithmetic, Redlog for arithmetic in \mathbb{R} , MONA for monodic second-order logic, etc. Adding fractional permissions required the development of an appropriate equational decision procedure for fractional shares.

Decision procedures for simple kinds of fraction share models (*e.g.*, rationals between 0 and 1) are nothing more than solvers for systems of linear equations. However, our logic uses the more sophisticated fractional share model of Dockins *et al* [11] and so we had to develop an equational solver from scratch.

Dockins *et. al* represents shares as binary trees with boolean-valued leaves. The full share \blacksquare is a tree with one true leaf \bullet and the empty share \square is a tree with one false leaf \circ . The left-half share \blacktriangleleft is a tree with two leaves, one true and one false: $\widehat{\bullet \circ}$; similarly, the right-half share \blacktriangleright is a tree with two leaves, one false and one true: $\widehat{\circ \bullet}$. The trees can continue to be split indefinitely: for example, the right half of \blacktriangleleft is $\widehat{\circ \bullet}$. Joining is defined by structural induction on the shape of the trees with the base cases $\circ \oplus \circ = \circ$, $\bullet \oplus \circ = \bullet$, and $\circ \oplus \bullet = \bullet$ (emphasis: \oplus is partial). When two trees do not have the same shape, they are unfolded according to the rules $\bullet \cong \widehat{\bullet \bullet}$ and $\circ \cong \widehat{\circ \circ}$; for example:

$$\widehat{\circ \bullet} \oplus \widehat{\bullet \circ} = \widehat{\circ \bullet \circ} \oplus \widehat{\bullet \circ \circ} = \widehat{\bullet \bullet \circ} = \widehat{\bullet \circ \bullet}$$

SLEEK takes a formula in separation logic with fractional shares and extracts a specialized formula over **strictly positive** shares whose syntax is as follows:

$$\phi ::= \exists v. \phi \mid \phi_1 \vee \phi_2 \mid \phi_1 \wedge \phi_2 \mid v_1 \oplus v_2 = v_3 \mid v_1 = v_2 \mid v = \chi$$

Our share formulae ϕ contain share variables v , existentials \exists , conjunctions \wedge , disjunctions \vee , join facts \oplus , equalities between variables, and assignments of variables to constants χ . Our tool also recognizes one additional form $v \in (\chi_1, \chi_2)$, pronounced “ v is bounded by χ_1 and χ_2 ”, which is semantically equal to:

$$((v = \chi_1) \vee (\exists v'. \chi_1 \oplus v' = v)) \wedge ((v = \chi_2) \vee (\exists v''. v \oplus v'' = \chi_2))$$

We require disjunctions because our share variables can only be instantiated with positive shares: $\forall v. \nexists v'. v \oplus v' = v$. It is more efficient to handle bounds checks “natively” rather than compile them into their semantic definition because we wish to avoid adding additional existentials and disjunctions.

SLEEK then asks our solver questions of the following forms:

- (1) (UNSAT) Is a given formula ϕ unsatisfiable?
- (2) (\exists -ELIM) Given a formula of the form $\exists v. \phi(v)$, is there a unique constant χ to which we can instantiate v without affecting the truth of $\phi(v)$?
- (3) (IMPL) Given two formulae ϕ_1 and ϕ_2 , does ϕ_1 entail ϕ_2 ?

Our solver is sound but incomplete. However, it is *complete enough* to help SLEEK check a wide variety of entailments involving fractional permissions, including all of those in the example from Figure 1.

As we will show, all of these questions can be reduced to solving a series of constraint systems whose equations are of the form $v_1 \oplus v_2 = v_3$, $v \in (\chi_1, \chi_2)$, and $v = \chi$. Solving constraint systems in separation algebras (*i.e.*, cancellative partial commutative monoids) is not as straightforward as it might seem because many of the traditional algebraic techniques

do not apply. Our lightweight constraint solver finds an overapproximation to the solution, returning either (a) the constant **UNSAT** or (b) for each variable v_i either an **assignment** $v_i = \chi$ or a **bound** $v_i \in (\chi_1, \chi_2)$ such that:

- (FALSE) If the algorithm returns **UNSAT**, then the formula is unsatisfiable. The algorithm will return **UNSAT** if it discovers a bound whose “lower value” is higher than its “upper value”, or if it discovers a falsehood (*e.g.*, after constant propagation one of the equations becomes $\blacksquare \oplus \blacksquare = \blacksquare$).
- (COMPLETE) All solutions to the system (if any) lie within the bounds.
- (SAT-PRECISE) We say that a solution is *precise* when all variables are given assignments. If a solution is precise, then the formula is satisfiable.

SLEEK queries are given in share formulae that must be transformed into the equational systems understood by our constraint solver. To do this transformation, first we put the relevant formulae into disjunctive normal form (DNF). Each disjunct becomes an independent system of equations. Given one disjunct we form this system by simply treating each basic constraint (*i.e.*, $v = v'$, $v = \chi$, $v \in (\chi_1, \chi_2)$, and $v_1 \oplus v_2 = v_3$) as an equation. Our solver approximates each system independently and can then answer SLEEK’s questions as follows:

- (UNSAT): We return **False** when the algorithm returns **UNSAT** for each constraint system obtained from the formula; otherwise return **True**.
- (\exists -ELIM): If the variable v has the same assignment in all constraint systems derived from the DNF, then return that value. We can soundly substitute that value for v and eliminate the existential. (If the formula is satisfiable, then that is the unique assignment that makes it so; if the formula is false then after the assignment it will still be false.)
- (IMPL): We return **True** only when either:
 - the solver returns **UNSAT** for all systems derived from the antecedent
 - the solver returns a precise solution for each system of equations derived from the antecedent, and the solver also returns **the same** precise solution for at least one of the consequent systems.

The constraint solver works by eliminating one class of constraints at a time:

- (1) First we substitute $v = cs$ constraints into the remaining equations.
- (2) We handle \oplus constraints with exactly one variable as follows:
 - $\chi_1 \oplus \chi_2 = v$: we check if the join is defined, and if so substitute the sum for v in the remaining equations; otherwise, we return **UNSAT**.
 - $\chi_l \oplus v = \chi_r$ or $v \oplus \chi_l = \chi_r$: we check if χ_r contains χ_l , and if so substitute the difference $\chi_r - \chi_l$ for v in the remaining equations; otherwise return **UNSAT**. (“ $-$ ” has the property that if $\chi_1 - \chi_2 = \chi_3$ then $\chi_3 \oplus \chi_2 = \chi_1$).
- (3) Constraints involving constants ($\chi_1 \oplus \chi_2 = \chi_3$ and $\chi \in (\chi_1, \chi_2)$) are dismissed if the equality/inequalities hold; otherwise return **UNSAT**.
- (4) We attempt to dismiss certain kinds of unsatisfiable systems via a consistency check as follows. We first compute the transitive closure of variable substitutions, resulting in facts of the form $v_1 \oplus \dots \oplus v_n \oplus \chi_1 \oplus \dots \oplus \chi_m = \chi$. Nonempty shares **cannot** join with themselves. Therefore, if the v_i contain duplicates we return **UNSAT**. We also return **UNSAT** if the constants χ_i do not join or if χ does not contain $\chi_1 \oplus \dots \oplus \chi_m$.
- (5) Variables in the remaining constraints are given initial domains of (\square, \blacksquare).
- (6) Each \in constraint is used to restrict the domain of its corresponding variable.

