

Local Reasoning for Storable Locks and Threads

Alexey Gotsman¹, Josh Berdine², Byron Cook²,
Noam Rinetzky^{3,*}, and Mooly Sagiv^{2,3}

¹ University of Cambridge

² Microsoft Research

³ Tel-Aviv University

Abstract. We present a resource oriented program logic that is able to reason about concurrent heap-manipulating programs with unbounded numbers of dynamically-allocated locks and threads. The logic is inspired by concurrent separation logic, but handles these more realistic concurrency primitives. We demonstrate that the proposed logic allows local reasoning about programs for which there exists a notion of dynamic ownership of heap parts by locks and threads.

1 Introduction

We are interested in modular reasoning, both manual and automatic, about concurrent heap-manipulating programs. Striking progress in this realm has recently been made by O’Hearn [10], who proposed concurrent separation logic as a basis for reasoning about such programs. Concurrent separation logic is a Hoare logic with two novel features: the assertion language of the logic contains the $*$ connective that splits the program state into disjoint parts, and the proof system has two important rules:

$$\frac{\{P\} C \{Q\}}{\{P * R\} C \{Q * R\}} \text{FRAME} \quad \frac{\{P_1\} C_1 \{Q_1\} \quad \{P_2\} C_2 \{Q_2\}}{\{P_1 * P_2\} C_1 \parallel C_2 \{Q_1 * Q_2\}} \text{PAR}$$

According to the FRAME rule, if P includes the part of the program state that C accesses, then executing C in the presence of additional program state R results in the same behavior, and C does not touch the extra state. The PAR rule says that if two processes access disjoint parts of the program state, they can safely execute in parallel and the final state is given by the $*$ -conjunction of the post-conditions of the processes. Therefore, to reason about a command (or a process) in a program, it is sufficient to consider only the part of the program state that the command actually accesses, a feature that greatly simplifies program proofs and is referred to as the principle of *local reasoning* [9].

In the PAR rule it is intended that the processes access a finite set of shared resources using conditional critical regions to synchronize access. Process interaction is mediated in the logic by assigning to every resource an assertion – its

* Supported by the German-Israeli Foundation for Scientific Research and Development (G.I.F.).

resource invariant – that describes the part of the heap owned by the resource and must be respected by every process. For any given process, resource invariants restrict how other processes can interfere with it, and hence, the process can be reasoned about in isolation. In this way the logic allows local reasoning about programs consistent with what O’Hearn terms the Ownership Hypothesis (“A code fragment can access only those portions of state that it owns.”) [10], i.e., programs that admit a notion of ownership of heap parts by processes and resources. At the same time, the ownership relation is not required to be static, i.e., it permits ownership transfer of heap cells between areas owned by different processes and resources. The resource-oriented flavor of the logic makes it possible to use it as a basis for thread-modular program analysis [7]: certain classes of resource invariants can automatically be inferred by an abstract interpretation that analyzes each process separately in contrast to a straightforward analysis that just enumerates all execution interleavings.

However, concurrent separation logic [10], its derivatives [1,12,3], and a corresponding program analysis [7] all suffer from a common limitation: they assume a bounded number of non-aliased and pre-allocated locks (resources) and threads (processes) and, hence, cannot be used to reason about concurrency primitives present in modern languages and libraries (e.g., POSIX threads) that use unbounded numbers of *storable* locks and threads. Here “storable” means that locks can be dynamically allocated and destroyed in the heap; threads can be dynamically created and can terminate themselves, and moreover, thread identifiers can be stored and subsequently used to wait for termination of the identified thread.

Reasoning about storable locks is especially difficult. The issue here is not that of expressiveness, but of modularity: storable locks can be handled by building a global invariant describing the shared memory as a whole, with all locks allocated in it. However, in this case the locality of reasoning is lost, which kicks back in global invariants containing lots of auxiliary state, proofs being extremely complex and program analyses for discovery of global invariants being infeasible. Recent efforts towards making proofs in this style of reasoning modular [4,15] use rely-guarantee reasoning to simplify the description of the global invariant and its possible changes (see Section 9 for a detailed comparison of such techniques with our work).

What we want is a logic that preserves concurrent separation logic’s local reasoning, even for programs that manipulate storable locks and threads. To this end, in this paper we propose a logic (Section 3), based upon separation logic, that treats storable locks along with the data structures they protect as resources, assigning invariants to them and managing their dynamic creation and destruction. The challenges of designing such a logic were (quite emotionally) summarized by Bornat et al. in [1]:

...the idea of semaphores in the heap makes theoreticians wince. The semaphore has to be available to a shared resource bundle:¹ that means a bundle will contain a bundle which contains resource, a notion which

¹ Here the term “resource bundle” is used to name what we, following O’Hearn’s original paper, call “resource invariant”.

makes everybody’s eyes water. None of it seems impossible, but it’s a significant problem, and solving it will be a small triumph.

Less emotionally, stored locks are analogous to stored procedures in that, unless one is very careful, they can raise a form of Russell’s paradox, circularity arising from what Landin called knots in the store. Stored locks can do this by referring to themselves through their resource invariants, and here we address this foundational difficulty by cutting the knots in the store with an indirection.

Our approach to reasoning about storable locks is to represent a lock in the assertion language by a *handle* whose denotation cuts knots in the store. A handle certifies that a lock allocated at a certain address exists and gives a thread owning the handle a permission to (try to) acquire the lock. By using the mechanism of permissions [1] the handle can be split among several threads that can then compete for the lock. Furthermore, a handle carries some information about the part of the program state protected by the lock (its resource invariant), which lets us mediate the interaction among threads, just as in the original concurrent separation logic. Handles for locks can be stored inside resource invariants, thereby permitting reasoning about the situation described in the quote above. In this way we extend the ability of concurrent separation logic to reason locally about programs that are consistent with the Ownership Hypothesis to the setting with storable locks and threads. As we show in Section 4, the class of such programs contains programs with coarse-grained synchronization and some, but not all, programs with fine-grained synchronization, including examples that were posed as challenges in the literature.

We prove the logic sound with respect to an interleaving operational semantics (Section 7). It happens that even formulating the soundness statement is non-trivial as we have to take into account resource invariants for locks not mentioned directly in the local states of threads.

The technical issues involved in reasoning about storable locks and storable threads are similar. To make the presentation more approachable, we first present a logic for programs consisting of one top-level parallel composition of several threads. In Section 8 we extend the logic to handle dynamic thread creation.

2 Technical Background

In this section we review some technical concepts of (sequential) separation logic that we reuse in ours. We consider a version of separation logic that is a Hoare logic for a heap-manipulating programming language with the following syntax:

V	$::= l, x, y, \dots$	variables
E, F	$::= \text{nil} \mid V \mid E + F \mid \dots$	expressions
G	$::= E = F \mid E \neq F$	branch guards
C	$::= V = E \mid V = [E] \mid [E] = F \mid V = \mathbf{new} \mid \mathbf{delete} E$	primitive commands
S	$::= C \mid S; S \mid \mathbf{if} G \mathbf{then} S \mathbf{else} S \mathbf{fi} \mid \mathbf{while} G \mathbf{do} S \mathbf{od}$	commands

Here square brackets denote pointer dereferencing; the meaning of the rest of the language is standard.

Formulae in the assertion language of separation logic denote program states represented by stack-heap pairs and have the following syntax:

$$\begin{aligned} \Phi ::= & \text{false} \mid \Phi \Rightarrow \Phi \mid \exists X. \Phi \mid \Phi * \Phi \mid \Phi \multimap \Phi \mid \text{emp}_s \mid \text{emp}_h \\ & \mid E = F \mid \pi = \mu \mid \text{Own}_\pi(x) \mid E \mapsto F \end{aligned}$$

We can define usual connectives not mentioned in the syntax definition using the provided ones. Note that we treat variables as resources [12] to avoid side conditions in proof rules, i.e., we treat the stack in the same way as the heap. Thus, the assertion $E \mapsto F$ denotes the set of stack-heap pairs such that the heap consists of one cell allocated at the address E and storing the value F , and the stack contains all variables mentioned in E and F . The assertion $\text{Own}_1(x)$ (the general form $\text{Own}_\pi(x)$ is explained later) restricts the stack to contain only the variable x and leaves the heap unconstrained. We can separate assertions about variable ownership $\text{Own}_1(x)$ with $*$ in the same way as assertions $E \mapsto F$ about ownership of heap cells. emp_s describes the empty stack and emp_h the empty heap. We distinguish integer program variables x, y, \dots (which may appear in programs) and logical variables X, Y, \dots (which do not appear in programs, only in formulae). In the assertion language definition E and F range over expressions, which are the same as in the programming language, but can contain logical variables. We write $E \mapsto _$ for $\exists X. E \mapsto X$ where X does not occur free in E .

