

Formal Reasoning about Concurrent Assembly Code with Reentrant Locks

Ming Fu^{1,2} Yu Zhang^{1,2} Yong Li^{1,2}

¹*Department of Computer Science & Technology
University of Science & Technology of China
Hefei, 230027, China
{fuming, liyong}@mail.ustc.edu.cn*

²*Software Security Laboratory
Suzhou Institute for Advanced Study, USTC
SuZhou, 215123, China
yuzhang@ustc.edu.cn*

Abstract—This paper focuses on the problem of reasoning about concurrent assembly code with reentrant locks. Our verification technique is based on concurrent separation logic (CSL). In CSL, locks are treated as non-reentrant locks and each lock is associated with a resource invariant, the lock-protected resources are obtained and released through acquiring and releasing the lock respectively. In order to accommodate for reentrancy, we introduce some additional notions into our specification language to describe reentrant level for each acquiring and releasing lock operation. Keeping track of the reentrant level for each lock in the pre- and post-conditions enables the program logic to ensure that resources are not re-acquired upon reentrancy, thus resources owned by a thread are prevented from reintroducing in the postcondition. Our framework is fully mechanized. Its soundness has been verified using the Coq proof assistant. We demonstrate the usage of our framework through giving a safety proof of a simple program.

Keywords-reentrant locks, concurrent separation logic, safety, program logic

I. INTRODUCTION

It is difficult to write correct concurrent programs due to potential inter-thread interference at every program point. In order to reduce the complexity of concurrent programming, most popular modern languages – Java and C# provide high-level reentrant locking primitives, which ease concurrent programming. However, it is difficult to use reentrant locks correctly and the incorrect usage can result in nasty concurrent errors like data races or deadlocks. Existing high-level languages do not provide any effective mechanisms to avoid such errors, thus it is important to develop a verification technique for reasoning about concurrent programs with reentrant locks. The reentrant mechanism allows a thread to re-acquire a lock that it already holds. It is important because it eliminates the possibility of a single thread deadlocking itself on a lock that it already holds.

Concurrent separation logic(CSL) [1] is an extension of *separation logic* [2] for reasoning about shared memory race-free concurrent programs. Separation logic is a program logic which is tailored to reason about the heap manipulating programs. In CSL, the shared memory is partitioned and each part is protected by a unique mutual exclusive lock. For each part of the partition, an invariant is assigned to specify its well-formedness. When a thread acquires one

of the mutual exclusive locks, it treats the part of shared memory protected by the lock as private. Before releasing the lock, it must ensure that the part of shared memory is well-formed with regard to the corresponding invariant. The ownership of lock-protected shared memory can be dynamically transferred among threads, the verification system ensures that a piece of shared memory is only accessed when the associated lock is held. However, in the invariants of CSL, locks are non-reentrant, we cannot directly apply CSL to reason about concurrent programs with reentrant locks.

In order to adapt CSL to reasoning about concurrent programs with reentrant locks, we build an abstract machine model based on an assembly language with RISC-style instructions and built-in "lock *l*" and "unlock *l*" primitives, and introduce additional specification constructs to trace the reentrant level for each lock. Instead of using the high-level parallel language proposed by Hoare [3], we use the assembly language because it has cleaner semantics, which makes our formulation much simpler. For instance, we do not use variables, instead we only use register files and memory. Therefore we can have a quick formulation in Coq [4] without worrying about variable renaming issues. Also we do not have to formalize the complicated syntactic constraints enforced in CSL over shared variables. Another important reason is that our work at low level can be easily applied to generate proof-carrying code (PCC) [5]. It seems unavoidable that the proof rule for acquiring a lock distinguishes between initial acquires and re-acquires. This is needed because it is quite obviously unsound to simply assume the resource invariant after a re-acquire. Thus, a verification system for reentrant locks must keep track of the reentrant level for each lock that the current thread holds in the pre- and post conditions, and we have to enrich our specification language to achieve this requirement. Our system addresses the safety issues at assembly-level as PCC systems do. So we do not need to trust the complicated compilation and optimization and can have a smaller trusted computing base to build executable PCC package for programs using reentrant locks. Furthermore, our formal model for reentrant locks is still general and similar to high-level ones. The verification technique we describe at assembly-level can be lifted up to higher levels. This paper makes the following contributions:

- 1) As far as we know, this paper first proposes a method to adapt CSL to fitting for reasoning about concurrent assembly code with reentrant locks. We present a program logic for reasoning about properties of concurrent assembly code with reentrant locks, and we prove it sound with respect to the semantics of reentrant locks.
- 2) We implement our framework using the Coq proof assistant, and prove an example under the framework. The result shows that the adapted inference rules can be easily applied to verify the concurrent assembly code with reentrant locks.

The rest of this paper is organized as follows: In section II, we explain the CSL and its limitation for verifying reentrant locks. We describe the abstract machine we model and the program logic for reasoning about concurrent assembly code with reentrant locks in section III. Section IV presents an example that are written and proved under our framework. Finally we discuss the related work and conclusion in section V and VI.

II. PRELIMINARIES

We give a brief description of separation logic. A more careful treatment is in [1] and [6]. A simplified syntax for separation logic is shown in Fig. 1.