- (7) At this point only $a_1 \oplus a_2 = a_3$ constraints involving at least two variables remain. The algorithm then proceeds by iteratively selecting an equation, checking it for consistency, and then refining the associated domains via a forward and backward propagation. The algorithm iterates until either a fixpoint is reached or a consistency check fails. To check an equation for consistency, the algorithm verifies that:
- for each variable, the lower bound is less than the upper bound
 - the current lower bounds of the LHS variables join together
 - the join of the LHS lower bounds is below by the RHS upper bound
 - the join of the LHS upper bounds is above by the RHS lower bound

Forward propagation consists of (Fa) lowering the upper bound of the RHS by intersecting away any subtree that does not appear in the upper bounds of the LHS, and (Fb) increasing the lower bound of the RHS by unioning all subtrees present in the lower bounds in the LHS. Backwards propagation consists of (Ba) lowering the upper bounds of the LHS by intersecting away any subtree that does not appear in the upper bound on the RHS. Increasing the lower bounds of the LHS (Bb) is trickier since we do not know *which operand* should be increased. There are several possibilities we could have taken, but we selected the simplest: we simply leave the bounds as they were unless one of the operands has been determined to be a constant, in which case we can calculate exactly what the lower bound for the other variable should be. This solution is can lead to over approximation, but a more refined solution would require a performance cost, which did not seem warranted by our experiments. After each forward/backwards propagation, if we have refined a domain to a single point, the variable is substituted for a constant value of that point in the remaining equations.

Once we reach a fixpoint, the resulting variable bounds represent an over approximation of the solution.

8.2. An introduction to SLEEK. SLEEK is an entailment checker for separation logic [21]. The antecedent is allowed to cover more heap resources than the consequent, and after the entailment succeeds SLEEK returns this residual heap state together with the pure state of the antecedent. SLEEK can also discover the correct instantiations for certain kinds of existential variables in the consequent.

One of SLEEK’s strong points is that it allows user defined inductive predicates. Predicates are defined as separation formulae that describe the shape of data structures and associated properties (*e.g.*, list length, tree height, and bag of values contained in a list). SLEEK uses the keyword `self` as a pointer variable to the current object. Predicate invariants can increase the precision of the verification (*e.g.*, `length > 0`). An invariant for a predicate instance has two components: a pure formula describing constraints on the arguments and a set of pointer arguments (*e.g.* the outward pointer for a list segment).

Figure 6 gives an outline of the specification language accepted by SLEEK with our extensions for the fractional permissions. The system accepts disjunctive separation logic formulae with both heap and pure constraints. The formulae syntax allows richer structures as well, *e.g.* directed case analysis, staged formulae, and uniformly handling pre/post conditions [12]; we use the latter feature to enable a program verifier like HIP to verify barrier calls.

The core of a separation logic formula is the heap formula. Heap constraints are sets of heap node descriptions connected by the separating conjunction. A node is either a data

Predicate	spred	$::= [self::]p\langle v^* \rangle \equiv \mathbf{Q} [\text{inv } \pi_0]$
Formula	\mathbf{Q}	$::= \mathbf{R} \mid \mathbf{R} \vee \mathbf{Q}$
	\mathbf{R}	$::= \text{case}\{\pi_1 \rightarrow \mathbf{Q}_1, \dots, \pi_n \rightarrow \mathbf{Q}_n\} \mid [w^*] \Phi [\mathbf{Q}]$
	Φ	$::= \bigvee \exists v^* \cdot (\kappa \wedge \pi \wedge \tau)$
Frac form.	τ	$::= v_{f1} \oplus v_{f2} = v_{f3} \mid v \in (\chi_1, \chi_2) \mid v = \chi \mid \tau \wedge \tau$
Pure form.	π	$::= \gamma \wedge \phi$
Pointer form.	γ	$::= v_1 = v_2 \mid v = \text{null} \mid v_1 \neq v_2 \mid v \neq \text{null} \mid \gamma_1 \wedge \gamma_2$
Heap form.	κ	$::= \text{emp} \mid v :: c^{vf} \langle v^* \rangle \mid v :: p^{vf} \langle v^* \rangle \mid \kappa_1 * \kappa_2$
Presburger arith.	ϕ	$::= \text{arith} \mid \phi_1 \wedge \phi_2 \mid \phi_1 \vee \phi_2 \mid \neg \phi \mid \exists v \cdot \phi \mid \forall v \cdot \phi$
	arith	$::= \mathbf{a}_1 = \mathbf{a}_2 \mid \mathbf{a}_1 \neq \mathbf{a}_2 \mid \mathbf{a}_1 < \mathbf{a}_2 \mid \mathbf{a}_1 \leq \mathbf{a}_2$
	\mathbf{a}	$::= k \mid v \mid k \times \mathbf{a} \mid \mathbf{a}_1 + \mathbf{a}_2 \mid -\mathbf{a} \mid \max(\mathbf{a}_1, \mathbf{a}_2) \mid \min(\mathbf{a}_1, \mathbf{a}_2)$
where	p	is a predicate name;
	v, w	are variable names;
	c	is a data type name;
	k	is an integer constant;
	τ	represents the fractional permission constraints
	χ	represents constant fractional shares

Figure 6: The Specification Language with Fractional Permissions.

structure instance or an instance of a user defined predicate. Here we use the same notation for both cases: $v :: c \langle v^* \rangle$, where v is the pointer to the structure, c indicates the data structure type or the predicate name, and v^* is the a list of arguments (either predicate arguments for a predicate instance or field values for a data structure). Separation logic formulae can also contain pure constraints over several domains: arithmetic, bag and list, etc. For brevity we only include arithmetic constraints in this presentation.

The syntax in Figure 6 contains two new extensions to SLEEK’s language. First, heap node descriptions can contain permission annotations for fractional ownership. A partially-owned heap node v_f is indicated by $v :: c^{vf} \langle v^* \rangle$. If p denotes a predicate, then the notation $v :: p^{vf} \langle v^* \rangle$ for a predicate instance indicates that all abstracted heap nodes are owned with permission v_f (e.g., a \blacksquare -owned list indicates each list cell is owned \blacksquare). A node/predicate without a permission annotation indicates full ownership. The second extension enables the expression of constraints over fractional permission variables using the syntax $v_{f1} \oplus v_{f2} = v_{f3}$, $v \in (\chi_1, \chi_2)$, $v_1 = v_2$, and $v = \chi$.

8.3. SLEEK entailment background. The core of the SLEEK entailment works by algorithmically discharging the heap obligations and then referring any remaining pure constraints to other provers. SLEEK discharges heap obligations in three ways: heap node matching, predicate folding, and predicate unfolding. To guarantee termination, SLEEK ensures that predicate folds/unfolds must be immediately followed by a match, which guarantees that heap nodes are never generated in the RHS of the entailment.