The assertion language includes fractional permissions [1] for variables, which are necessary for getting a complete (in the sense of [12]) proof system when variables are treated as resources. For clarity of presentation we omit the treatment of permissions for heap cells. Permissions are denoted with permission expressions (ranged over by π and μ), which are expressions evaluating to numbers from $(0, 1]$. A permission shows “how much” of a variable is owned by the assertion. For example, variable x represented by $\text{Own}_1(x)$ can be split into two permissions $\text{Own}_{1/2}(x)$, each of which permits reading the variable, but not writing to it. Two permissions $\text{Own}_{1/2}(x)$ can later be recombined to obtain the full permission $\text{Own}_1(x)$, which allows both reading from and writing to x . We make the convention that \Vdash binds most loosely, use $\pi_1 x_1, \dots, \pi_k x_k \Vdash P$ to denote $\text{Own}_{\pi_1}(x_1) * \dots * \text{Own}_{\pi_k}(x_k) \wedge P$ and abbreviate $1x$ to x .

The proof rules (see [6]) are the same as in [13,12] modulo treating variables as resources in heap-manipulating commands. In the rules and the following, O ranges over assertions of the form $\pi_1 x_1, \dots, \pi_k x_k$. We also allow O to be empty, in which case we interpret $O \Vdash P$ as $\text{emp}_s \wedge P$.

3 Logic

We now consider a concurrent programming language based on the sequential one presented in Section 2:

$$\begin{aligned} C ::= & \dots \mid \mathbf{init}(E) \mid \mathbf{finalize}(E) \mid \mathbf{acquire}(E) \mid \mathbf{release}(E) && \text{primitive commands} \\ P ::= & S \parallel \dots \parallel S && \text{programs} \end{aligned}$$

We assume that each program consists of one parallel composition of several threads. Synchronization is performed using locks, which are dynamically created

and destroyed in the heap. $\mathbf{init}(E)$ converts a location allocated at the address E to a lock. After the completion of $\mathbf{init}(E)$ the thread that executed it holds the lock. $\mathbf{acquire}(E)$ and $\mathbf{release}(E)$ try to acquire, respectively, release the lock allocated at the address E . $\mathbf{finalize}(E)$ converts the lock into an ordinary heap cell containing an unspecified value provided that the lock at the address E is held by the thread that is executing the command.

As in concurrent separation logic [10], with each lock we associate a resource invariant – a formula that describes the part of the heap protected by the lock. (This association is considered to be part of the proof, rather than of the program.) To deal with unbounded numbers of locks we assume that each lock has a *sort* that determines its invariant. Formally, we assume a fixed set \mathcal{L} of function symbols with positive arities representing lock sorts, and with each $A \in \mathcal{L}$ of arity k we associate a formula $I_A(L, \vec{X})$ containing k free logical variables specified as parameters – the resource invariant for the sort A . The meaning of the first parameter is fixed as the address at which the lock is allocated. Other parameters can have arbitrary meaning. In Sections 5 and 7 we give certain restrictions that resource invariant formulae must satisfy for the logic to be sound.

We extend the assertion language of separation logic with two extra forms: $\Phi ::= \dots \mid \pi A(E, \vec{F}) \mid \mathbf{Locked}_A(E, \vec{F})$. An expression of the form $A(E, \vec{F})$, where $A \in \mathcal{L}$, is a *handle* for the lock of the sort A allocated at the address E . It can be viewed as an existential permission for the lock: a thread having $A(E, \vec{F})$ knows that the heap cell at the address E is allocated and is a lock, and can try to acquire it. $A(E, \vec{F})$ does not give permissions for reading from or writing to the cell at the address E . Moreover, it does not ensure that the part of the heap protected by the lock satisfies the resource invariant until the thread successfully acquires the lock. We allow using $A(E, \vec{F})$ with fractional permissions [1] writing $\pi A(E, \vec{F})$. The intuition behind the permissions is that a handle for a lock with the full permission 1 can be split among several threads, thereby allowing them to compete for the lock. A thread having a permission for the handle less than 1 can acquire the lock; a thread having the full permission can in addition finalize the lock. We abbreviate $1A(E, \vec{F})$ to $A(E, \vec{F})$. Assertions in the code of threads can also use a special form $\mathbf{Locked}_A(E, \vec{F})$ to represent the fact that the lock at the address E is held by the thread in the surrounding code of the assertion. $\mathbf{Locked}_A(E, \vec{F})$ also ensures that the cell at the address E is allocated and is a lock of the sort A with parameters \vec{F} .

Our logic includes the proof rules of sequential separation logic and four new rules for lock-manipulating commands shown in Figure 1. We do not provide a rule for parallel composition as our programs consist of only one top-level parallel composition, and in Section 8 we instead treat dynamic thread creation. We write $\vdash \{P\} C \{Q\}$ to denote that the triple $\{P\} C \{Q\}$ is provable in our logic.

Initializing a lock (\mathbf{INIT}) converts a cell in the heap at the address E to a lock. Upon completion of $\mathbf{init}(E)$ the thread that executed it gets both the ownership (with the full permission) of the handle $A(E, \vec{F})$ for the lock and the knowledge that it holds the lock, represented by $\mathbf{Locked}_A(E, \vec{F})$. Note that for

$$\begin{array}{c}
\frac{(O \Vdash E \mapsto _) \Rightarrow \vec{F} = \vec{F}}{\{O \Vdash E \mapsto _ \} \mathbf{init}_{A, \vec{F}}(E) \{O \Vdash A(E, \vec{F}) * \mathbf{Locked}_A(E, \vec{F})\}} \text{INIT} \\
\frac{}{\{O \Vdash A(E, \vec{F}) * \mathbf{Locked}_A(E, \vec{F})\} \mathbf{finalize}(E) \{O \Vdash E \mapsto _ \}} \text{FINALIZE} \\
\frac{}{\{(O \Vdash \pi A(L, \vec{X})) \wedge L = E\} \mathbf{acquire}(E) \{(O \Vdash \pi A(L, \vec{X}) * \mathbf{Locked}_A(L, \vec{X})) * I_A(L, \vec{X})\}} \text{ACQUIRE} \\
\frac{}{\{((O \Vdash \mathbf{Locked}_A(L, \vec{X})) * I_A(L, \vec{X})) \wedge L = E\} \mathbf{release}(E) \{O \Vdash \mathbf{emp}_h\}} \text{RELEASE}
\end{array}$$

Fig. 1. Proof rules for lock-manipulating commands

the precondition $O \Vdash E \mapsto _$ to be consistent O must contain variables mentioned in E . In this and other rules we use O to supply the permissions for variables necessary for executing the command. For $\mathbf{init}_{A, \vec{F}}(E)$ commands to be safe the stack must contain variables mentioned in E and \vec{F} , hence, the premiss $(O \Vdash E \mapsto _) \Rightarrow \vec{F} = \vec{F}$ additionally requires that variables be contained in O (see [12]). An implicit side condition in the INIT rule is that in all branches of a proof of a program, the sort A of the lock and the values of parameters \vec{F} have to be chosen consistently for each \mathbf{init} command (as otherwise the conjunction rule of Hoare logic becomes unsound). This is formally enforced by annotating each \mathbf{init} command with the sort of the lock that is being created and its parameters (defined by arbitrary expressions \vec{F} over program variables). In general, the lock sort can also be computed as a function of program variables. To simplify notation we assume the sort of the lock is fixed for each \mathbf{init} command in the program. We note that although we use the sort of the lock and its parameters for conceptually different purposes (see the examples in Section 4), technically they are merely pieces of auxiliary state associated with the handle for the lock that carry some information about the resource invariant of the lock. Therefore, the annotations of lock sorts and parameters at \mathbf{init} commands can be viewed as just assignments to auxiliary cells in memory.

Finalizing a lock results in it being converted into an ordinary cell. To finalize a lock (FINALIZE) a thread has to have the full permission for the handle $A(E, \vec{F})$ associated with the lock. Additionally, the lock has to be held by the thread, i.e., $\mathbf{Locked}_A(E, \vec{F})$ has to be in its local state.