Here we briefly demonstrate the logical semantics for each construct in the syntax. Both A and B are assertions that describe the heap. $1 \mapsto v$ holds if the heap consists entirely of the binding of location 1 to value v . emp holds only on the empty heap. $A * B$ holds if the heap can be split into two disjoint parts such that A holds on one and B on the other. $A \wedge B$ holds if both A and B hold on the entire heap. $A \vee B$ holds if either A or B holds on the heap. $\exists x. B$ holds if there exists an x that B holds on the heap. $\forall x. A$ holds on a heap that satisfies A for all x .

$$A, B ::= 1 \mapsto v \mid \text{emp} \mid A * B \mid A \wedge B \mid A \vee B \\ \mid \exists x. B \mid \forall x. A$$

Figure 1. Syntax of Separation Logic

The *frame property* of separation logic requires that if a program does not go wrong in a particular state with heap \mathbb{H} , then it will not go wrong in a larger state with heap $\mathbb{H} \uplus \mathbb{H}'$ (the notion " \uplus " is used to merge two disjoint heaps into a larger one, we give its formal definition in Fig. 6); the effect will still be taken on \mathbb{H} , leaving the added heap \mathbb{H}' completely unaffected. Thus the separation conforms to the following frame rule:

$$\frac{\{Q\}C\{R\}}{\{P * Q\}C\{P * R\}}$$

(no variable occurring free in P is modified by C)
If C cannot modify the variables of P , and if the heap it

manipulates is disjoint from that of P , then we can reason about C and its effect separately from P .

CSL introduces the concurrency rule based on separation logic for reasoning about concurrent programs. The concurrency rule given below

$$\frac{\{Q_1\}C_1\{R_1\} \quad \{Q_2\}C_2\{R_2\}}{\{Q_1 * Q_2\}C_1 \parallel C_2\{R_1 * R_2\}}$$

describes how concurrent threads with disjoint heap resources can be treated separately. As a concurrent program executes, heap resources must remain separated but the separation need not be fixed : the ownership can be transferred among threads through locking operation. The rule below is used to deal with the non-reentrant locks for transferring the ownership of shared resources.

$$\frac{I \text{ is a resource invariant associated with a lock } l}{\Gamma, l \rightsquigarrow I \vdash \{\text{emp}\} \text{lock } l\{I\}}$$

However, we can not directly apply this rule to reason about concurrent programs with reentrant locks. The main problem is that a verification system for reentrant locks has to distinguish between initial lock entries and reentries, because only after initial entries is it sound to assume a lock's resource invariant. This means that initial lock entries need precondition requiring that the current thread has not yet held the acquired lock. In Fig. 2, a simple code sequence is made up of two consecutive statements that acquire the same lock l . Both the first and the second acquiring lock l operations lead to obtain additional resource satisfying invariant I . According to the frame rule and the above rule, the second acquiring lock operation requires that the postcondition be $I * I$. However separation logic treats $I * I$ as a false assertion, and this leads to incorrect verification. The following sections show our technique for solving this problem and adapting CSL to reasoning about concurrent programs with reentrant locks.

$$\begin{array}{l} \{\text{emp}\} \\ (1) \text{lock } l ; \\ \quad \{I\} \quad (I \text{ is lock } l \text{ 's resource invariant}) \\ (2) \text{lock } l ; \\ \quad \{I * I\} \quad (\text{Wrong!!!}) \\ \dots \end{array}$$

Figure 2. CSL does not Support Reentrant Locks

III. THE FRAMEWORK

A. Abstract Machine

Fig. 3 defines the abstract machine and the syntax of an assembly language. We extend CAP [7], [8] by adding built-in primitives "lock l " and "unlock l " for reentrant locks. The whole world \mathbb{W} consists of a code heap \mathbb{C} , a shared data heap

(World)	W	$::= (\mathbb{C}, \mathbb{H}, \mathbb{T}\mathbb{S}, \mathbb{L})$
(ThreadSet)	$\mathbb{T}\mathbb{S}$	$::= (\mathbb{T}_1, \dots, \mathbb{T}_n)$
(ThreadState)	\mathbb{S}	$::= (\mathbb{H}, \mathbb{T}, \mathbb{L})$
(Thread)	\mathbb{T}_i	$::= (\mathbb{R}, \text{pc}, \text{tid})$
(ThrID)	tid	$::= m \text{ (nat nums, and } m > 0)$
(CodeHeap)	\mathbb{C}	$::= (\mathbb{f} \rightsquigarrow \mathbb{t})^*$
(Heap)	\mathbb{H}	$::= \{\mathbb{l} \rightsquigarrow \mathbb{w}\}^*$
(LockMap)	\mathbb{L}	$::= \{\mathbb{l} \rightsquigarrow (\text{tid}, \mathbb{n})\}^*$
(ReentrantLevel)	\mathbb{n}	$::= i \text{ (nat nums, and } i > 0)$
(RegFile)	\mathbb{R}	$::= \{\mathbb{r} \rightsquigarrow \mathbb{w}\}^*$
(Register)	\mathbb{r}	$::= \mathbb{r}_0 \mid \dots \mid \mathbb{r}_{31}$
(Labels)	$\mathbb{f}, \mathbb{l}, \text{pc}$	$::= i \text{ (nat nums)}$
(Locks)	\mathbb{l}	$::= i \text{ (nat nums)}$
(Word)	\mathbb{w}	$::= i \text{ (nat nums)}$
(Instr)	\mathbb{t}	$::= \text{add } \mathbb{r}_d, \mathbb{r}_s, \mathbb{r}_t \mid \text{addi } \mathbb{r}_d, \mathbb{r}_s, \mathbb{w}$ $\mid \text{sub } \mathbb{r}_d, \mathbb{r}_s, \mathbb{r}_t \mid \text{ld } \mathbb{r}_d, \mathbb{w}(\mathbb{r}_s)$ $\mid \text{st } \mathbb{r}_d, \mathbb{w}(\mathbb{r}_s) \mid \text{beq } \mathbb{r}_s, \mathbb{r}_t, \mathbb{f}$ $\mid \text{lock } \mathbb{l} \mid \text{unlock } \mathbb{l}$
(InstrSeq)	\mathbb{I}	$::= \mathbb{t}; \mathbb{I} \mid \mathbb{j} \mathbb{f} \mid \mathbb{j}\mathbb{r} \mathbb{r}_s$

Figure 3. The Abstract Machine

\mathbb{H} , a thread set $\mathbb{T}\mathbb{S}$ which contains n threads $(\mathbb{T}_1, \dots, \mathbb{T}_n)$ and a shared lock mapping \mathbb{L} .