Entailments in SLEEK are written as follows: $\Delta_A \vdash_V^{\kappa} \Delta_C * \Delta_R$, which is shorthand for $\kappa * \Delta_A \vdash \exists V \cdot (\kappa * \Delta_C) * \Delta_R$. The entailment establishes if all the consequent heap nodes Δ_C are covered by heap nodes in antecedent Δ_A if so, SLEEK returns a residual heap Δ_R , consisting of the antecedent nodes that are not used in covering Δ_C . κ is the history of nodes from the antecedent that have been used to match nodes from the consequent, V is the list of existentially quantified variables from the consequent. Note that k and V are discovered iteratively: the entailment checking procedure begins with $\kappa = \text{emp}$ and $V = \phi$.

$$\begin{array}{c}
XPure(\text{emp}) \equiv (\text{true}, \phi) \\
\text{IsData}(c) \\
\hline
XPure(p :: c < v^* >) \equiv (p \neq 0; \{p\}) \\
XPure(\kappa_1) \equiv (f_1, s_1) \quad XPure(\kappa_2) \equiv (f_2, s_2) \\
\hline
XPure(\kappa_1 * \kappa_2) \equiv (f_1 \wedge f_2, s_1 \cup s_2) \\
\text{IsPred}(c) \quad (c < v^* > \equiv \Phi \text{ inv } (\pi_1, \pi_2)) \in \text{P} \\
\hline
XPure(p :: c < v_p^* >) \equiv ([p/self, v_p^*/v^*]\pi_1, [p/self, v_p^*/v^*]\pi_2) \\
XPure(\kappa_1) \equiv (f_1, s_1) \quad XPure(\kappa_2) \equiv (f_2, s_2) \\
\hline
XPure(\kappa_1 * \kappa_2) \equiv (f_1 \wedge f_2, s_1 \cup s_2)
\end{array}$$

Figure 7: XPure : Translating to Pure Form

The initial system behavior was described in detail in [21, 12]. Here we outline the main rules for matching, folding, unfolding, and discharging of pure constraints. The main entailment checking rules are given in Fig 8.

Entailment between separation formulae is reduced to entailment between pure formulae by matching heap nodes in the RHS to heap nodes in the LHS (possibly after a fold/unfold). Once the RHS is pure, the remaining LHS heap formula is soundly approximated to a pair of pure formula and set of disjoint pointers by function XPure as defined in Fig 7. The functions IsData(c) and IsPred(c) decide respectively if c is a data structure or a predicate. The procedure successively pairs up heap nodes that it proves are aliased. SLEEK keeps the successfully matched nodes from the antecedent in κ for better precision in the next iteration.

$$\begin{array}{c}
\text{EMP} \qquad \qquad \qquad \text{MATCH} \\
\frac{(\rho, S) = XPure(\kappa_1 * \kappa) \quad \rho \wedge (\forall x, y \in S \cdot x \neq y) \implies \rho \exists V \cdot \pi_2}{\kappa_1 \wedge \pi_1 \vdash_V^{\kappa} \pi_2 * (\kappa_1 \wedge \pi_1)} \quad \frac{fst(XPure(p_1 :: c < v_1^* > * \kappa_1)) \wedge \pi_1 \implies p_1 = p_2 \quad \rho = [v_1^*/v_2^*] \quad \kappa_1 \wedge \pi_1 \wedge \text{freeEqn}(\rho, V) \vdash_{V - \{v_2^*\}}^{\kappa * p_1 :: c < v_1^* >} \rho(\kappa_2 \wedge \pi_2) * \Delta}{p_1 :: c < v_1^* > * \kappa_1 \wedge \pi_1 \vdash_V^{\kappa} (p_2 :: c < v_2^* > * \kappa_2 \wedge \pi_2) * \Delta} \\
\\
\text{FOLD} \\
\frac{\text{IsPred}(c_2) \wedge \text{IsData}(c_1) \quad c_2 < v^* > \equiv \Phi \in \text{P} \quad W_i = V_i - \{v^*, p\} \quad fst(XPure(p_1 :: c_1 < v_1^* > * \kappa_1)) \wedge \pi_1 \implies p_1 = p_2 \quad p_1 :: c_1 < v_1^* > * \kappa_1 \wedge \pi_1 \vdash_{\{p, v^*\}}^{\kappa} [p/self] \Phi * \{(\Delta_i, \kappa_i, V_i, \pi_i)\}_{i=1}^n \quad (\Delta^r, \kappa^r, \pi^r) \in \{(\Delta_i, \kappa_i, V_i, \pi_i)\}_{i=1}^n \quad (\pi^a, \pi^c) = \text{split}_V^{\{v_2^*\}}(\pi^r) \quad \Delta^r \wedge \pi^a \vdash_V^{\kappa^r} (\kappa_2 \wedge \pi_2 \wedge \pi^c) * \Delta}{p_1 :: c_1 < v_1^* > * \kappa_1 \wedge \pi_1 \vdash_V^{\kappa} (p_2 :: c_2 < v_2^* > * \kappa_2 \wedge \pi_2) * \Delta} \\
\\
\text{UNFOLD} \\
\frac{fst(XPure(p_1 :: c_1 < v_1^* > * \kappa_1)) \wedge \pi_1 \implies p_1 = p_2 \quad c < v^* > \equiv \Phi \in \text{P} \quad \text{IsPred}(c_1) \wedge \text{IsData}(c_2) \quad [p/self, v_1^*/v^*] \Phi * \kappa_1 \wedge \pi_1 \vdash_V^{\kappa} (p_2 :: c_2 < v_2^* > * \kappa_2 \wedge \pi_2) * \Delta}{p_1 :: c_1 < v_1^* > * \kappa_1 \wedge \pi_1 \vdash_V^{\kappa} (p_2 :: c_2 < v_2^* > * \kappa_2 \wedge \pi_2) * \Delta}
\end{array}$$

Figure 8: Separation Constraint Entailment

All three heap reducing steps start by establishing that there is a heap node on the LHS of the entailment that is aliased with the RHS heap node that is to be reduced ($p_1 = p_2$). In order to

$$\begin{array}{c}
FXPure(\mathbf{emp}, \tau) \equiv (\mathbf{true}, \phi) \\
\frac{\text{IsData}(c) \quad \tau \Rightarrow v_f = cs}{FXPure(p :: c^{v_f} < v^* >, \tau) \equiv (p \neq 0; \{(p, cs)\})} \\
\frac{FXPure(\kappa_1, \tau) \equiv (f_1, s_1) \quad FXPure(\kappa_2, \tau) \equiv (f_2, s_2)}{FXPure(\kappa_1 * \kappa_2, \tau) \equiv (f_1 \wedge f_2, s_1 \cup s_2)} \\
\frac{\text{IsPred}(c) \quad (c^{v^*}) \equiv \Phi \text{ inv } (\pi_1, \pi_2) \in \mathbf{P}}{\tau \Rightarrow v_f = cs \quad \pi'_1 = [p/self]\pi_1 \quad \pi'_2 = \{\forall v \in \pi_2, ([p/self]v, cs)\}} \\
\hline
XPure(p :: c^{v_f} < v^* >, \tau) \equiv (\pi'_1, \pi'_2)
\end{array}$$

Figure 9: FXPure: XPure with shares

prove the aliasing, the LHS heap together with the previously consumed nodes are approximated to a pure formula and together with the LHS pure formula the $p_1 = p_2$ implication is checked. When a match occurs (rule MATCH), all the free arguments of the matched node from the RHS are instantiated and those instantiations are moved to the LHS. The free RHS arguments are bound to the corresponding argument in the LHS and the equalities are moved to the LHS.