A thread can acquire a lock if it has a permission for the handle of the lock. Acquiring a lock (ACQUIRE) results in the resource invariant of the lock (with appropriately instantiated parameters) being *-conjoined to the local state of the thread. The thread also obtains the corresponding \mathbf{Locked} fact, which guarantees that it holds the lock. A thread acquiring the same lock twice deadlocks, which is enforced by $\mathbf{Locked}_A(E, \vec{F}) * \mathbf{Locked}_A(E, \vec{F})$ being inconsistent (see Section 5). Conversely, a thread can release a lock (RELEASE) only if it holds the lock, i.e., the corresponding \mathbf{Locked} fact is present in the local state of the thread. Upon releasing the lock the thread gives up both this knowledge and the ownership of the resource invariant associated with the lock. The fact that resource invariants

<pre> struct RECORD { LOCK Lock; int Data; }; main() { RECORD *x; {x ⊢ emp_h} x = new RECORD; {x ⊢ x ↦₋ * x.Data ↦₋} init_R(x); {x ⊢ x.Data ↦₋ * R(x) * Locked_R(x)} </pre>	<pre> x->Data = 0; {x ⊢ x.Data ↦₀ * R(x) * Locked_R(x)} release(x); {x ⊢ R(x)} // ... {x ⊢ R(x)} // ... {x ⊢ R(x)} acquire(x); {x ⊢ x.Data ↦₋ * R(x) * Locked_R(x)} x->Data++; {x ⊢ x.Data ↦₋ * R(x) * Locked_R(x)} </pre>	<pre> release(x); {x ⊢ R(x)} // ... {x ⊢ R(x)} acquire(x); {x ⊢ x.Data ↦₋ * R(x) * Locked_R(x)} finalize(x); {x ⊢ x ↦₋ * x.Data ↦₋} delete x; {x ⊢ emp_h} </pre>
$I_R(L) \triangleq \text{emp}_s \wedge L.Data \mapsto _$		

Fig. 2. A very simple example of reasoning in the logic

can claim ownership of program variables complicates the rules ACQUIRE and RELEASE. E.g., in the postcondition of ACQUIRE we cannot put $I_A(L, \vec{X})$ inside the expression after \vdash as it may claim ownership of variables not mentioned in O . This requires us to use a logical variable L in places where the expression E would have been expected.

4 Examples of Reasoning

We first show (in Example 1 below) that straightforward application of rules for lock-manipulating commands allows us to handle programs in which locks protect parts of the heap without other locks allocated in them. We then present two more involved examples of using the logic, which demonstrate how extending the logic with storable locks has enabled reasoning more locally than was previously possible in some interesting cases (Examples 2 and 3).

Instead of the minimalistic language presented in Section 3, in our examples we use a language with some additional C-like syntax (in particular, C structures) that can easily be desugared to the language of Section 3. For an address x of a structure, we use $x.F$ in the assertion language as syntactic sugar for $x + d$, where d is the offset of the field F in the structure. We assume that each field in a structure takes one memory cell. We also use an obvious generalization of **new** and **delete** that allocate and deallocate several memory cells at once.

Example 1: A simple situation. Figure 2 shows a proof outline for a program with a common pattern: a lock-field in a structure protecting another field in the same structure. We use a lock sort R with invariant $I_R(L)$. The proof outline shows how the “life cycle” of a lock is handled in our proof system: creating a cell, converting it to a lock, acquiring and releasing the lock, converting it to an ordinary cell, and disposing the cell. For simplicity we consider a program with only one thread.

Example 2: “Last one disposes”. This example was posed as a challenge for local reasoning in [1]. The program in Figure 3 represents a piece of multicasting code:

```

struct PACKET {
  LOCK Lock;
  int Count;
  DATA Data;
};

PACKET *p;

initialize() {
  {p ⊢ emph}
  p = new PACKET;
  {p ⊢ p ↦ _ * p.Count ↦ _ * p.Data ↦ _}
  p->Count = 0;
  {p ⊢ p ↦ _ * p.Count ↦ 0 * p.Data ↦ _}
  initP,M(p);
  {p ⊢ p.Count ↦ 0 * p.Data ↦ _ * P(p, M) *
   LockedP(p, M)}
  // ...Initialize data...
  release(p);
  {p ⊢ P(p, M)}
}

thread() {
  {(1/M)p ⊢ (1/M)P(p, M)}
  acquire(p);
  {(1/M)p ⊢ ∃X. 0 ≤ X < M ∧ p.Count ↦ X *
   p.Data ↦ _ * ((X+1)/M)P(p, M) * LockedP(p, M)}
  // ...Process data...
  p->Count++;
  {(1/M)p ⊢ ∃X. 1 ≤ X ≤ M ∧ p.Count ↦ X *
   p.Data ↦ _ * (X/M)P(p, M) * LockedP(p, M)}
  if (p->Count == M) {
    {(1/M)p ⊢ p.Count ↦ M * p.Data ↦ _ *
     P(p, M) * LockedP(p, M)}
    // ...Finalize data...
    finalize(p);
    {(1/M)p ⊢ p.Count ↦ M * p.Data ↦ _ * p ↦ _}
    delete p;
  } else {
    {(1/M)p ⊢ ∃X. 1 ≤ X < M ∧ p.Count ↦ X *
     p.Data ↦ _ * (X/M)P(p, M) * LockedP(p, M)}
    release(p);
  }
  {(1/M)p ⊢ emph}
}

```

$$I_P(L, M) \triangleq \text{emp}_s \wedge \exists X. X < M \wedge L.Count \rightarrow X * L.Data \rightarrow _ * \\ ((X = 0 \wedge \text{emp}_h) \vee (X \geq 1 \wedge (X/M)P(L, M)))$$

Fig. 3. Proof outline for the “Last one disposes” program

a single packet p (of type *PACKET*) with *Data* inside the packet is distributed to M threads at once. For efficiency reasons instead of copying the packet, it is shared among threads. A *Count* of access permissions protected by *Lock* is used to determine when everybody has finished and the packet can be disposed. The program consists of a top level parallel composition of M calls to the procedure *thread*. Here M is a constant assumed to be greater than 0. For completeness, we also provide the procedure *initialize* that can be used to initialize the packet and thereby establish the precondition of the program.

To prove the program correct the resource invariant for the lock at the address p has to contain a partial permission for the handle of *the same lock*. This is formally represented by a lock sort P with the resource invariant $I_P(L, M)$. Initially the resource invariant contains no permissions of this kind and the handle $P(p, M)$ for the lock is split among M threads (hence, the precondition of each thread is $(1/M)p \vdash (1/M)P(p, M)$). Each thread uses the handle to acquire the lock and process the packet. When a thread finishes processing and releases the lock, it transfers the permission for the handle it owned to the resource invariant of the lock. The last thread to process the packet can then get the full permission for the lock by combining the permission in the invariant with its own one and can therefore dispose the packet.