The code heap \mathbb{C} is a partial mapping from code labels to instructions. The global shared heap \mathbb{H} is modeled as a finite partial mapping from heap locations \mathbb{l} (natural numbers) to word values \mathbb{w} (natural numbers). The locking map \mathbb{L} is a finite mapping from reentrant locks to lock value pairs (tid, \mathbb{n}) , where the integer tid identifies the thread holding the lock exclusively and the integer \mathbb{n} is the reentrant level counting how often it currently holds the lock.

The abstract machine has a fixed number of threads. Each thread \mathbb{T}_i contains a register file \mathbb{R} , a program counter pc and its thread identifier tid . Here we allow each thread to have its own register file and program counter, which is consistent with most implementation of thread library where the register file is saved to the execution context when a thread is preempted. The register file \mathbb{R} is represented as a total function from registers to words. Each thread's program counter pc points to its current command in a shared code heap \mathbb{C} . The set of instructions we present here are the commonly used subset in RISC machines with additional reentrant "lock \mathbb{l} " and "unlock \mathbb{l} " primitives for synchronization.

We define the instruction sequence \mathbb{I} as a sequence of sequential instructions ending with jump or return instructions. $\mathbb{C}[\text{pc}]$ extracts an instruction sequence starting from pc in \mathbb{C} , as defined in Fig. 4. $(F\{a \rightsquigarrow b\})(x)$ is used to formalize memory update in our operational semantics. Macros $\mathbb{S} \mid_{\mathbb{H}'}$ and $\mathbb{S} \mid_{\mathbb{L}'}$ are defined for constructing thread states by replacing the heap and the lock set respectively.

$$\mathbb{C}[\text{pc}] \stackrel{\text{def}}{=} \begin{cases} \mathbb{t} & \mathbb{t} = \mathbb{C}(\text{pc}) \text{ and } \mathbb{t} = \mathbb{j} \mathbb{f} \text{ or } \mathbb{j}\mathbb{r} \mathbb{r}_s \\ \mathbb{t}; \mathbb{I} & \mathbb{t} = \mathbb{C}(\text{pc}) \text{ and } \mathbb{I} = \mathbb{C}[\text{pc}+1] \end{cases}$$

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

$$\mathbb{S} \mid_{\mathbb{H}'} \stackrel{\text{def}}{=} (\mathbb{H}', \mathbb{S}, \mathbb{T}, \mathbb{S}, \mathbb{L})$$

$$\mathbb{S} \mid_{\mathbb{L}'} \stackrel{\text{def}}{=} (\mathbb{S}, \mathbb{H}, \mathbb{S}, \mathbb{T}, \mathbb{L}')$$

Figure 4. Definition of Representations

B. Operational Semantics

The operational semantics of each instruction is defined in Fig. 5. The relation NextS_t shows the transition of thread states by executing instruction \mathbb{t} . The operational semantics for most instructions are quite straightforward. Note the execution of instruction for acquiring locks. It allows a lock to be re-acquired by one thread locks and does not lead to deadlock. There exist three different cases for executing lock \mathbb{l} : when \mathbb{l} is not in the domain of \mathbb{L} (means that \mathbb{l} is free), the current thread tid can exclusively and successfully acquire the lock \mathbb{l} , and set lock \mathbb{l} with pair value $(\text{tid}, 1)$. tid denotes that lock \mathbb{l} is held by thread tid and the reentrant level "1" shows that the current program point is at the initial lock entry of the lock \mathbb{l} . Shared resource can be obtained by the thread through the initial lock acquiring. When \mathbb{l} is in the domain of \mathbb{L} and held by the current thread tid , thread tid tries to re-acquired a held lock. Non-reentrant locking mechanism makes the current thread block and leads to deadlock, while our model avoids deadlock through setting the reentrant level in the lock pair value with the increment of "1". When \mathbb{l} is held by the other thread, the current thread blocks. The semantics for releasing locks is straightforward, the reentrant level makes acquiring and releasing operation on the same lock keep in pair.

Fig. 5 also defines $(\mathbb{C}, \mathbb{S}) \rightsquigarrow (\mathbb{C}, \mathbb{S}')$ and $(\mathbb{W} \mapsto \mathbb{W}')$ for the thread execution and the whole world execution respectively. Note that relation $(\mathbb{C}, \mathbb{S}) \rightsquigarrow (\mathbb{C}, \mathbb{S}')$ is deterministic but our semantics of the abstract machine $(\mathbb{W} \mapsto \mathbb{W}')$ is not deterministic: the state transition may be made by executing any thread in \mathbb{W} . Also, given a \mathbb{W} , there may not always exist a \mathbb{W}' such that $(\mathbb{W} \mapsto \mathbb{W}')$ holds. If there is no such \mathbb{W}' , we say the program gets stuck at \mathbb{W} . One important goal of our program logic is to show that verified programs never get stuck.

C. Program Logic

1) *Assertion Language*: Fig. 6 shows the syntax and semantics of the assertion language. We use the predicate \mathbb{m} over a heap and separation logic connectors $*$ in our assertion language. The assertion \mathbb{a} is a predicate over a thread state.