These instantiations are particularly useful for the entailments constructed by program verifiers like HIP because free variables in the RHS are usually variables from method preconditions and these bindings are useful in subsequent program verification steps. By moving the bindings to the LHS, they can be carried forward in the residual state and thus returned to the program verifier. This process is formalized by the function `freeEqn` below, where V is the set of existentially quantified variables:

$$\text{freeEqn}([u_i/v_i]_{i=1}^n, V) \equiv \bigwedge_{i=1}^n (\text{if } v_i \in V \text{ then true else } v_i = u_i)$$

Unfold/fold operations handle inductive predicates in a deductive manner. SLEEK can unfold a predicate instance that appears in the LHS if the unfolding exposes a heap node that matches immediately with a node in the RHS. Similarly, several LHS nodes can be folded into a predicate instance if the resulting predicate instance can be immediately matched with a RHS node. Well-formedness conditions imposed on the predicate definitions ensure that after a fold or unfold a matching always takes place; these conditions have been elided for this presentation. The unfold rule presents the replacement of a predicate instance with the predicate definition in which the arguments have been substituted. The fold step requires the LHS to entail the predicate definition. The residue of this entailment is then used as the new LHS for the rest of the original entailment. The folding rule needs to carefully deal with the instantiations: some move to the LHS while others still need to be proven and thus are moved to the RHS. This partitioning of the pure constraints is done by the `split` function. For a more detailed explanation of the SLEEK entailment process please consult Nguyen *et. al* [21].

8.4. Entailment Procedure for Separation Logic with shares. Adding fractional permissions required several modifications to the entailment process.

- **Empty heap.** In a separation logic without shares, whenever $(\exists a, b. x \mapsto a * y \mapsto b)$ then $x \neq y$. In SLEEK, this fact is captured in the EMP rule which tries to prove the pure part of the consequent after enriching the antecedent pure formula with pure information collected from the previously consumed heap and the remaining LHS heap. It extracts both the invariants of the heap nodes and also constructs a formula that enforces that all pointers in the heap are distinct.

Introducing fractional permissions requires the relaxation of this constraint because $\exists a, b. x \xrightarrow{x_f} a * y \xrightarrow{y_f} b$ implies $x \neq y$ only if the x_f and y_f shares overlap. We changed the XPure

<p style="margin: 0;">FOLD</p> $\frac{\begin{array}{l} \text{IsPred}(c_2) \wedge \text{IsData}(c_1) \quad c_2 < v^* > \equiv \Phi \in P \quad W_i = V_i - \{v^*, p\} \\ \text{fst}(\text{FXPure}(p_1 :: c_1^{f_1} < v_1^* > * \kappa_1, \tau_1)) \wedge \pi_1 \implies p_1 = p_2 \\ \Phi' = \text{set_shares}([p/\text{self}]\Phi, f_2) \\ p_1 :: c_1^{f_1} < v_1^* > * \kappa_1 \wedge \pi_1 \wedge \tau_1 \vdash_{\{p, v^*\}}^{\kappa} \Phi' * \{(\Delta_i, \kappa_i, V_i, \pi_i, \tau_i)\}_{i=1}^n \\ (\Delta^r, \kappa^r, \pi^r, \tau^r) \in \{(\Delta_i, \kappa_i, V_i, \pi_i, \tau_i)\}_{i=1}^n \\ (\pi^a, \pi^c) = \text{split}_{V-\{v_2^*\}}^{\{v_2^*\}}(\pi^r) \quad (\tau^a, \tau^c) = \text{split}_{V-\{v_2^*\}}^{\{v_2^*\}}(\tau^r) \\ \Delta^r \wedge \pi^a \wedge \tau^a \vdash_V^{\kappa^r} (\kappa_2 \wedge \pi_2 \wedge \tau_2 \wedge \pi^c \wedge \tau^c) * \Delta \end{array}}{p_1 :: c_1^{f_1} < v_1^* > * \kappa_1 \wedge \pi_1 \wedge \tau_1 \vdash_V^{\kappa} (p_2 :: c_2^{f_2} < v_2^* > * \kappa_2 \wedge \pi_2 \wedge \tau_2) * \Delta}$
<p style="margin: 0;">UNFOLD</p> $\frac{\begin{array}{l} \text{fst}(\text{FXPure}(p_1 :: c_1^{f_1} < v_1^* > * \kappa_1, \tau_1)) \wedge \pi_1 \implies p_1 = p_2 \\ c < v^* > \equiv \Phi \in P \quad \text{IsPred}(c_1) \wedge \text{IsData}(c_2) \\ \Phi' = \text{set_shares}([p/\text{self}, v_1^*/v^*]\Phi, f_1) \\ \Phi' * \kappa_1 \wedge \pi_1 \wedge \tau_1 \vdash_V^{\kappa} (p_2 :: c_2^{f_2} < v_2^* > * \kappa_2 \wedge \pi_2 \wedge \tau_2) * \Delta \end{array}}{p_1 :: c_1^{f_1} < v_1^* > * \kappa_1 \wedge \pi_1 \wedge \tau_1 \vdash_V^{\kappa} (p_2 :: c_2^{f_2} < v_2^* > * \kappa_2 \wedge \pi_2 \wedge \tau_2) * \Delta}$

Figure 10: Folding/Unfolding in the presence of shares

function to return a pair of a pure formula, and pairs of pointers and associated fractional shares. Our new version of XPure allowed the EMP rule to be rewritten to enforce inequality only between pointers that have conflicting shares:

$$\frac{(\rho, S) = \text{FXPure}(\kappa_1 * \kappa, \tau) \quad \rho \wedge (\forall (x, x_f), (y, y_f) \in S, (\neg \exists z \cdot x_f \oplus y_f = z) \cdot x \neq y) \implies \rho \exists V \cdot \pi_2}{\kappa_1 \wedge \pi_1 \wedge \tau_1 \vdash_V^{\kappa} \pi_2 * (\kappa_1 \wedge \pi_1 \wedge \tau_1)} \text{EMP}$$

- **Folding/unfolding.** Given the convention that all the heap nodes abstracted by a predicate instance are owned with the same fractional permission as the predicate instance, the unfold operation replaces the permissions of all the heap nodes in the predicate definition with the permission of the node that is unfolded (the LHS node) before the actual replacement of the predicate instance with the predicate definition. Similarly for folding, before entailing the predicate definition, the permissions of all nodes in the definition are replaced with the permission of the RHS node that needs to be discharged. This is achieved by the use of the `set_shares`(Φ, v) function which sets the permissions of all heap nodes in Φ to v . The new set of rules can be formalized in Figure 10.
- **Matching.** In order to properly handle a match in the presence of fractional shares the entailment process needs to explore three matching options: (a) it can reduce entirely both LHS and RHS nodes, or (b) split in two the LHS node and reduce one, or (c) split the RHS and reduce one.