Example 3: Lock coupling list. We next consider a fine-grained implementation of a singly-linked list with concurrent access, whose nodes store integer keys. The program (Figures 4 and 5) consists of M operations *add* and *remove* running


```

locate(int e) {
  NODE *prev, *curr;
  {O ⊢ -∞ < e ∧ (1/M)H(head)}
  prev = head;
  {O ⊢ -∞ < e ∧ prev = head ∧ (1/M)H(head)}
  acquire(prev);
  {O ⊢ ∃V'. -∞ < e ∧ -∞ < V' ∧ (1/M)H(head) * LockedH(prev) *
   ∃X. prev.Val → -∞ * prev.Next → X * N(X, V')}
  curr = prev → Next;
  {O ⊢ ∃V'. -∞ < e ∧ -∞ < V' ∧ (1/M)H(head) * LockedH(prev) *
   prev.Val → -∞ * prev.Next → curr * N(curr, V')}
  acquire(curr);
  {O ⊢ ∃V'. -∞ < e ∧ -∞ < V' ∧ (1/M)H(head) * N(curr, V') * LockedH(prev) *
   LockedN(curr, V') * prev.Val → -∞ * prev.Next → curr * curr.Val → V' *
   ((curr.Next → nil ∧ V' = +∞) ∨ (∃X, V''. curr.Next → X * N(X, V'') ∧ V' < V''))}
  while (curr → Val < e) {
    {O ⊢ ∃V, V'. V' < e ∧ (1/M)H(head) * N(curr, V') * LockedN(curr, V') *
     (LockedH(prev) ∧ V = -∞ ∨ LockedN(prev, V)) * prev.Val → V * prev.Next → curr *
     ∃X, V''. curr.Val → V' * curr.Next → X * N(X, V'') ∧ V < V' < V''}
    release(prev);
    {O ⊢ ∃X, V', V''. V' < e ∧ V' < V'' ∧ (1/M)H(head) * LockedN(curr, V') *
     curr.Val → V' * curr.Next → X * N(X, V'')}
    prev = curr;
    curr = curr → Next;
    {O ⊢ ∃V, V'. V < e ∧ V < V' ∧ (1/M)H(head) * LockedN(prev, V) *
     prev.Val → V * prev.Next → curr * N(curr, V')}
    acquire(curr);
    {O ⊢ ∃V, V'. V < e ∧ V < V' ∧ (1/M)H(head) * LockedN(prev, V) *
     LockedN(curr, V') * N(curr, V') * prev.Val → V * prev.Next → curr * curr.Val → V' *
     ((V' = +∞ ∧ curr.Next → nil) ∨ ∃X, V''. curr.Next → X * N(X, V'') ∧ V' < V'')}
  }
  {O ⊢ ∃V, V'. V < e ≤ V' ∧ (1/M)H(head) * LockedN(prev, V) * LockedN(curr, V') * N(curr, V') *
   prev.Val → V * prev.Next → curr * curr.Val → V' * ((V' = +∞ ∧ curr.Next → nil) ∨
   ∃X, V''. curr.Next → X * N(X, V'') ∧ V' < V'')}
  return (prev, curr);
}

```

$$\begin{aligned}
I_H(L) &\triangleq \text{emp}_s \wedge \exists X, V'. L.Val \rightarrow -\infty * L.Next \rightarrow X * N(X, V') \wedge -\infty < V' \\
I_N(L, V) &\triangleq \text{emp}_s \wedge ((L.Val \rightarrow V * L.Next \rightarrow \text{nil} \wedge V = +\infty) \vee \\
&\quad (\exists X, V'. L.Val \rightarrow V * L.Next \rightarrow X * N(X, V') \wedge V < V'))
\end{aligned}$$

Fig. 4. Proof outline for a part of the lock coupling list program. Here O is e , $prev$, $curr$, $(1/M)head$.

in parallel. The operations add and remove an element with the given key to or from the list. Traversing the list uses lock coupling: the lock on one node is not released until the next node is locked. The list is sorted and the first and last nodes in it are sentinel nodes that have values $-\infty$, respectively, $+\infty$. It is initialized by the code in procedure *initialize*. We only provide a proof outline for the procedure *locate* (Figure 4), which is invoked by other procedures to traverse the list. We use lock sorts H (for the head node) and N (for all other nodes) with the invariants $I_H(L)$ and $I_N(L, V)$. In this example the resource invariant for the lock protecting a node in the list holds a handle for the lock protecting the next node in the list. The full permission for $N(X, V')$ in the invariants above essentially means that the only way a thread can lock a node is by first locking its predecessor: here the invariant enforces a particular locking policy.

```

struct NODE { LOCK Lock;
              int Val;
              NODE *Next; }

NODE *head;

initialize() {
  NODE *last;
  last = new NODE;
  last->Val = INFINITY;
  last->Next = NULL;
  initN,+∞(last);
  release(last);
  head = new NODE;
  head->Val = -INFINITY;
  head->Next = last;
  initH(head);
  release(head);
}

add(int e) {
  NODE *n1, *n2, *n3, *result;
  (n1, n3) = locate(e);
  if (n3->Val != e) {
    n2 = new NODE;
    n2->Val = e;
    n2->Next = n3;
    initN,e(n2);
    release(n2);
    n1->Next = n2;
    result = true;
  } else {
    result = false;
  }
  release(n1);
  release(n3);
  return result;
}

remove(int e) {
  NODE *n1, *n2, *n3;
  NODE *result;
  (n1, n2) = locate(e);
  if (n2->Val == e) {
    n3 = n2->Next;
    n1->Next = n3;
    finalize(n2);
    delete n2;
    result = true;
  } else {
    result = false;
  }
  release(n1);
  return result;
}

```

Fig. 5. Lock coupling list program. The procedure *locate* is shown in Figure 4.

$$\begin{array}{ll}
\text{nil} = 0 & \text{Values} = \{\dots, -1, 0, 1, \dots\} \\
\text{Perms} = [0, 1] & \text{Vars} = \{x, y, \dots\} \\
\text{Stacks} = \text{Vars} \rightarrow_{\text{fin}} (\text{Values} \times \text{Perms}) & \text{Locs} = \{1, 2, \dots\} \\
\text{LockPerms} = [0, 1] & \text{ThreadIDs} = \{1, 2, \dots\} \\
\text{LockVals} = \{\mathbf{U}, 0\} \cup \text{ThreadIDs} & \text{States} = \text{Stacks} \times \text{Heaps} \\
\text{Heaps} = \text{Locs} \rightarrow_{\text{fin}} (\text{Cell}(\text{Values}) \cup \text{Lock}(\mathcal{L} \times \text{LockVals} \times \text{LockPerms}) \\
& \quad \setminus \text{Lock}(\mathcal{L} \times \{\mathbf{U}\} \times \{0\}))
\end{array}$$

Fig. 6. Model of the assertion language

We were able to present modular proofs for the programs above because in each case we could associate with every lock a part of the heap such that a thread accessed the part only when it held the lock, that is, the lock owned the part of the heap. We note that we would not be able to give modular proofs to programs that do not obey this policy, for instance, to optimistic list [14] – another fine-grained implementation of the list from Example 3 in which the procedure *locate* first traverses the list without taking any locks and then validates the result by locking two candidate nodes and re-traversing the list to check that they are still present and adjacent in the list.

5 Model of the Assertion Language

As usual, assertion language formulae denote sets of pairs of a stack and a heap, both represented by finite partial functions. They are interpreted over the domain in Figure 6. However, in contrast to the standard domain used in separation logic, here cells in the heap can be of two types: ordinary cells (*Cell*) and locks (*Lock*). A lock has a sort, a value, and is associated with a permission from $[0, 1]$. To simplify notation, here and in the further semantic development we assume that lock sorts have no parameters other than the address of the lock. Our results can straightforwardly be adjusted to the general case (parameters can be treated in

$$\begin{aligned}
(s, h, i) \models_k E \mapsto F &\Leftrightarrow \llbracket E \rrbracket_{(s,i)} \downarrow \wedge \llbracket F \rrbracket_{(s,i)} \downarrow \wedge h = \llbracket [E] \rrbracket_{(s,i)} : \mathbf{Cell}(\llbracket [F] \rrbracket_{(s,i)}) \\
(s, h, i) \models_k \mathbf{Own}_\pi(x) &\Leftrightarrow \exists u. \llbracket \pi \rrbracket_{(s,i)} \downarrow \wedge s = [x : (u, \llbracket \pi \rrbracket_{(s,i)})] \wedge 0 < \llbracket \pi \rrbracket_{(s,i)} \leq 1 \\
(s, h, i) \models_k \pi A(E) &\Leftrightarrow \\
&\llbracket E \rrbracket_{(s,i)} \downarrow \wedge \llbracket \pi \rrbracket_{(s,i)} \downarrow \wedge h = \llbracket [E] \rrbracket_{(s,i)} : \mathbf{Lock}(A, \mathbf{U}, \llbracket \pi \rrbracket_{(s,i)}) \wedge 0 < \llbracket \pi \rrbracket_{(s,i)} \leq 1 \\
(s, h, i) \models_k \mathbf{Locked}_A(E) &\Leftrightarrow \llbracket E \rrbracket_{(s,i)} \downarrow \wedge h = \llbracket [E] \rrbracket_{(s,i)} : \mathbf{Lock}(A, k, 0) \\
(s, h, i) \models_k \mathbf{emp}_s &\Leftrightarrow s = [] \\
(s, h, i) \models_k \mathbf{emp}_h &\Leftrightarrow h = [] \\
(s, h, i) \models_k E = F &\Leftrightarrow \llbracket E \rrbracket_{(s,i)} \downarrow \wedge \llbracket F \rrbracket_{(s,i)} \downarrow \wedge \llbracket E \rrbracket_{(s,i)} = \llbracket F \rrbracket_{(s,i)} \\
(s, h, i) \models_k \pi = \mu &\Leftrightarrow \llbracket \pi \rrbracket_{(s,i)} \downarrow \wedge \llbracket \mu \rrbracket_{(s,i)} \downarrow \wedge \llbracket \pi \rrbracket_{(s,i)} = \llbracket \mu \rrbracket_{(s,i)} \\
(s, h, i) \models_k P \Rightarrow Q &\Leftrightarrow ((s, h, i) \models_k P) \Rightarrow ((s, h, i) \models_k Q) \\
(s, h, i) \models_k \mathbf{false} &\Leftrightarrow \mathbf{false} \\
(s, h, i) \models_k P * Q &\Leftrightarrow \\
&\exists s_1, h_1, s_2, h_2. s = s_1 * s_2 \wedge h = h_1 * h_2 \wedge (s_1, h_1, i) \models_k P \wedge (s_2, h_2, i) \models_k Q \\
(s, h, i) \models_k P \multimap Q &\Leftrightarrow \\
&\forall s', h'. s \# s' \wedge h \# h' \wedge ((s', h', i) \models_k P) \Rightarrow ((s * s', h * h', i) \models_k Q) \\
(s, h, i) \models_k \exists X. P &\Leftrightarrow \exists u. (s, h, i[X : u]) \models_k P
\end{aligned}$$