Most of the definitions are simple and straightforward. Here we explain some special ones. The assertion $"\mathbb{l} \mapsto \mathbb{v}"$

NextS _t S S' where S = (H, (R, pc, tid), L)	
if $\iota =$	S' =
add r_d, r_s, r_t	(H, (R{ $r_d \rightsquigarrow R(r_s) + R(r_t)$ }, pc+1, tid), L)
addi r_d, r_s, w	(H, (R{ $r_d \rightsquigarrow R(r_s) + w$ }, pc+1, tid), L)
sub r_d, r_s, r_t	(H, (R{ $r_d \rightsquigarrow R(r_s) - R(r_t)$ }, pc+1, tid), L)
ld $r_d, w(r_s)$	(H, (R{ $r_d \rightsquigarrow H(R(r_s) + w)$ }, pc+1, tid), L) if $R(r_s) + w \in \text{dom}(H)$
st $r_d, w(r_s)$	(H{ $R(r_s) + w \rightsquigarrow R(r_d)$ }, (R, pc+1, tid), L) if $R(r_s) + w \in \text{dom}(H)$
lock l	(H, (R, pc+1, tid), L{ $l \rightsquigarrow (\text{tid}, 1)$ }) if $l \notin \text{dom}(L)$
	(H, (R, pc+1, tid), L{ $l \rightsquigarrow (\text{tid}, n+1)$ }) if $L(l) = (\text{tid}, n)$
	(H, (R, pc, tid), L) otherwise
unlock l	(H, (R, pc+1, tid), L{ $l \rightsquigarrow (\text{tid}, n-1)$ }) if $L(l) = (\text{tid}, n) \wedge n > 1$
	(H, (R, pc+1, tid), L/{ l }) if $L(l) = (\text{tid}, n) \wedge n = 1$
j f	(H, (R, f, tid), L)
jr r_s	(H, (R, R(r_s), tid), L)
beq r_s, r_t, f	(H, (R, f, tid), L) if $R(r_s) = R(r_t)$
	(H, (R, pc+1, tid), L) if $R(r_s) \neq R(r_t)$

$$\frac{\iota = C(\text{pc}) \quad \text{NextS}_t \text{ S S'}}{(C, S) \rightsquigarrow (C, S')} \text{ THREADSTEP}$$

$$\frac{\exists T_k. T_k \in \text{TS} \wedge (C, (H, T_k, L)) \rightsquigarrow (C, (H', T'_k, L'))}{(C, H, \text{TS}, L) \mapsto (C, H', (T_1 \dots, T'_k, \dots, T_n), L')} \text{ WORLDSTEP}$$

Figure 5. Operational Semantics

(ThrdStatePred) $a \in \text{ThreadState} \rightarrow \text{Prop}$
(HeapPred) $m \in \text{Heap} \rightarrow \text{Prop}$

$m ::= 1 \mapsto v \mid \text{true} \mid \text{emp} \mid m_1 * m_2$
 $\mid m_1 \wedge m_2 \mid m_1 \vee m_2 \mid \exists x. m \mid \forall x. m$
 $a ::= [m] \mid \text{own}_k(l, n) \mid r = v$
 $\mid a_1 \wedge a_2 \mid a_1 \vee a_2 \mid \exists x. a \mid \forall x. a$

$\text{true} \stackrel{\text{def}}{=} \lambda H. \text{True}$
 $\text{emp} \stackrel{\text{def}}{=} \lambda H. \text{dom}(H) = \emptyset$
 $H_1 \perp H_2 \stackrel{\text{def}}{=} \text{dom}(H_1) \cap \text{dom}(H_2) = \emptyset$
 $1 \mapsto v \stackrel{\text{def}}{=} \lambda H. H = \{1 \rightsquigarrow v\}$
 $H_1 \uplus H_2 \stackrel{\text{def}}{=} \begin{cases} H_1 \cup H_2 & \text{if } H_1 \perp H_2 \\ \text{undefined} & \text{otherwise} \end{cases}$
 $m_1 * m_2 \stackrel{\text{def}}{=} \lambda H. \exists H_1, H_2. (H_1 \uplus H_2 = H) \wedge m_1 H_1 \wedge m_2 H_2$
 $[m] \stackrel{\text{def}}{=} \lambda S. m \text{ S.H}$
 $\text{own}_k(l, n) \stackrel{\text{def}}{=} \lambda S. (k = \text{S.T.tid}) \wedge \text{S.L}(l) = (k, n)$
 $r = v \stackrel{\text{def}}{=} \lambda S. \text{S.T.R}(r) = v$

Figure 6. Syntax and Semantics of the Assertion Language

into two disjoint parts, and m_1 and m_2 hold over each of them respectively. $[m]$ means predicate over a thread state containing a heap satisfying m , we use this syntax to lift predicates over heap to assertions specifying a thread state. Predicate $\text{own}_k(l, n)$ is used to specify that l is held by the thread k with corresponding reentrant level n . Here, we omit the semantics of some straightforward connectors, such as $\wedge, \vee, \text{etc.}$

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

(WorldSpec) $\phi := ([\Psi_1, \dots, \Psi_n], \Gamma)$
(CdHpspec) $\Psi := \{(f_1, a_1), \dots, (f_n, a_n)\}$
(LockINV) $\Gamma := \{l \rightsquigarrow m\}^*$
(Well-formed World) $\phi, [a_1, \dots, a_n] \vdash (C, H, \text{TS}, L)$
(Well-formed Thread) $\Psi, \Gamma \vdash \{a\} (C, H, T, L)$
(Well-formed Code Heap) $\Psi, \Gamma \vdash C : \Psi'$
(Well-formed Instr. Seq.) $\Psi, \Gamma \vdash \{a\} \text{pc} : \iota; \text{I}$
(Well-formed Instruction) $\Psi, \Gamma \vdash \{a\} \text{pc} : \iota$

Figure 7. Specification Constructs for the Program Logic

holds only if the heap has only one cell at location 1 containing value v . $m_1 * m_2$ means the heap can be split

The specification constructs of our logic are presented in Fig. 7. The world specification ϕ contains a collection of code heap specifications for each thread and a specification

Γ for lock-protected heap. Code heap specification ψ maps a code label to a predicate a over thread state \mathbb{S} as the precondition of corresponding instruction sequence. The specification Γ of a lock-protected heap maps a lock to an invariant m specifying the shared heap.