$$\frac{\begin{array}{l} \rho = [v_1^*/v_2^*, f_1/f_2] \\ (\text{fst}(\text{FXPure}(p_1 :: c_1^{f_1} < v_1^* > * \kappa_1, \tau_1)) \wedge \pi_1) \implies p_1 = p_2 \\ \kappa_1 \wedge \pi_1 \wedge \tau_1 \wedge \text{freeEqn}(\rho, V) \vdash_{V-\{v_2^*, f_2\}}^{\kappa * p_1 :: c_1^{f_1} < v_1^* >} \rho(\kappa_2 \wedge \pi_2 \wedge \tau_2) * \Delta \end{array}}{p_1 :: c_1^{f_1} < v_1^* > * \kappa_1 \wedge \pi_1 \wedge \tau_1 \vdash_V^{\kappa} (p_2 :: c_2^{f_2} < v_2^* > * \kappa_2 \wedge \pi_2 \wedge \tau_2) * \Delta} \text{FULL-MATCH (a)}$$

$$\begin{array}{c}
\rho = [v_1^*/v_2^*, f_{c1}/f_2] \\
\text{fst}(\text{FXPure}(p_1 :: c^{f_1} < v_1^* > * \kappa_1, \tau_1)) \wedge \pi_1 \implies p_1 = p_2 \\
\tau_1' = \tau_1 \wedge f_{c1} \oplus f_{r1} = f_1 \\
p_1 :: c^{f_{r1}} < v_1^* > * \kappa_1 \wedge \pi_1 \wedge \tau_1' \wedge \text{freeEqn}(\rho, V) \\
\frac{\vdash_{V - \{v_2^*, f_2\}}^{\kappa * p_1 :: c^{f_{c1}} < v_1^* >} \rho(\kappa_2 \wedge \pi_2 \wedge \tau_2) * \Delta}{p_1 :: c^{f_1} < v_1^* > * \kappa_1 \wedge \pi_1 \wedge \tau_1 \vdash_V^{\kappa} (p_2 :: c^{f_2} < v_2^* > * \kappa_2 \wedge \pi_2 \wedge \tau_2) * \Delta} \text{LEFT-SPLIT-MATCH (b)}
\end{array}$$

$$\begin{array}{c}
\rho = [v_1^*/v_2^*, f_1/f_{c2}] \\
\text{fst}(\text{FXPure}(p_1 :: c^{f_1} < v_1^* > * \kappa_1, \tau_1)) \wedge \pi_1 \implies p_1 = p_2 \\
V' = \text{if } f_2 \in V \text{ then } V \cup \{f_{c2}, f_{r2}\} \text{ else } V \\
\tau_1' = \text{if } f_2 \in V \text{ then } \tau_1 \text{ else } (\tau_1 \wedge f_{c2} \oplus f_{r2} = f_2) \\
\tau_2' = \text{if } f_2 \in V \text{ then } (\tau_2 \wedge f_{c2} \oplus f_{r2} = f_2) \text{ else } \tau_2 \\
\kappa_1 \wedge \pi_1 \wedge \tau_1' \wedge \text{freeEqn}(\rho, V) \\
\frac{\vdash_{V' - \{v_2^*, f_2, f_{c2}\}}^{\kappa * p_1 :: c^{f_1} < v_1^* >} \rho(p_2 :: c^{f_{r2}} < v_2^* > * \kappa_2 \wedge \pi_2 \wedge \tau_2') * \Delta}{p_1 :: c^{f_1} < v_1^* > * \kappa_1 \wedge \pi_1 \wedge \tau_1 \vdash_V^{\kappa} (p_2 :: c^{f_2} < v_2^* > * \kappa_2 \wedge \pi_2 \wedge \tau_2) * \Delta} \text{RIGHT-SPLIT-MATCH (c)}
\end{array}$$

Because this search process can be computationally expensive, we have devised an aggressive pruning technique. We try to determine if the fractional constraints restrict enough the fractional variables such that either : $f_1 = f_2$ in which case only FULL-MATCH is feasible or if f_1 is included in f_2 leading to RIGHT-SPLIT-MATCH or if f_1 includes f_2 allowing only LEFT-SPLIT-MATCH.

8.5. Proving barrier soundness. The fractional share solver and enhancements to SLEEK’s entailment procedures discussed above help with any program logic that needs fractional shares (e.g., concurrent separation logic with locks or sequential separation logic with read-only data). In contrast, our remaining work on HIP/SLEEK is specific to our logic for Pthreads-style barriers. Our initial goal is to automatically check the consistency of barrier definitions—that is, if a barrier definition meets the side conditions presented in §4. To do this first we must be able to describe a barrier diagram to SLEEK.

Although the barrier diagrams presented in §4 are intuitive and concise, we need a more textual representation for programs. Barrier diagrams describe all the possible transitions a barrier state can make and the specifications associated with said transitions. In a sense, a barrier definition can be viewed as a disjunctive predicate definition where the body is a disjunction of possible transitions.

SLEEK already contains user defined predicates so it is easy to introduce the “is a barrier” predicate $\text{barrier}(bn, v_f, s)$ as required by the barrier logic, with a slight change to notation to accommodate for the syntax presented in Figure 6 to $bn^{v_f} < s >$.

We extended SLEEK’s language to accept barrier diagrams in the form:

```

bdef           ::= barrier (b_name , thread_cnt , v* , transition*)
transition    ::= (from_state , to_state , pre-post-spec*)
pre-post-spec ::= ( $\Phi_{pre}$  ,  $\Phi_{post}$ )

```

We can now automatically check the well-formed conditions on the barrier definitions:

- all transitions must have exactly thread_cnt specifications, one for each thread
- For each transition, let from and to be the state labels, then:
 - for each specification (Φ_{pre} , Φ_{post})
 - * Φ_{pre} contains a fraction of the barrier in state from :
 $\Phi_{pre} \vdash \text{self} :: bn^{v_f} < \text{from} >$
 - * Φ_{post} contains a fraction of the barrier in state to :
 $\Phi_{pre} \vdash \text{self} :: bn^{v_f} < \text{to} >$

* $\Phi_{pre} * \Phi_{pre} \vdash \text{False}$

Our soundness proof assumes that each precondition P is precise. Unfortunately, precision is not very easy to verify automatically. As indicated in footnote 7, we believe that the logic will be sound if we can assume the (strictly) weaker property “token”: $P * P \vdash \text{False}$ instead of precision. At this stage, our prototype extension to SLEEK verifies that preconditions are tokens rather than that they are precise. We are in the process of attempting to update our soundness proof to require that preconditions be tokens rather than precise; if we are unable to do so then one solution would be for SLEEK to output a Coq file stating lemmas regarding the precision of each precondition. Users would then be required to prove these lemmas manually to be sure that their barrier definitions were sound. In our example barrier, the Coq proofs of precision were only a small part of the 2,700 total lines of Coq script, so the savings from using SLEEK to verify the soundness of a barrier definition should still be quite substantial.

- the star of all the precondition contains the full barrier (recall that the entailment \vdash check in SLEEK can produce a residue)

$$*_{i=1}^{i=\text{thread_cnt}} \Phi_{pre}^i \vdash \text{self} :: \text{bn}^{\blacksquare} \langle \text{from} \rangle$$
- the star of all the postconditions contains the full barrier

$$*_{i=1}^{i=\text{thread_cnt}} \Phi_{post}^i \vdash \text{self} :: \text{bn}^{\blacksquare} \langle \text{to} \rangle$$
- the star of all the preconditions equals the star of all postconditions modulo the barrier state change for a transition. We check this constraint by carving the full barrier out of the total heap using the residues Δ_{pre} and Δ_{post} of the entailments given in the previous constraints. Δ_{pre} and Δ_{post} are then tested for equality by requiring bi-entailment with empty residue. That is, given

$$*_{i=1}^{i=\text{thread_cnt}} \Phi_{pre}^i \vdash \text{self} :: \text{bn}^{\blacksquare} \langle \text{from} \rangle * \Delta_{pre}$$
 and

$$*_{i=1}^{i=\text{thread_cnt}} \Phi_{post}^i \vdash \text{self} :: \text{bn}^{\blacksquare} \langle \text{to} \rangle * \Delta_{post},$$
 we check

$$\Delta_{pre} \vdash \Delta_{post} \text{ and } \Delta_{post} \vdash \Delta_{pre} \text{ with empty residues.}$$
- For states with more than one successor, we check mutual exclusion for the preconditions as required by §4 by verifying that for any two preconditions of two distinct transitions must entail False . This check was extremely tedious to do for the example barrier in Coq but SLEEK can do it easily.