Fig. 7. Satisfaction relation for the assertion language formulae: $(s, h, i) \models_k \Phi$

the same way as lock sorts). The permission 0 is used to represent the existential permission for a lock that is carried by $\mathbf{Locked}_A(E, \vec{F})$. Locks are interpreted as follows: 0 represents the fact that the lock is not held by any thread (i.e., is *free*), values from ThreadIDs represent the identifier of the thread that holds the lock, and \mathbf{U} means that the status of the lock is unknown. \mathbf{U} is not encountered in the states obtained in the operational semantics we define in Section 6, but is used for interpreting formulae representing parts of complete states. The semantics of formulae and commands never encounter locks of form $\mathbf{Lock}(A, \mathbf{U}, 0)$ for any A , and so the definition of Heaps removes them in order to make the $*$ operation on states cancellative [3].

Note how Heaps in the domain of Figure 6 is not defined recursively, but instead uses an indirection through \mathcal{L} , whose elements are associated with resource invariants, and hence indirectly to Heaps. It is this indirection that deals with the foundational circularity issue raised by locks which may refer to themselves.

In this paper we use the following notation for partial functions: $f(x) \downarrow$ means that the function f is defined on x , $f(x) \uparrow$ that the function f is undefined on x , and $[]$ denotes a nowhere-defined function. Furthermore, we denote with $f[x : y]$ (defined only if $f(x) \uparrow$) the function that has the same value as f everywhere, except for x , where it has the value y . We abbreviate $[] [x : y]$ to $[x : y]$.

We now define $*$ on states in our domain, which interprets the $*$ -connective in the logic. We first define the $*$ operation on values of locks in the following way: $\mathbf{U} * \mathbf{U} = \mathbf{U}$, $k * \mathbf{U} = \mathbf{U} * k = k$, and $k * j$ is undefined for $k, j \in \{0\} \cup \text{ThreadIDs}$. Note that $k * k$ is undefined as it arises in the cases when a thread tries to acquire a lock twice (recall that we specify that a thread deadlocks in this case).

For $s_1, s_2 \in \text{Stacks}$ let

$$s_1 \# s_2 \Leftrightarrow \forall x. s_1(x) \downarrow \wedge s_2(x) \downarrow \Rightarrow (\exists v, \pi_1, \pi_2. s_1(x) = (v, \pi_1) \wedge s_2 = (v, \pi_2) \wedge \pi_1 + \pi_2 \leq 1).$$

If $s_1 \# s_2$, then

$$s_1 * s_2 = \{(x, (v, \pi)) \mid (s_1(x) = (v, \pi) \wedge s_2(x) \uparrow) \vee (s_2(x) = (v, \pi) \wedge s_1(x) \uparrow) \vee (s_1(x) = (v, \pi_1) \wedge s_2(x) = (v, \pi_2) \wedge \pi = \pi_1 + \pi_2)\} ,$$

otherwise $s_1 * s_2$ is undefined. For $h_1, h_2 \in \text{Heaps}$ let

$$h_1 \# h_2 \Leftrightarrow \forall u. h_1(u) \downarrow \wedge h_2(u) \downarrow \Rightarrow ((\exists v. h_1(u) = h_2(u) = \text{Cell}(v)) \vee (\exists A, v_1, v_2, \pi_1, \pi_2. h_1(u) = \text{Lock}(A, v_1, \pi_1) \wedge h_2(u) = \text{Lock}(A, v_2, \pi_2) \wedge v_1 * v_2 \downarrow \wedge \pi_1 + \pi_2 \leq 1)) .$$

If $h_1 \# h_2$, then

$$h_1 * h_2 = \{(u, \text{Cell}(v)) \mid h_1(u) = \text{Cell}(v) \vee h_2(u) = \text{Cell}(v)\} \cup \{(u, \text{Lock}(A, v, \pi)) \mid (h_1(u) = \text{Lock}(A, v, \pi) \wedge h_2(u) \uparrow) \vee (h_2(u) = \text{Lock}(A, v, \pi) \wedge h_1(u) \uparrow) \vee (h_1(u) = \text{Lock}(A, v_1, \pi_1) \wedge h_2(u) = \text{Lock}(A, v_2, \pi_2) \wedge \pi = \pi_1 + \pi_2 \wedge v = v_1 * v_2)\} ,$$

otherwise $h_1 * h_2$ is undefined. We lift $*$ to states and sets of states pointwise.

The satisfaction relation for the assertion language formulae is defined in Figure 7. A formula is interpreted with respect to a thread identifier $k \in \{0\} \cup \text{ThreadIDs}$, a stack s , a heap h , and an interpretation i mapping logical variables to Values. Note that in this case it is convenient for us to consider 0 as a dummy thread identifier. We assume a function $\llbracket E \rrbracket_{(s,i)}$ that evaluates an expression with respect to the stack s and the interpretation i . We consider only interpretations that define the value of every logical variable used. We omit i when s suffices to evaluate the expression. We let $\llbracket P \rrbracket_i^k$ denote the set of states in which the formula P is valid with respect to the thread identifier k and the interpretation i and let $\mathcal{I}_k(A, u) = \llbracket I_A(L) * \text{Locked}_A(L) \rrbracket_{[L:u]}^k$.

We say that a predicate $p \subseteq \text{States}$ is *precise* [10] if for any state σ , there exists at most one substate σ_0 (i.e., $\sigma = \sigma_0 * \sigma_1$ for some σ_1) satisfying p . We say that a predicate p is *intuitionistic* [8] if it is closed under stack and heap extension: if p is true of a state σ_1 , then for any state σ_2 , such that $\sigma_1 * \sigma_2$ is defined, p is also true of $\sigma_1 * \sigma_2$. We say that a predicate p *has an empty lockset* if the value of any lock in every state satisfying p is \mathbf{U} . A formula is precise, intuitionistic, or has an empty lockset if its denotation with respect to any thread identifier and interpretation of logical variables is precise, intuitionistic, or has an empty lockset. We require that formulae representing resource invariants be precise and have an empty lockset, i.e., that for each u and k the predicate $\llbracket I_A(L) \rrbracket_{[L:u]}^k$ be precise and have an empty lockset. The former requirement is inherited from concurrent separation logic, where it is required for soundness of the conjunction rule. The latter requirement is necessary for soundness of our logic and stems from the fact that in our semantics we do not allow a thread that did not acquire a lock to release it (in agreement with the semantics of mutexes in the POSIX threads library). If we were to allow this (i.e., if we treated locks as binary semaphores rather than mutexes), then this requirement would not be necessary. It is easy to check that the invariants for lock sorts R , P , H , and N from Section 4 satisfy these constraints.