The last five judgments are used to define the well-formed world, well-formed thread, well-formed code heap, well-formed instruction sequence and well-formed instruction respectively. The inference rules for these judgments will be presented in the following subsection.

$$\begin{aligned} a_1 \Rightarrow a_2 &\stackrel{\text{def}}{=} \lambda \mathbb{S}. a_1 \mathbb{S} \rightarrow a_2 \mathbb{S} \\ a * m &\stackrel{\text{def}}{=} \lambda \mathbb{S}. \exists \mathbb{H}_1, \mathbb{H}_2. (\mathbb{H}_1 \uplus \mathbb{H}_2 = \mathbb{S}. \mathbb{H}) \wedge a \mathbb{S}_{|\mathbb{H}_1} \wedge m \mathbb{H}_2 \\ \psi \circ \text{NextS}_t &\stackrel{\text{def}}{=} \lambda \mathbb{S}. \exists \mathbb{S}', \text{NextS}_t \mathbb{S} \mathbb{S}' \wedge \psi(\mathbb{S}'. \text{T.pc}) \mathbb{S}' \\ \forall_* x \in S. P(x) &\stackrel{\text{def}}{=} \begin{cases} \text{emp} & \text{if } S = \emptyset \\ P(x_i) * \forall_* x \in S'. P(x) & \text{if } S = S' \uplus \{x_i\} \end{cases} \end{aligned}$$

Figure 8. Auxiliary Definition

3) *Inference Rules:* The inference rules for a program and instructions are presented in Fig. 9.

A world is well-formed with regard to a world specification ϕ and thread state predicates a_1, \dots, a_n for each thread when the following conditions hold:

- There is a partition of the global heap into $n + 1$ disjoint parts, where the shared heap \mathbb{H}_s satisfies the invariants specified in Γ and $\mathbb{H}_1, \dots, \mathbb{H}_n$ satisfy each thread state predicate a_k respectively. As in O’Hearn’s original work on CSL [1], we also require invariants specified in Γ to be precise, denoted as $\text{Precise}(\Gamma)$ defined as below. Every thread of the world is required to be well-formed. Thus our system support thread-modular verification by decomposing the verification of multi-threaded program into that of its component threads.

$$\begin{aligned} \text{Precise}(m) &\stackrel{\text{def}}{=} \forall \mathbb{H}_1, \mathbb{H}_2, \mathbb{H}. \mathbb{H}_1 \subseteq \mathbb{H} \rightarrow \mathbb{H}_2 \subseteq \mathbb{H} \rightarrow \\ &\quad m \mathbb{H}_1 \wedge m \mathbb{H}_2 \rightarrow \mathbb{H}_1 = \mathbb{H}_2 \\ \text{Precise}(\Gamma) &\stackrel{\text{def}}{=} \forall l \in \text{dom}(\Gamma). \text{Precise}(\Gamma(l)) \end{aligned}$$

- The shared state $(\mathbb{H}_s, \underline{\cdot}, \mathbb{L})$ satisfies the predicate a_Γ , which is defined below. The definition of a_Γ is the separating conjunction of invariants assigned to the locks which are free (not in the domain of the global \mathbb{L}). It ensures that the shared heap are well-formed outside critical region. Here \forall_* is an indexed, finitely iterated separating conjunction, which is formalized in Fig. 8.

$$a_\Gamma \stackrel{\text{def}}{=} \lambda \mathbb{S}. (\forall_* l \in \{l \mid l \notin \text{dom}(\mathbb{S}. \mathbb{L})\}. \Gamma(l)) \mathbb{S}. \mathbb{H}$$

A thread is well-formed if the current thread state satisfies the precondition a and both the code heap and the instruction sequence are required to be well-formed. Since a only specifies the private resource, we use “filter” operator “ $\mathbb{L}|_{\text{tid}}$ ” formalized below to prevent a from having access

to the ownership information of locks not held by the current thread.

$$(\mathbb{L}|_{\text{tid}})(l) \stackrel{\text{def}}{=} \begin{cases} (\text{tid}, n) & \text{if } \mathbb{L}(l) = (\text{tid}, n) \\ \text{undefined} & \text{otherwise} \end{cases}$$

The rule CDHP shows that a code heap is well-formed only if each instruction sequence specified in ψ' is well-formed with respect to the imported interfaces specified with ψ and the lock specification Γ .

The rule INSQ shows that an instruction sequence is well-formed if it is composed of a single instruction ι and another instruction sequence \mathbb{I} , both of which are well-formed.

A well-formed instructions includes the following cases with the order of Fig. 9.

- The rule INSN - If the instruction ι is not lock l or unlock l , it can execute for all thread states specified by the current thread state predicate a , and the new modified thread state must satisfy the thread state predicate for the target address of instruction ι given by ψ . We use $\psi \circ \text{NextS}_t$ defined in Fig. 8 to specify the new modified state generated by executing instruction ι .
- The rule LOCK - We have a unified rule for reasoning about instruction lock l which may be executed at either the initial entry or reentry. The reentrant locks are handy in the presence of polymorphism, i.e. where a given routine that executes lock is called both in a context where the lock is free and where the lock was previously acquired. In that sense, whether the locking operation happens at the initial entry or reentry cannot be established statically, and the unified rule LOCK can support reasoning about the either case automatically. The rule applies when lock l is acquired at the initial entry or reentry. The thread predicate $((\text{En}_{\text{lk}}^l \wedge m) \vee (\text{En}_{\text{re lk}}^l \wedge \text{emp}))$ is used to enforces the ownership transfer under the following two cases:
 - If the current state satisfies the predicate En_{lk}^l (defined in Fig. 10) which ensures the lock is free and enables safely locking operation at the initial entry, we can carry the knowledge m in the postcondition given by ψ at the target address of instruction ι . The global invariant ensures that the part of heap protected by l satisfies the invariant m .
 - If the current state satisfies the predicate $\text{En}_{\text{re lk}}^l$ (defined in Fig. 10) which ensures the lock is held by itself and enables safely locking operation at the reentry, we can use the empty heap predicate emp to represent nothing is acquired at the reentry. The part of heap protected by l will not be reintroduced into the result state.
- The rule UNLOCK is similar with the rule LOCK , we use the predicate $((\text{En}_{\text{unlk}}^l \wedge m) \vee (\text{En}_{\text{re unlk}}^l \wedge \text{emp}))$ to represent two different cases, one is that the invariant gets established and the lock l ’s current reentrancy level