Once SLEEK has verified each of the above conditions, we know that the barrier definition is well-formed according to the constraints described in §4 (modulo precision).

8.6. Extension to program verification. Integrating our Hoare rule for barriers into HIP was the easiest part of adding our program logic to HIP/SLEEK. Following the concept of structured specifications [12], we transform our barrier diagrams into disjunctions of the form

$$\text{bn} ::= \bigvee (\text{requires } \Phi_{pre} \text{ ensures } \Phi_{post}),$$

where the disjunction spans all specifications in all transitions in the barrier definition. Verification of barrier calls trivially reduces to an entailment check of the disjunction.

8.7. Tool performance outline. We have developed a small set of benchmarks for our HIP/SLEEK with barriers prototype. Our SLEEK tests divide into two categories: entailment checks for separation logic formulae containing fractional permissions and checking barrier consistency checks as in §8.5. Individual entailment checks are quite speedy and our benchmark covers a number of interesting cases (*e.g.*, inductively defined predicates). Barrier consistency checks take more time but the

performance is more than adequate, as follows:

Sleek Examples	Test details	Entailments	Time (s)
<code>fractions.slk</code>	fractional entailments	54	0.08
<code>barrier.slk</code>	6 barrier definitions	279	2.3

One of the barrier definitions in `barrier.slk` is the example barrier given in Figure 1. It took **2,700 (highly tedious) lines of code** and 48 seconds of verification time (Figure 5) to convince Coq that the example barrier definition met the soundness requirements. SLEEK verifies this example barrier definition and analyzes 5 others (some sound, others not) in 2.3 seconds **without any interaction from the user**¹³.

We have also benchmarked HIP with a slightly modified variant of the example from §3. We made two modifications due to certain existing weaknesses in the HIP/SLEEK toolchain. First, we substituted recursive functions for loops due to the convenience of specification of recursive functions in HIP; and second, we changed the way x_1 , x_2 , y_1 , and y_2 are modified in lines 6 and 10 to enable the numerical decision procedures (*i.e.*, Omega) to be able to discharge the associated obligations. Both of these changes are orthogonal to our logic for barriers: for example, a more powerful decision procedure for numerical equalities would allow us to return to the original program.

We verified our modified code against three specifications. In `barrier-paper.ss`, we verify a trivial correctness property for the exact barrier definition from 3—*i.e.*, we verify the postcondition of `True`, meaning that the program does not get stuck. We also verified two more complex postconditions by using two more finely-grained barrier definitions: in `barrier-weak.ss`, we verified the relationship between x_1 , x_2 , and n ; finally, in `barrier-strong.ss` we verified the precise value in x_1 after the loop terminates (*i.e.*, $x_1 = 59$). The code and specification for `barrier-strong.ss` is given in Appendix A. We recorded the following timings from HIP:

File	Postcondition	LOC (code + specification)	Time(s)
<code>barrier-paper.ss</code>	<code>True</code>	73	2.55
<code>barrier-weak.ss</code>	lax bounds	73	2.91
<code>barrier-strong.ss</code>	exact bounds	73	3.04

As expected, the tighter bounds require more verification time; however, the differences are relatively small because most of the work is spent dealing with the heap constraints as opposed to the pure constraints. Part of the time for each example is spent verifying the correctness of the included barrier definition; all three barrier definitions from the HIP examples were also included in the `barrier.slk` benchmark.

Although we think that the HIP verification times are decent, barrier calls are fairly computationally expensive to verify due to the need to check multiple entailments. We believe that performance can be further improved by adding optimizations to SLEEK in the style of [10]. Since barrier calls are fairly rare in actual code, we believe that the performance of HIP/SLEEK on larger examples will be acceptable.

9. LIMITATIONS AND FUTURE WORK

We can extend the logic by making the barriers first-class (*i.e.*, dynamic barrier creation/destruction). In the present work we thought we could simplify the proofs by having statically declared barriers in the style of O’Hearn [22]. This turned out to be somewhat of a mistake, at least as far as the soundness proof went: since we were forced to track the barrier states (and partial shares) explicitly in the Hoare logic, we estimate that 90% of the work required to make the barriers first-class has already been done in the present work; moreover, a further 8% (the intrinsic contravariant circularity) would be easy to handle via indirection theory [17]. With perfect foresight (or if it were nontrivial to restart a large mechanized proof), we would have certainly made the barriers first-class.

¹³As explained in §8.5, SLEEK verifies slightly properties from those verified in Coq.

On the other hand, first-class barriers might have made the verification problem considerably more complicated.

We suspect that our SLEEK prototype could be improved in numerous ways. For example, our decision procedure for share formulae is quite incomplete and we believe that several performance enhancements to SLEEK would speed up the consistency checks. Finally, we need to resolve the precision/token issue.

We also do not address the tricky problem of barrier definition inference.

10. RELATED WORK

Calcagno *et al.* proposed separation algebras as models of separation logic [9]; fractional permissions were discussed by Bornat *et al.* [5]. In our work we use the share model and separation algebra development of Dockins *et al.* [11, 1].

O’Hearn’s concurrent separation logic focused on programs that used critical regions [22, 6]; subsequent work by Hobor *et al.* and Gotsman *et al.* added first-class locks and threads [16, 14, 15]. Our basic soundness techniques (unerased semantics tracks resource accounting; oracle semantics isolates sequential and concurrent reasoning from each other; etc.) follow Hobor *et al.* Recently both Villard *et al.* and Bell *et al.* extended concurrent separation logic to channels [3, 23]. The work on channels is similar to ours in that both Bell and Villard track additional dynamic state in the logic and soundness proof. Bell tracks communication histories while Villard tracks the state of a finite state automata associated with each communication channel. Of all of the previous soundness results, only Hobor *et al.* had a machine-checked soundness proof; it was incomplete.

An interesting question is whether it is possible to reason about barriers in a setting with locks or channels. The question has both an operational and a logical flavor. Speaking operationally, in a practical sense the answer is no: for performance reasons barriers are not implemented with channels or locks. If we ignore performance, however, it is possible to implement barriers with channels or locks¹⁴. The logical part of the question then becomes, are the program logics defined by O’Hearn, Hobor, Gotsman, Villard, or Bell (including their coauthors) strong enough to reason about the (implementation of) barriers in the style of the logic we have presented? As far as we can tell each previous solution is missing at least one required feature, so in a strict sense, the answer here is again no.

For illustration we examine what seems to be the closest solution to ours: the copyless message passing channels of Villard *et al.* Operationally speaking, the best way to implement barriers seems to be by adding a central authority that maintains a channel with each thread using a barrier. When a thread hits a barrier, it sends “waiting” to the central authority, and then waits until it receives “proceed”. In turn, the central authority waits for a “waiting” message from each thread, and then sends each of them a “proceed” message. Fortunately Villard allows the central authority to wait on multiple channels simultaneously.

The question then becomes a logical one. Although it should not pose any fundamental difficulty, their logic would first need to be enhanced with fractional permissions; in fact we believe that Villard’s Heap-Hop tool already uses the same fractional permission model (by Dockins *et al.*) that we do¹⁵. Since Villard uses automata to track state, we think it probable, but not certain, that our barrier state machines can be encoded as a series of his channel state machines.