$x = E, (s[x : (u, 1)], h)$	$\rightsquigarrow_k (s[x : (\llbracket E \rrbracket_{s[x:(u,1)]}, 1)], h)$
$x = [E], (s[x : (u, 1)], h[e : \text{Cell}(v)])$	$\rightsquigarrow_k (s[x : (v, 1)], h[e : \text{Cell}(v)]), e = \llbracket E \rrbracket_{s[x:(u,1)]}$
$[E] = F, (s, h[\llbracket E \rrbracket_s : \text{Cell}(u)])$	$\rightsquigarrow_k (s, h[\llbracket E \rrbracket_s : \text{Cell}(\llbracket F \rrbracket_s)])$
$x = \mathbf{new}, (s[x : (u, 1)], h)$	$\rightsquigarrow_k (s[x : (v, 1)], h[v : \text{Cell}(w)]), \text{ if } h(v) \uparrow$
$\mathbf{delete} E, (s, h[\llbracket E \rrbracket_s : \text{Cell}(u)])$	$\rightsquigarrow_k (s, h)$
$\mathbf{assume}(G), (s, h)$	$\rightsquigarrow_k (s, h), \text{ if } \llbracket G \rrbracket_s = \mathbf{true}$
$\mathbf{assume}(G), (s, h)$	$\not\rightsquigarrow_k \text{ if } \llbracket G \rrbracket_s = \mathbf{false}$
$\mathbf{init}_A(E), (s, h[\llbracket E \rrbracket_s : \text{Cell}(u)])$	$\rightsquigarrow_k (s, h[\llbracket E \rrbracket_s : \text{Lock}(A, k, 1)])$
$\mathbf{finalize}(E), (s, h[\llbracket E \rrbracket_s : \text{Lock}(A, k, 1)])$	$\rightsquigarrow_k (s, h[\llbracket E \rrbracket_s : \text{Cell}(u)])$
$\mathbf{acquire}(E), (s, h[\llbracket E \rrbracket_s : \text{Lock}(A, 0, \pi)])$	$\rightsquigarrow_k (s, h[\llbracket E \rrbracket_s : \text{Lock}(A, k, \pi)])$
$\mathbf{acquire}(E), (s, h[\llbracket E \rrbracket_s : \text{Lock}(A, j, \pi)])$	$\not\rightsquigarrow_k \text{ if } j > 0$
$\mathbf{release}(E), (s, h[\llbracket E \rrbracket_s : \text{Lock}(A, k, \pi)])$	$\rightsquigarrow_k (s, h[\llbracket E \rrbracket_s : \text{Lock}(A, 0, \pi)])$
$C, (s, h)$	$\rightsquigarrow_k \top, \text{ otherwise}$

Fig. 8. Transition relation for atomic commands. $\not\rightsquigarrow_k$ is used to denote that the command does not fault, but gets stuck. \top indicates that the command faults.

6 Interleaving Operational Semantics

Consider a program S consisting of a parallel composition of n threads. We abstract away from the particular syntax of the programming language and represent each thread by its control-flow graph (CFG). A CFG over a set \mathcal{C} of atomic commands is defined as a tuple $(N, F, \mathbf{start}, \mathbf{end})$, where N is the set of program points, $F \subseteq N \times \mathcal{C} \times N$ the control-flow relation, \mathbf{start} and \mathbf{end} distinguished start and end program points. We note that a command in our language can be translated to a CFG. Conditional expressions in **if** and **while** commands are translated using the **assume**(G) statement that acts as a filter on the state space of programs – G is assumed to be true after **assume**(G) is executed. We let the set of atomic commands consist of primitive commands and the **assume** command. Let $(N_k, F_k, \mathbf{start}_k, \mathbf{end}_k)$ be the CFG of thread with identifier k and let $N = \bigcup_{k=1}^n N_k$ and $F = \bigcup_{k=1}^n F_k$. First, for each thread $k = 1..n$ and atomic command C we define a transition relation \rightsquigarrow_k shown in Figure 8.

The interleaving operational semantics of the program S is defined by a transition relation \rightarrow_S that transforms pairs of program counters (represented by mappings from thread identifiers to program points) $\mathbf{pc} \in \{1, \dots, n\} \rightarrow N$ and states $\sigma \in \text{States} \cup \{\top\}$. The relation \rightarrow_S is defined as the least one satisfying:

$$\frac{(v, C, v') \in F \quad k \in \{1, \dots, n\} \quad C, (s, h) \rightsquigarrow_k \sigma}{\mathbf{pc}[k : v], (s, h) \rightarrow_S \mathbf{pc}[k : v'], \sigma} .$$

We denote with \rightarrow_S^* the reflexive and transitive closure of \rightarrow_S . Let us denote with \mathbf{pc}_0 the initial program counter $[1 : \mathbf{start}_1] \dots [n : \mathbf{start}_n]$ and with \mathbf{pc}_f the final one $[1 : \mathbf{end}_1] \dots [n : \mathbf{end}_n]$. We say that the program S is *safe* when run from an initial state σ_0 if it is not the case that for some \mathbf{pc} we have $\mathbf{pc}_0, \sigma_0 \rightarrow_S^* \mathbf{pc}, \top$.

```

{x, y ⊢ x ↦ * y ↦ *}
initA,y(x);
initB,x(y);
{x, y ⊢ A(x, y) * LockedA(x, y) * B(y, x) * LockedB(y, x)}
release(x);
{x, y ⊢ A(x, y) * LockedB(y, x)}
release(y);
{x, y ⊢ emph}

```

$$I_A(X, Y) \triangleq \text{emp}_s \wedge B(Y, X) \quad \text{and} \quad I_B(X, Y) \triangleq \text{emp}_s \wedge A(Y, X)$$

Fig. 9. A pathological situation

7 Soundness

As it stands now, the logic allows some unpleasant situations to happen: in certain cases the proof system may not be able to detect a memory leak. Figure 9 shows an example of this kind. We assume defined lock sorts A and B with invariants $I_A(X, Y)$ and $I_B(X, Y)$. In this case the knowledge that the locks at the addresses x and y exist is lost by the proof system: the invariant for the lock x holds the full permission for the handle of the lock y and vice versa, hence, local states of the threads are then left without any permissions for the locks whatsoever.

Situations such as the one described above make the formulation of the soundness statement for our logic non-trivial. We first formulate a soundness statement (Theorem 1) showing that every final state of a program (according to the operational semantics of Section 6) can be obtained as the $*$ -conjunction of the postconditions of threads and the resource invariants *for the free locks allocated in the state*. Note that here a statement about a state uses the information about the free locks allocated in the same state. We then put restrictions on resource invariants that rule out situations similar to the one shown in Figure 9 and formulate a soundness statement (Theorem 4) in which the set of free locks in a final state is computed solely from the postconditions of threads.

For a state σ let $\text{Free}(\sigma)$, respectively, $\text{Unknown}(\sigma)$ be the set of pairs from $\mathcal{L} \times \text{Locs}$ consisting of sorts and addresses of locks allocated in the state that have value 0, respectively, \mathbf{U} . We denote with \otimes iterated separating conjunction [13]: $\otimes_{j=1}^k P_j = (\text{emp}_s \wedge \text{emp}_h) * P_1 * \dots * P_k$. The soundness of the logic with respect to the interleaving operational semantics from Section 6 is established by:

Theorem 1. *Let S be the program $C_1 \parallel \dots \parallel C_n$ and suppose $\vdash \{P_k\} C_k \{Q_k\}$ for $k = 1..n$. Then for any interpretation i and state σ_0 such that $\sigma_0 \in (\otimes_{k=1}^n \llbracket P_k \rrbracket_i^k) * (\otimes_{(A,u) \in \text{Free}(\sigma_0)} \mathcal{I}_0(A, u))$ the program S is safe when run from σ_0 and if $\text{pc}_0, \sigma_0 \rightarrow_S^* \text{pc}_f, \sigma$, then $\sigma \in (\otimes_{k=1}^n \llbracket Q_k \rrbracket_i^k) * (\otimes_{(A,u) \in \text{Free}(\sigma)} \mathcal{I}_0(A, u))$.*

The proof is given in a companion Technical Report [6]. We do not follow Brookes's original proof of soundness of concurrent separation logic [2]. Instead, we prove soundness with the aid of an intermediate thread-local semantics defined by fixed-point equations that can be viewed as the scheme of a thread-modular program analysis in the style of [7]. This method of proving soundness

should facilitate designing program analyses based on our logic. The idea of our proof, however, is close to that of Brookes’s and consists of establishing what is called the Separation Property in [10] and is formalized as the Parallel Decomposition Lemma in [2].²At any time, the state of the program can be partitioned into that owned by each thread and each free lock. As a direct consequence of the Separation Property, we can also show that provability of a program in our proof system ensures the absence of data races (see [6] for details).