$$\boxed{\phi, [a_1, \dots, a_n] \vdash (\mathbb{C}, \mathbb{H}, \mathbb{T}\mathbb{S}, \mathbb{L})} \quad \text{(Well-formed World)}$$

$$\frac{\phi = ([\Psi_1, \dots, \Psi_n], \Gamma) \quad \mathbb{H} = \mathbb{H}_s \uplus \mathbb{H}_1 \uplus \dots \uplus \mathbb{H}_n \quad \text{Precise}(\Gamma) \quad \text{a}_\Gamma(\mathbb{H}_s, \underline{\quad}, \mathbb{L}) \quad \Psi_k, \Gamma \vdash \{a_k\}(\mathbb{C}, \mathbb{H}_k, \mathbb{T}_k, \mathbb{L}) \quad \text{for all } k \in \{1, \dots, n\}}{\phi. [a_1, \dots, a_n] \vdash (\mathbb{C}, \mathbb{H}, [\mathbb{T}_1, \dots, \mathbb{T}_n], \mathbb{L})} \quad \text{WORLD}$$

$$\boxed{\Psi, \Gamma \vdash \{a\}(\mathbb{C}, \mathbb{H}, \mathbb{T}, \mathbb{L})} \quad \text{(Well-formed Thread)}$$

$$\frac{\text{a}(\mathbb{H}, (\mathbb{R}, \text{pc}, \text{tid}), \mathbb{L}|_{\text{tid}}) \quad \Psi, \Gamma \vdash \mathbb{C} : \Psi \quad \Psi, \Gamma \vdash \{a\} \text{pc} : \mathbb{C}[\text{pc}]}{\Psi, \Gamma \vdash \{a\}(\mathbb{C}, \mathbb{H}, (\mathbb{R}, \text{pc}, \text{tid}), \mathbb{L})} \quad \text{THRD}$$

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

$$\frac{\forall (\text{pc}, a) \in \Psi' : \Psi, \Gamma \vdash \{a\} \text{pc} : \mathbb{C}[\text{pc}]}{\Psi, \Gamma \vdash \mathbb{C} : \Psi'} \quad \text{CDHP}$$

$$\boxed{\Psi, \Gamma \vdash \{a\} \text{pc} : \mathbb{I}} \quad \text{(Well-formed Instr. Sequence)}$$

$$\frac{\Psi, \Gamma \vdash \{a'\} \text{pc} + 1 : \mathbb{I} \quad \Psi\{\text{pc} + 1 \rightsquigarrow a'\}, \Gamma \vdash \{a\} \text{pc} : \mathfrak{t}}{\Psi, \Gamma \vdash \{a\} \text{pc} : \mathfrak{t}; \mathbb{I}} \quad \text{INSQ}$$

$$\boxed{\Psi, \Gamma \vdash \{a\} \text{pc} : \mathfrak{t}} \quad \text{(Well-formed Instruction)}$$

$$\frac{\mathfrak{t} \notin \{\text{lock } l, \text{unlock } l\} \quad a \Rightarrow \Psi \circ \text{NextS}_{\mathfrak{t}}}{\Psi, \Gamma \vdash \{a\} \text{pc} : \mathfrak{t}} \quad \text{INSN}$$

$$\frac{\text{a} * ((\text{En}_{\text{lk}}^l \wedge \text{m}) \vee (\text{En}_{\text{relk}}^l \wedge \text{emp})) \Rightarrow \Psi \circ \text{NextS}_{\text{lock } l}}{\Psi, \Gamma\{l \rightsquigarrow \text{m}\} \vdash \{a\} \text{pc} : \text{lock } l} \quad \text{LOCK}$$

$$\frac{\text{a} \Rightarrow (\Psi \circ \text{NextS}_{\text{unlock } l}) * ((\text{En}_{\text{unlk}}^l \wedge \text{m}) \vee (\text{En}_{\text{reunlk}}^l \wedge \text{emp}))}{\Psi, \Gamma\{l \rightsquigarrow \text{m}\} \vdash \{a\} \text{pc} : \text{unlock } l} \quad \text{UNLOCK}$$

Figure 9. Inference Rules

is 1; the other is that the specified heap is empty and the lock l 's current reentrancy level is bigger than 1. The predicate enforces that the ownership of the well-formed shared heap protected by the lock l only be transferred from private part to the shared part at the last releasing and the middle unlocking operations do not change the domain of thread private heap. The predicate $\text{En}_{\text{unlk}}^l$ and $\text{En}_{\text{reunlk}}^l$ defined in Fig. 10 enables safely unlocking action taken on the current state.

$$\begin{aligned}
\text{En}_{\text{lk}}^l &\stackrel{\text{def}}{=} \lambda \mathbb{S}. l \notin \text{dom}(\mathbb{S}. \mathbb{L}) \\
\text{En}_{\text{relk}}^l &\stackrel{\text{def}}{=} \lambda \mathbb{S}. (\mathbb{S}. \mathbb{T}. \text{tid}, \underline{\quad}) = \mathbb{S}. \mathbb{L}(l) \\
\text{En}_{\text{unlk}}^l &\stackrel{\text{def}}{=} \lambda \mathbb{S}. (\mathbb{S}. \mathbb{T}. \text{tid}, 1) = \mathbb{S}. \mathbb{L}(l) \\
\text{En}_{\text{reunlk}}^l &\stackrel{\text{def}}{=} \lambda \mathbb{S}. \exists n. (\mathbb{S}. \mathbb{T}. \text{tid}, n) = \mathbb{S}. \mathbb{L}(l) \wedge n > 1
\end{aligned}$$

Figure 10. Predicates for Enabling Instructions

4) *Soundness*: The soundness of these inference rules with respect to the operational semantics of the abstract

machine is proved following the syntactic approach [9]. From the "progress" and "preservation" lemmas, we can guarantee that given a well-formed program under the compatible preconditions, the current instruction sequence will be able to execute without getting "stuck". The soundness of our framework is formally stated as Theorem III.3.