There are some problems to solve. Villard requires certain side conditions on his channels; we require other kinds of side conditions on our barriers; these conditions do not seem fully compatible¹⁶. Assuming that we can weaken/strengthen conditions appropriately, we reach a second problem with

¹⁴Indeed, it is possible to implement channels and locks in terms of each other.

¹⁵To be precise, Heap-Hop uses the code extracted from the fractional permission Coq proof development by Dockins *et al.*

¹⁶For example, Villard requires determinacy whereas we do not; he would also require that the postconditions of barriers be precise whereas we do not; etc.

the side conditions: some of our side conditions (*e.g.*, mutual exclusion) are restrictions on the shape of the entire diagram; in Villard’s setting the barrier state diagram has been partitioned into numerous separate channel state machines. Verifying our side conditions seems to require verification of the relationships that these channel state machines have to each other; the exact process is unclear.

Once the matter of side conditions is settled, there remains the issue of verifying the individual threads and the central authority. Villard’s logic seems to have all that is required for the individual threads; the question is how difficult it would be to verify the central authority. Here we are less sure but suspect that with enough ghost state/instructions it can be done.

There remains a question as to whether it is a good idea to reason about barriers via channels (or locks). We suspect that it is not a good idea, even ignoring the fact that actual implementations of barriers do not use channels. The main problem seems to be a loss of intuition: by distributing the barrier state machine across numerous channel state machines and the inclusion of necessary ghost state, it becomes much harder to see what is going on. We believe that one of the major contributions of our work is that our barrier rule is extremely simple; with a quick reference to the barrier state diagram it is easy to determine what is going on. There is a secondary problem: we believe that our barrier rule will look and behave essentially the same way in a setting with first-class barriers in which it is possible to define functions that are polymorphic over the barrier diagram; even assuming a channel logic enriched in a similar way, the verification of a polymorphic central authority seems potentially formidable.

Finally, work on concurrent program analysis is in the early stages; Gotsman *et al.*, Calcagno *et al.*, and Villard *et al.* give techniques that cover some use cases involving locks and channels but much remains to be done [13, 8, 24].

Connection to a result by Jacobs and Piessens. We recently learned that Jacobs and Piessens have an impressive result on modular fine-grained concurrency [19]. Jacobs was able to reason about our example program using his VeriFast tool by designing an implementation of barriers using locks and reducing our barrier diagram to a large disjunction for a resource invariant.

However, VeriFast has some disadvantages compared to the HIP/SLEEK approach we presented. First, HIP/SLEEK required far less input from the user. In the case of our 30-line example program, more than 600 lines of annotation were required in VeriFast, not including the code/annotations for the barrier implementation itself. HIP/SLEEK were able to verify the same program with approximately 30 lines of annotation (mostly the barrier definition). Second, it was harder to gain insight into the program from the disjunction-form of the invariant; in contrast we find our barrier diagrams straightforward to understand. Finally, it is unclear to us whether the reduction is always possible or whether it was only enabled by the relative simplicity of our example program. That said, Jacobs and Piessens have the only logic and tool proven to be able to reason about barriers as derived from a more general mechanism.

11. CONCLUSION

We have designed and proved sound a program logic for Pthreads-style barriers. Our development includes a formal design for barrier definitions and a series of soundness conditions to verify that a particular barrier can be used safely. Our Hoare rules can verify threads independently, enabling a thread-modular approach. Our soundness proof defines an operational semantics that explicitly tracks permission accounting during barrier calls and is machine-checked in Coq. We have modified the verification toolset HIP/SLEEK to use our logic to verify concurrent programs that use barriers.

Our soundness results are machine-checked in Coq and are available at:

www.comp.nus.edu.sg/~hobor/barrier

Our prototype HIP/SLEEK verification tool is available at:

www.comp.nus.edu.sg/~cristian/projects/barriers/tool.html

Acknowledgements. We thank Christian Bienia for showcasing numerous example programs containing barriers, Christopher Chak for help on an early version of this work, Jules Villard for useful comments in general and in particular on the relation of our logic to the logic of his Heap-Hop tool, and Bart Jacobs for discovering how to verify our example program in his VeriFast tool.

REFERENCES

- [1] Andrew Appel, Robert Dockins, and Aquinas Hobor. Mechanized Semantic Library. Available at <http://msl.cs.princeton.edu>, 2009–2010.
- [2] Andrew W. Appel and Sandrine Blazy. Separation logic for small-step C minor. In *TPHOLS*, pages 5–21, 2007.
- [3] Christian J. Bell, Andrew W. Appel, and David Walker. Concurrent separation logic for pipelined parallelization. In *SAS*, 2010.
- [4] Christian Bienia. *Benchmarking Modern Multiprocessors*. PhD thesis, Princeton University, Department of Computer Science, Princeton, NJ, December 2010.
- [5] Richard Bornat, Cristiano Calcagno, Peter O’Hearn, and Matthew Parkinson. Permission accounting in separation logic. In *POPL*, pages 259–270, 2005.
- [6] Stephen D. Brookes. A semantics for concurrent separation logic. In *CONCUR*, pages 16–34, 2004.
- [7] David R. Butenhof. *Programming with POSIX Threads*. Addison-Wesley, 1997.
- [8] Cristiano Calcagno, Dino Distefano, and Viktor Vafeiadis. Bi-abductive resource invariant synthesis. In *APLAS*, 2009.
- [9] Cristiano Calcagno, Peter W. O’Hearn, and Hongseok Yang. Local action and abstract separation logic. In *Symposium on Logic in Computer Science*, 2007.
- [10] Wei-Ngan Chin, Cristian Gherghina, Razvan Voicu, Quang Loc Le, Florin Craciun, and Shengchao Qin. A specialization calculus for pruning disjunctive predicates to support verification. In *CAV*, pages 293–309, 2011.
- [11] Robert Dockins, Aquinas Hobor, and Andrew W. Appel. A fresh look at separation algebras and share accounting. In *APLAS*, 2009.
- [12] Cristian Gherghina, Cristina David, Shengchao Qin, and Wei-Ngan Chin. Structured specifications for better verification of heap-manipulating programs. In *FM*, pages 386–401, 2011.
- [13] Alexey Gotsman, Josh Berdine, and Byron Cook. Interprocedural Shape Analysis with Separated Heap Abstractions. In *SAS*, pages 240–260, 2006.
- [14] Alexey Gotsman, Josh Berdine, Byron Cook, Noam Rinetzkzy, and Mooly Sagiv. Local reasoning for storable locks and threads. In *APLAS*, pages 19–37, 2007.
- [15] Aquinas Hobor. Oracle semantics. Technical Report TR-836-08, Princeton, 2008.
- [16] Aquinas Hobor, Andrew W. Appel, and Francesco Zappa Nardelli. Oracle semantics for concurrent separation logic. In *ESOP*, pages 353–367, 2008.
- [17] Aquinas Hobor, Robert Dockins, and Andrew W. Appel. A theory of indirection via approximation. In *POPL ’10*, pages 171–185, 2010.
- [18] Aquinas Hobor and Cristian Gherghina. Barriers in concurrent separation logic. In *20th European Symposium of Programming (ESOP 2011)*, pages 276–296, 2011.
- [19] Bart Jacobs and Frank Piessens. Expressive modular fine-grained concurrency specification. In *POPL ’11*, page To appear, 2011.
- [20] Huu Hai Nguyen and Wei-Ngan Chin. Enhancing program verification with lemmas. In *CAV*, pages 355–369, 2008.
- [21] Huu Hai Nguyen, Cristina David, Shengchao Qin, and Wei-Ngan Chin. Automated verification of shape and size properties via separation logic. In *VMCAI*, pages 251–266, 2007.
- [22] Peter W. O’Hearn. Resources, concurrency and local reasoning. *Theoretical Computer Science*, 375(1):271–307, May 2007.
- [23] Jules Villard, Étienne Lozes, and Cristiano Calcagno. Proving copyless message passing. In *APLAS*, pages 194–209, 2009.