We now proceed to formulate a soundness statement in which the component $\otimes_{(A,u) \in \text{Free}(\sigma)} \mathcal{I}_0(A, u)$ from Theorem 1 representing the resource invariants for free locks in the final state is obtained directly from the thread postconditions Q_k . To this end, we introduce an auxiliary notion of closure. Intuitively, closing a state amounts to $*$ -conjoining it to the invariants of all free locks whose handles are reachable via resource invariants from the handles present in the state.

Definition 2 (Closure). For $p \subseteq \text{States}$ let $c(p) \subseteq \text{States}$ be the least predicate such that $p \cup \{\sigma_1 * \sigma_2 \mid \sigma_1 \in c(p) \wedge \sigma_2 \in \otimes_{(A,u) \in \text{Unknown}(\sigma_1)} \mathcal{I}_0(A, u)\} \subseteq c(p)$. The closure $\langle p \rangle$ of p is the set of states from $c(p)$ that do not contain locks with the value \mathbf{U} .

In general, the closure is not guaranteed to add invariants for all the free locks allocated in the state. For example, the closure of the postcondition of the program in Figure 9 still has an empty heap while in the final states obtained by executing the operational semantics there are locks allocated at addresses x and y . The problem is that there may exist a “self-contained” set of free locks (containing the locks at the addresses x and y in our example) such that the corresponding resource invariants hold full permissions for all the locks from the set. Local states of threads are then left without any permissions for the locks in the set, and hence, closure is not able to reach to their invariants. The following condition on resource invariants ensures that this does not happen.

Definition 3 (Admissibility of resource invariants). Resource invariants for a set of lock sorts \mathcal{L} are admissible if there do not exist non-empty set $L \subseteq \mathcal{L} \times \text{Locs}$ and state $\sigma \in \otimes_{(A,u) \in L} \mathcal{I}_0(A, u)$ such that for all $(A, u) \in L$ the permission associated with the lock at the address u in σ is 1.

Definitions 2 and 3 generalize to the case when resource invariants have more than one parameter in the obvious way. Revisiting Example 3 of Section 4, we can check that any state satisfying the closure of $\llbracket O \Vdash (1/M)H(\text{head}) \rrbracket^k$ for any thread identifier k represents an acyclic sorted list starting at head . It is easy to check that resource invariants for the set of lock sorts $\{R, P, H, N\}$ from Section 4 are admissible whereas those for $\{A, B\}$ from this section are not. The admissibility of N is due to the fact that I_N implies sortedness of lists built out of resource invariants for N , hence, the invariants cannot form a cycle.

We say that a state is *complete* if permissions associated with all the locks allocated in it are equal to 1. Note that according to the semantics in Section 6,

² We call it the Over-approximation Lemma in [6] due to the analogy between our proof and proofs of soundness of program analyses based on abstract interpretation.

if σ_0 is complete and $\text{pc}_0, \sigma_0 \rightarrow_S^* \text{pc}, \sigma$, then σ is also complete. We can now formulate and prove the desired soundness statement.

Theorem 4. *Let S be the program $C_1 \parallel \dots \parallel C_n$ and suppose $\vdash \{P_k\} C_k \{Q_k\}$ for $k = 1..n$. Suppose further that either at least one of Q_k is intuitionistic or resource invariants for lock sorts used in the proofs are admissible. Then for any interpretation i and complete state σ_0 such that $\sigma_0 \in \langle \otimes_{k=1}^n \llbracket P_k \rrbracket_i^k \rangle$ the program S is safe when run from σ_0 and if $\text{pc}_0, \sigma_0 \rightarrow_S^* \text{pc}_f, \sigma$, then $\sigma \in \langle \otimes_{k=1}^n \llbracket Q_k \rrbracket_i^k \rangle$.*

Proof. Consider an interpretation i and a complete state $\sigma_0 \in \langle \otimes_{k=1}^n \llbracket P_k \rrbracket_i^k \rangle$. Therefore $\sigma_0 \in (\otimes_{k=1}^n \llbracket P_k \rrbracket_i^k) * (\otimes_{(A,u) \in \text{Free}(\sigma_0)} \mathcal{I}_0(A, u))$ from the definition of closure. Then by Theorem 1 the program S is safe when run from σ_0 and if $\text{pc}_0, \sigma_0 \rightarrow_S^* \text{pc}_f, \sigma$, then $\sigma \in (\otimes_{k=1}^n \llbracket Q_k \rrbracket_i^k) * (\otimes_{(A,u) \in \text{Free}(\sigma)} \mathcal{I}_0(A, u))$. Hence, by the definition of closure, we have $\sigma \in \sigma_1 * \sigma_2$ where $\sigma_1 \in \langle \otimes_{k=1}^n \llbracket Q_k \rrbracket_i^k \rangle$ and $\sigma_2 \in \otimes_{(A,u) \in L} \mathcal{I}_0(A, u)$ for some $L \subseteq \text{Free}(\sigma)$. If one of Q_k is intuitionistic, then from this it directly follows that $\sigma \in \langle \otimes_{k=1}^n \llbracket Q_k \rrbracket_i^k \rangle$.

Suppose now that $L \neq \emptyset$ and the resource invariants for lock sorts mentioned in L are admissible. Consider any $(A, u) \in L$. The state σ is complete, therefore, the permission associated with the lock at the address u in σ is 1. Besides, since $L \subseteq \text{Free}(\sigma)$, the value associated with u in σ is 0. Hence, if the permission associated with u in σ_2 were less than 1, then u would have to be allocated in σ_1 with a non-zero permission and the value U , which would contradict the definition of closure (a state in a closure cannot contain locks with the value U). So, for any $(A, u) \in L$ the permission associated with u in σ_1 is 1, which contradicts the admissibility of resource invariants for lock sorts used in the proof of the program. Therefore, $L = \emptyset$ and, hence, $\sigma \in \langle \otimes_{k=1}^n \llbracket Q_k \rrbracket_i^k \rangle$. \square

Note that for garbage-collected languages we can use the intuitionistic version of the logic [8] (i.e., one in which every assertion is intuitionistic) and, hence, do not have to check admissibility. Also, admissibility does not have to be checked if we are not interested in detecting memory leaks, as then Theorem 1 can be used.

8 Dynamic Thread Creation

We now extend the programming language with dynamically created threads:

$$\begin{array}{ll} T ::= f, f_1, f_2, \dots & \text{procedure names} \\ C ::= \dots \mid V = \mathbf{fork}(T) \mid \mathbf{join}(E) & \text{primitive commands} \\ P ::= \mathbf{let } T = S, \dots, T = S \mathbf{ in } S & \text{programs} \end{array}$$

We represent the code of threads by parameterless procedures (passing parameters to threads at the time of their creation is orthogonal to our concerns here and can be handled in a way similar to the one used for handling procedure calls when variables are treated as resources [12]; see [6] for details). A program consists of several procedure declarations along with the code of the main thread. We consider only well-formed programs in which all declared procedure names are distinct and all procedures used are declared. $x = \mathbf{fork}(f)$ creates a

new thread executing the code of the procedure f and stores the corresponding thread identifier into the variable x . $\mathbf{join}(E)$ waits until the thread with the identifier E finishes executing. In our semantics we allow at most one thread to wait for the termination of a given thread.

We add two new forms to our assertion language: $\Phi ::= \dots \mid \mathbf{tid}_f(E) \mid \mathbf{emp}_t$. A formula $\mathbf{tid}_f(E)$, which we call a thread handle, represents the knowledge that the thread with the identifier E exists and executes the code of the procedure f , and gives its owner a permission to join the thread. \mathbf{emp}_t denotes that the assertion does not contain any permissions of this kind. Note that a thread is deallocated (only) when it is joined.