Lemma III.1 (Progress) *For any world $\mathbb{W} = (\mathbb{C}, \mathbb{H}, (\mathbb{T}_1, \dots, \mathbb{T}_n), \mathbb{L})$, and if $\Psi, [a_1, \dots, a_n] \vdash \mathbb{W}$, then for any thread \mathbb{T}_k , there exist $\mathbb{H}'_k, \mathbb{T}'_k$ and \mathbb{L}' such that $(\mathbb{C}, (\mathbb{H}, \mathbb{T}_k, \mathbb{L})) \rightsquigarrow (\mathbb{C}, (\mathbb{H}', \mathbb{T}'_k, \mathbb{L}'))$.*

Lemma III.2 (Preservation) $\phi = ([\Psi_1, \dots, \Psi_n], \Gamma)$. *If $\phi, [a_1, \dots, a_n] \vdash \mathbb{W}$ and $\mathbb{W} \mapsto \mathbb{W}'$, then there exist a'_1, \dots, a'_n such that $\phi, [a'_1, \dots, a'_n] \vdash \mathbb{W}'$.*

Theorem III.3 (Soundness) $\phi = ([\Psi_1, \dots, \Psi_n], \Gamma)$. *If there exist a_1, \dots, a_n , such that $\phi, [a_1, \dots, a_n] \vdash \mathbb{W}$, then for any $n \geq 0$, there exist a world \mathbb{W}' and a'_1, \dots, a'_n such that $\mathbb{W} \mapsto^n \mathbb{W}'$ and $\phi, [a'_1, \dots, a'_n] \vdash \mathbb{W}'$.*

We have mechanized the complete soundness proof in the Coq proof assistant. Interested readers can check out our Coq implementation [10] for more detail.

IV. EXAMPLE

In this section, we give an example to demonstrate the mechanized verification of safety properties (usually the shared memory invariant in parallel program) for concurrent assembly code with reentrant locks.

A simple example is present in Fig. 11, which is the concurrent code that computes the next even number according to the current value stored in the shared memory location x . The shared location x is protected by a reentrant lock l and the value stored in it is initialized with 0. In high level code, we unfold the inlined synchronized method located from line 3 to line 5. The inlined method leads to acquiring the same lock which has been held by the caller. The lock l is a reentrant lock, so this code will run correctly without deadlock.

The corresponding assembly code and assertions are given in Fig. 12. We verify the code in our framework. Following MIPS convention, we assume r_0 always contains 0. Assertions are shown as annotations enclosed in " $\{-\}$ ", the shared memory location x protected by the reentrant lock l specified by the invariants m that requires the value stored in shared location is even ($\exists a, b. x \mapsto a \wedge a = 2b$). According to CSL, the shared memory is well-formed and conforms to the invariant m when the corresponding lock is free. The precondition and postcondition for instructions in the example are straightforward, it is trivial to apply the inference rules in our framework to verify the assembly code with assertions. Note that the rule `LOCK` is applied to reason about the acquiring lock operation at the initial entry and the rule `RELOCK` is applied to reason about the second reacquiring lock operation. Only the first locking operation transfers the shared memory location x from the shared part to its private part and the second locking operation acquires nothing but increasing the reentrancy level by 1. We use the rules `REUNLOCK` and `UNLOCK` to reason about the first and second releasing operations respectively.

```
Initially : [x] = 0;
Thread ID : k
  //x protected by lock l
  1: lock l;
  2: [x] := [x] + 1;
  3: lock l;
  4: [x] := [x] + 1;
  5: unlock l;
  6: unlock l;
```

Figure 11. Reentrant Lock Example

$$\begin{aligned}
m &\stackrel{\text{def}}{=} \exists a, b. x \mapsto a \wedge a = 2b \\
\Gamma &\stackrel{\text{def}}{=} \{l \rightsquigarrow m\} \\
&\{-\{\text{emp}\}\} \\
&\text{lock } l; \\
&\{-\{m\} \wedge \text{own}_k(l, 1)\} \\
&\text{ld } r_1, x(r_0); \\
&\{-\{\exists a, b. [x \mapsto a] \wedge \text{own}_k(l, 1) \wedge r_1 = a \wedge a = 2b\}\} \\
&\text{addi } r_1, r_1, 1; \\
&\{-\{\exists a, b. [x \mapsto a] \wedge \text{own}_k(l, 1) \wedge r_1 = a + 1 \wedge a = 2b\}\} \\
&\text{st } r_1, x(r_0); \\
&\{-\{\exists a, b. [x \mapsto a + 1] \wedge \text{own}_k(l, 1) \wedge r_1 = a + 1 \wedge a = 2b\}\} \\
&\text{lock } l; \\
&\{-\{\exists a, b. [x \mapsto a + 1] \wedge \text{own}_k(l, 2) \wedge r_1 = a + 1 \wedge a = 2b\}\} \\
&\text{addi } r_1, r_1, 1; \\
&\{-\{\exists a, b. [x \mapsto a + 1] \wedge \text{own}_k(l, 2) \wedge r_1 = a + 2 \wedge a = 2b\}\} \\
&\text{st } r_1, x(r_0); \\
&\{-\{\exists a, b. [x \mapsto a + 2] \wedge \text{own}_k(l, 2) \wedge r_1 = a + 2 \wedge a = 2b\}\} \\
&\{-\{\exists a', b'. [x \mapsto a'] \wedge \text{own}_k(l, 2) \wedge r_1 = a' \wedge a' = 2b'\}\} \\
&\{-\{m\} \wedge \text{own}_k(l, 2)\} \\
&\text{unlock } l; \\
&\{-\{m\} \wedge \text{own}_k(l, 1)\} \\
&\text{unlock } l; \\
&\{-\{\text{emp}\}\}
\end{aligned}$$

Figure 12. Assembly Code with Assertions

V. RELATED WORK

Many approaches have been proposed for reasoning about properties of both sequential and concurrent programs [3], [11], [12], [13]. But most efforts on concurrent programs focus on the non-reentrant lock-based programs and do not consider the reentrant locks. As we present in this paper, there exist some differences between reasoning about concurrent programs with non-reentrant locks and those with reentrant locks.