- [24] Jules Villard, Étienne Lozes, and Cristiano Calcagno. Tracking heaps that hop with heap-hop. In *TACAS*, pages 275–279, 2010.

APPENDIX A. EXAMPLE FROM §3 REVISITED

Below we present a slight variation on the example from section §3 that we verified with our HIP/SLEEK toolchain. In this example, we specify exact postconditions. Starting the execution with $x_1 = x_2 = 1$ will lead to $x_1 = x_2 = 59$. The example is expressed in the HIP/SLEEK input language (where [L] and [R] respectively denote the left and right half of the full share). It makes use of recursive functions instead of while loops, but this is only for aesthetic reasons.

```

data cl {int val;}

barrier bn, 2,x1 x2 y1 y2 i, // bn=barrier name, 2=thread count, x1..i shared heap
[(0,1, // transition description, start/end state
 [ requires
   x1::cl@[L]<A1>*x2::cl@[L]<B1>* y1::cl@[L]<C1>*y2::cl@[L]<D1>*
   i::cl@[L]<T1>*self::bn@[L]<0>
 ensures
   x1::cl@[L]<A1>*x2::cl@[L]<B1>* y1::cl@[L]<C1>*i::cl@[L]<T1>*
   self::bn@[L]<1> & T1 < 30;, // one pre-post
 requires
   x1::cl@[R]<A2>*x2::cl@[R]<B2>*y1::cl@[R]<C2>*y2::cl@[R]<D2>*
   i::cl@[R]<T2>*self::bn@[R]<0>
 ensures
   x1::cl@[R]<A2>*x2::cl@[R]<B2>*y2::cl@[R]<D2>*i::cl@[R]<T2>*
   self::bn@[R]<1> & T2 < 30;]),

(1,2,[
  requires
   x1::cl@[L]<A>*x2::cl@[L]<A>*y1::cl@[L]<C>* i::cl@[L]<T>* self::bn@[L]<1>&
   T<30 & A=2*T-1 & C = 3*A+2
 ensures
   x1::cl@[L]<A>*y1::cl@[L]<C>*y2::cl@[L]<D>*i::cl@[L]<T>*self::bn@[L]<2>&
   T<30 & A=2*T-1 & D=2*A & C = 3*A+2;,
 requires
   x1::cl@[R]<A>*x2::cl@[R]<A>*y2::cl@[R]<D>*i::cl@[R]<T>*self::bn@[R]<1>&
   T<30 & D=2*A & A=2*T-1
 ensures
   x2::cl@[R]<A>*y1::cl@[R]<C>*y2::cl@[R]<D>* self::bn@[R]<2> &
   D=2*A & C = 3*A+2 & A=2*T-1;]),

(2,1,[
  requires
   x2::cl@[R]<B>*y1::cl@[R]<C>*y2::cl@[R]<D>* self::bn@[R]<2> & B=C-D
 ensures
   x1::cl@[R]<A>*x2::cl@[R]<B>*y2::cl@[R]<D>*i::cl@[R]<T>*self::bn@[R]<1>&
   A=C-D & A=B & A=2*T-1 & T <= 30;,
 requires
   x1::cl@[L]<A>*y1::cl@[L]<C>*y2::cl@[L]<D>*i::cl@[L]<T>*self::bn@[L]<2>&

```

```

    A=C-D & A=2*T-1 & T <= 30
ensures
  x1::cl@[L]<A>*x2::cl@[L]<B>*y1::cl<C>*i::cl@[L]<T>*self::bn@[L]<1>&
  A=C-D & A=B & A=2*T-1 & T <= 30;)] ,

(1,3,[
  requires
    x1::cl@[L]<A>*x2::cl@[L]<B>*i::cl@[L]<T>*self::bn@[L]<1>& T=30
  ensures
    x1::cl@[L]<A>*x2::cl@[L]<B>*i::cl<T>*self::bn@[L]<3> & T=30; ,
  requires
    x1::cl@[R]<A>*x2::cl@[R]<B>*i::cl@[R]<T>*self::bn@[R]<1>& T=30
  ensures
    x1::cl@[R]<A>*x2::cl@[R]<B> *self::bn@[R]<3>;)]);

// end barrier definition, begin code

void th1 (cl x1, cl x2, cl y1, cl y2, cl i, bn b)
  requires
    x1::cl@[L]<1>*x2::cl@[L]<1>*y1::cl@[L]<_>*
    y2::cl@[L]<_>*i::cl@[L]<1>*b::bn@[L]<0>
  ensures
    x1::cl@[L]<v>*x2::cl@[L]<v>*b::bn@[L]<3>& v=59;
{
  // stage 0
  barrier b; // stage 0->1
  th1_loop (x1,x2,y1,y2,i,b);
}

void th1_loop(cl x1, cl x2, cl y1, cl y2, cl i, bn b)
  requires
    x1::cl@[L]<v>*x2::cl@[L]<v>*y1::cl<_>*i::cl@[L]<a>*
    b::bn@[L]<1> & v=2*a -1 & a <= 30
  ensures
    x1::cl@[L]<v1>*x2::cl@[L]<v1>*b::bn@[L]<3>& v1=59;
{
  if (i.val<30)
  {
    // stage 1
    y1.val = x1.val + 2*x2.val+2;
    barrier b; // stage 1->2
    x1.val = y1.val - y2.val;
    i.val= i.val+1;
    barrier b; // stage 2->1
    th1_loop (x1,x2,y1,y2,i,b);
  }
  else barrier b; // stage 1->3
}

void th2 (cl x1, cl x2, cl y1, cl y2, cl i, bn b)
  requires
    x1::cl@[R]<1>*x2::cl@[R]<1>*y1::cl@[R]<_>*y2::cl@[R]<_>*
    i::cl@[R]<1>*b::bn@[R]<0>

```

```

    ensures
      x1::cl@[R]<v>*x2::cl@[R]<v>*b::bn@[R]<3>& v=59;
  {
    barrier b; // stage 0
    th2_loop (x1,x2,y1,y2,i,b); // stage 0->1
  }

void th2_loop(cl x1, cl x2, cl y1, cl y2, cl i, bn b)
  requires
    x1::cl@[R]<v>*x2::cl@[R]<v>*y2::cl[_]>*i::cl@[R]<a>*
    b::bn@[R]<1> & v=2*a -1 & a <= 30
  ensures
    x1::cl@[R]<v1>*x2::cl@[R]<v1>*b::bn@[R]<3>& v1=59;
  {
    if (i.val<30)
    {
      y2.val = x1.val + x2.val; // stage 1
      barrier b; // stage 1->2
      x2.val = y1.val - y2.val;
      barrier b; // stage 2->1
      th2_loop (x1,x2,y1,y2,i,b);
    }
    else barrier b; // stage 1->3
  }
}

```