Judgements are now of the form $\Gamma \vdash \{P\} C \{Q\}$ where Γ is a context consisting of a set of procedure specifications, each of the form $\{P\} f \{Q\}$. We consider only contexts in which there is at most one specification for each procedure. As procedures are parameterless, we restrict our attention here to contexts in which pre- and postconditions do not contain free logical variables. We add $\Gamma \vdash$ to all the triples in the rules from Figure 1 as well as in the standard rules of separation logic. In addition, we \wedge -conjoin \mathbf{emp}_t to every pre- and postcondition in the axioms for primitive commands (except for the postcondition of ACQUIRE and the precondition of RELEASE: in the postcondition of ACQUIRE and the precondition of RELEASE we \wedge -conjoin \mathbf{emp}_t right after $\mathbf{Locked}_A(L, \vec{X})$). To reason about \mathbf{fork} and \mathbf{join} we introduce two new axioms:

$$\frac{}{\Gamma, \{P\} f \{Q\} \vdash \{(x \Vdash \mathbf{emp}_h \wedge \mathbf{emp}_t) * P\} x = \mathbf{fork}(f) \{x \Vdash \mathbf{emp}_h \wedge \mathbf{tid}_f(x)\}} \text{FORK}$$

$$\frac{}{\Gamma, \{P\} f \{Q\} \vdash \{O \Vdash \mathbf{emp}_h \wedge \mathbf{tid}_f(E)\} \mathbf{join}(E) \{(O \Vdash \mathbf{emp}_h \wedge \mathbf{emp}_t) * Q\}} \text{JOIN}$$

That is, upon creating a new thread executing the code of procedure f , the thread that executed \mathbf{fork} obtains the thread handle $\mathbf{tid}_f(x)$ for the newly-created thread and gives up ownership of the precondition of f . Joining a thread with the identifier E requires the joining thread to own the handle $\mathbf{tid}_f(E)$. When \mathbf{join} succeeds, the thread exchanges the handle for the postcondition of f .

The model of the assertion language has to be adapted to account for thread handles. A state of the program is now represented by a triple of a stack, a heap, and a thread pool, the latter represented by a finite partial function from thread identifiers to procedure names. Now a lock can be free, held by the main thread, held by a thread, or its status may be unknown. Assertions are then interpreted with respect to a thread identifier, a stack, a heap, a thread pool, and an interpretation of logical variables. We define the $*$ operation on thread pools as disjoint union of the partial functions representing them, and the semantics of $P * Q$ and $P \multimap Q$ are then straightforwardly adjusted so as to partition thread pools. In addition, we add clauses for the new forms in our assertion language such that $\mathbf{tid}_f(E)$ describes singleton thread pools and \mathbf{emp}_t describes the empty thread pool. The satisfaction relation for all other formulae just ignores the thread pool. The notion of precision of formulae does not change.

Due to space constraints, we omit the detailed development of semantics and soundness for the extended logic, which can be found in [6]. Instead, we just

state the conditions under which the logic is sound. A proof of the program **let** $f_1 = C_1, \dots, f_n = C_n$ **in** C is given by triples $\Gamma \vdash \{P_1\} C_1 \{Q_1\}, \dots, \Gamma \vdash \{P_n\} C_n \{Q_n\}, \Gamma \vdash \{P\} C \{Q\}$, where $\Gamma = \{P_1\} f_1 \{Q_1\}, \dots, \{P_n\} f_n \{Q_n\}$. For the proof to be sound, P_k must be precise, and P_k and Q_k must have empty locksets, for all $k = 1..n$. We note that the operational semantics of Section 6 can be adjusted to our setting and a soundness statement similar to Theorem 1 can then be formulated. The proof of soundness is then done in the same style as that of Theorem 1. The notions of closure and admissibility can also be generalized to the new setting and a theorem similar to Theorem 4 can be proved.

9 Conclusions and Related Work

We have presented a logic that allows reasoning about concurrent heap-manipulating programs with realistic concurrency primitives including unbounded numbers of locks dynamically allocated and destroyed in the heap and threads dynamically created and terminating themselves. We have demonstrated that the logic makes it possible to reason locally about programs with a notion of dynamic ownership of heap parts by locks and threads. We believe that in the future this aspect of the logic will produce some additional pay-offs. First, the resource-oriented flavor of the logic should make it easy to design program analyses on the basis of it following the lines of [7]. In fact, the fixed-point equations defining the thread-local semantics used in the proof of soundness of our logic can be seen as a scheme of a thread-modular program analysis in the style of [7]. Second, lock handles in our logic’s assertion language are somewhat reminiscent of abstract predicates used for modular reasoning in separation logic about object-oriented programs [11]. This is not a coincidence as object-oriented programs use information hiding extensively in their locking mechanisms, and hence, often satisfy the Ownership Hypothesis. For this reason, we believe that our logic, combined with the techniques from [11], should be convenient for reasoning about concurrent object-oriented programs. Note, however, that lock handles and abstract predicates are different, in particular, we cannot see a way in which the former can be encoded in terms of the latter.

Two papers [4,15] have recently suggested combinations of separation logic and rely-guarantee reasoning that, among other things, can be used to reason about storable locks. For example, in [15] locks are not treated natively in the logic, but are represented as cells in memory storing the identifier of the thread that holds the lock; rely-guarantee is then used to simplify reasoning about the global shared heap with locks allocated in it. The logic allows modular reasoning about complex fine-grained concurrency algorithms (e.g., about the optimistic list mentioned in Section 4), but loses locality of reasoning for programs that allocate and deallocate many simple data structures protected by locks, which results in awkward proofs. In other words, as the original concurrent separation logic, the logics in [4,15] are designed for reasoning about the concurrent control of bounded numbers of data structures whereas our logic is designed to reason about the concurrent control of unboundedly many data structures that are

dynamically created and destroyed. Ideally, one wants to have a combination of both: a logic in which on the higher-level the reasoning is performed in a resource-oriented fashion and on the lower-level rely-guarantee is applied to deal with complex cases. Achieving this is another direction of our future research.

Feng and Shao [5] presented a rely-guarantee logic for reasoning about concurrent assembly code with dynamic thread creation. They do not have analogs of our rules for ownership transfer at **fork** and **join** commands. On a higher level, our logic for storable threads relates to theirs in the same way as separation logic relates to rely-guarantee reasoning: the former is good at describing ownership transfer, the latter at describing interference. As in the case of storable locks, investigating possible combinations of the two approaches would be fruitful.

Acknowledgments. We would like to thank Richard Bornat, Cristiano Calcagno, Peter O’Hearn, and Matthew Parkinson for comments and discussions that helped to improve the paper.

References

1. Bornat, R., Calcagno, C., O’Hearn, P.W., Parkinson, M.: Permission accounting in separation logic. In: POPL (2005)
2. Brookes, S.D.: A semantics of concurrent separation logic. Theoretical Computer Science 375(1-3), 227–270 (2007) Preliminary version appeared in CONCUR, 2004
3. Calcagno, C., O’Hearn, P., Yang, H.: Local action and abstract separation logic. In: LICS (2007)
4. Feng, X., Ferreira, R., Shao, Z.: On the relationship between concurrent separation logic and assume-guarantee reasoning. In: ESOP (2007)
5. Feng, X., Shao, Z.: Modular verification of concurrent assembly code with dynamic thread creation and termination. In: ICFP (2005)
6. Gotsman, A., Berdine, J., Cook, B., Rinetzky, N., Sagiv, M.: Local reasoning for storable locks and threads. Technical Report MSR-TR-2007-39, Microsoft Research (April 2007)
7. Gotsman, A., Berdine, J., Cook, B., Sagiv, M.: Thread-modular shape analysis. In: PLDI (2007)
8. Ishtiaq, S., O’Hearn, P.W.: BI as an assertion language for mutable data structures. In: POPL (2001)
9. O’Hearn, P., Reynolds, J., Yang, H.: Local reasoning about programs that alter data structures. In: CSL (2001)
10. O’Hearn, P.W.: Resources, concurrency and local reasoning. Theoretical Computer Science 375(1-3), 271–307 (2007) Preliminary version appeared in CONCUR 2004
11. Parkinson, M., Bierman, G.: Separation logic and abstraction. In: POPL (2005)
12. Parkinson, M., Bornat, R., Calcagno, C.: Variables as resource in Hoare logics. In: LICS (2006)
13. Reynolds, J.C.: Separation logic: A logic for shared mutable data structures. In: LICS (2002)
14. Vafeiadis, V., Herlihy, M., Hoare, T., Shapiro, M.: Proving correctness of highly-concurrent linearisable objects. In: PPOPP (2006)
15. Vafeiadis, V., Parkinson, M.J.: A marriage of rely/guarantee and separation logic. In: CONCUR (2007)