Peter O'Hearn [1], [6] proposed CSL, which applies the local-reasoning idea from separation logic [14], [2] to verify shared-state concurrent programs with memory pointers. Separation logic assertions are used to capture ownerships of resources. Separating conjunction enforces the partition of resources. Verification of sequential threads in CSL is no different from verification of sequential programs. Memory modularity is supported by using separating conjunction and frame rules. However, the rule for acquiring and releasing resource in CSL cannot be directly applied to verify concurrent programs with reentrant mutual exclusive locks. We adapt CSL to a concurrent assembly language with reentrant locks.

In recent years, Shao *et al.* have developed CCAP[8], CMAP[15] and SAGL[13] as extensions to the PCC frame-

work to verify properties of concurrent programs using locks, which are treated as non-reentrant locks. And we present an extension to enable verification of concurrent program using reentrant locks.

A recent work [16] proposes a verification technique for a concurrent Java-like language with reentrant locks. The verification technique is based on permission accounting separation logic. The essential differences between [16] and our paper are: we focus on verifying concurrent assembly code with reentrant locks and develop an extension to the PCC framework; instead of using hand-writing proof, we provide machine-checkable proof for our framework.

VI. CONCLUSION

In this paper we have presented a system for verifying concurrent programs using reentrant locks. We modeled an assembly level machine with built-in reentrant locking primitives. We adapted concurrent separation logic to verifying concurrent assembly code with reentrant locks. We also used a simple example to demonstrate the effectiveness of our framework.

ACKNOWLEDGEMENT

We thank Prof. Zhong Shao (Yale University) and anonymous referees for their suggestions and comments on an earlier version of this paper. This research was supported in part by the National Natural Science Foundation of China under Grant No. 90718026 and gifts from Intel Corporation. Any opinions, findings, and conclusions contained in this document are those of the authors and do not reflect the views of these agencies.

REFERENCES

- [1] P. W. O’Hearn, “Resources, concurrency and local reasoning,” in *Proc. 15th Int’l Conf. on Concurrency Theory (CONCUR’04)*, ser. LNCS, vol. 3170, 2004, pp. 49–67.
- [2] J. C. Reynolds, “Separation logic: A logic for shared mutable data structures,” in *Proc. LICS’02*, Jul. 2002, pp. 55–74.
- [3] C. A. R. Hoare, “Towards a theory of parallel programming,” in *Operating Systems Techniques*, C. A. R. Hoare and R. H. Perrott, Eds. Academic Press, 1972, pp. 61–71.
- [4] The Coq Development Team, “The Coq proof assistant reference manual,” The Coq release v8.0, Oct. 2004.
- [5] G. Necula, “Proof-carrying code,” in *Proc. 24th ACM Symp. on Principles of Prog. Lang.* ACM Press, Jan. 1997, pp. 106–119.
- [6] P. W. O’Hearn, “Resources, concurrency, and local reasoning,” *Theoretical Computer Science*, vol. 375, no. 1-3, pp. 271–307, 2007.
- [7] D. Yu, N. A. Hamid, and Z. Shao, “Building certified libraries for PCC: Dynamic storage allocation,” in *Proc. 2003 European Symposium on Programming (ESOP’03)*, April 2003.
- [8] D. Yu and Z. Shao, “Verification of safety properties for concurrent assembly code,” in *Proc. 2004 ACM SIGPLAN Int’l Conf. on Functional Prog.*, September 2004, pp. 175–188.
- [9] A. K. Wright and M. Felleisen, “A syntactic approach to type soundness,” *Information and Computation*, vol. 115, no. 1, pp. 38–94, 1994.
- [10] M. Fu, Y. Zhang, , and Y. Li, “Formal reasoning about concurrent assembly code with reentrant locks. Coq implementation.” <http://sbg.ustcsz.edu.cn/vsync/papers/frcacrl>, Jan 2009.
- [11] S. Owicki and D. Gries, “Verifying properties of parallel programs: an axiomatic approach,” *Commun. ACM*, vol. 19, no. 5, pp. 279–285, 1976.
- [12] A. Cotsman, J. Berdine, B. Cook, N. Rinetzky, and M. Sagiv, “Local reasoning for storable locks and threads,” in *Proc. Fifth Asian Symp. on Prog. Lang. and Sys. (APLAS’07)*, LNCS 4807. Springer-Verlag, November 2007.
- [13] X. Feng, R. Ferreira, and Z. Shao, “On the relationship between concurrent separation logic and assume-guarantee reasoning,” Dept. of Computer Science, Yale University, New Haven, CT, Tech. Rep. YALEU/DCS/TR-1374 and Formulation in Coq, January 2007.
- [14] S. S. Ishtiaq and P. W. O’Hearn, “BI as an assertion language for mutable data structures,” in *Proc. 28th ACM Symp. on Principles of Prog. Lang.*, 2001, pp. 14–26.
- [15] X. Feng and Z. Shao, “Modular verification of concurrent assembly code with dynamic thread creation and termination,” in *Proc. ICFP’05*, 2005, pp. 254–267.
- [16] C. Haack, M. Huisman, and C. Hurlin, “Reasoning about java’s reentrant locks,” in *APLAS ’08: Proceedings of the 6th Asian Symposium on Programming Languages and Systems*. Berlin, Heidelberg: Springer-Verlag, 2008, pp. 171–187.