On the Relationship between Concurrent Separation Logic and Assume-Guarantee Reasoning*

Xinyu Feng Rodrigo Ferreira Zhong Shao

Department of Computer Science, Yale University
New Haven, CT 06520-8285, U.S.A.
{feng, rodrigo, shao}@cs.yale.edu

Abstract. We study the relationship between Concurrent Separation Logic (CSL) and the assume-guarantee (A-G) method (a.k.a. rely-guarantee method). We show in three steps that CSL can be treated as a specialization of the A-G method for well-synchronized concurrent programs. First, we present an A-G based program logic for a low-level language with built-in locking primitives. Then we extend the program logic with explicit separation of "private data" and "shared data", which provides better memory modularity. Finally, we show that CSL (adapted for the low-level language) can be viewed as a specialization of the extended A-G logic by enforcing the invariant that "*shared resources are well-formed outside of critical regions*". This work can also be viewed as a different approach (from Brookes') to proving the soundness of CSL: our CSL inference rules are proved as lemmas in the A-G based logic, whose soundness is established following the syntactic approach to proving soundness of type systems.

1 Introduction

It is hard to prove non-interference and correctness of shared-state concurrent programs because of the exponential state space. Memory aliasing makes concurrency verification even harder. Therefore a program logic supporting both thread modularity and memory modularity is the key to practical concurrency verification.

Peter O'Hearn [11, 10] proposed concurrent separation logic (CSL), which applies the local-reasoning idea from separation logic [7, 15] 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, following Owicki and Gries [13], CSL works only for *well-synchronized programs* in the sense that transfer of resource ownerships can only occur at entry and exit points of critical regions. It is unclear how to apply CSL to support general concurrent programs with ad-hoc synchronizations.

* This research is based on work supported in part by gifts from Intel and Microsoft, and NSF grants CCR-0524545. Any opinions, findings, and conclusions contained in this document are those of the authors and do not reflect the views of these agencies.

Another approach to modular verification of shared-state concurrent programs is the assume-guarantee method (a.k.a. rely-guarantee method) [8]. In this approach, invariants of state transitions are specified using assumptions and guarantees. Each thread ensures that its atomic transitions satisfy its guarantee to the environment (*i.e.,* the collection of all other threads) as long as its assumption is satisfied by the environment. Non-interference is guaranteed as long as threads have compatible specifications, *i.e.,* the guarantee of each thread satisfies the assumptions of all other threads. The A-G method supports thread modular verification in the sense that each thread is verified with regard to its own specifications, and without looking into code of other threads. It is very general and does not require language constructs for synchronizations. However, in each individual step of the verification, we need to prove that the state transition satisfies the guarantee. This makes proofs more complicated in A-G reasoning than in CSL. Also, assumptions and guarantees are usually complicated and hard to define, because they specify global invariants for all shared resources during the program execution.

In this paper we study the relationship between CSL and A-G reasoning. We propose the Separated A-G Logic (SAGL), which extends A-G reasoning with the local-reasoning idea in separation logic. Instead of treating all resources as shared, SAGL partitions resources into shared and private. Like in CSL, each thread has full access to its private resources, which are invisible to its environments. Shared resources can be accessed in two ways in SAGL: they can be accessed directly, or be converted into private first and then accessed. Conversions between shared and private can occur at any program point, instead of being coupled with critical regions. Both direct accesses and conversions are governed by guarantees, so that non-interference is ensured following A-G reasoning. Private resources are not specified in assumptions and guarantees, therefore specifications in SAGL are simpler and more modular than A-G reasoning.

We then show that CSL can be viewed as a specialization of SAGL with the invariant that *shared resources are well-formed outside of critical regions*. The specialization is pinned down by formalizing the CSL invariant as a specific assumption and guarantee in SAGL. Our formulation can also be viewed as a novel approach to proving the soundness of CSL. Different from Brookes' proof based on an action-trace semantics [2], we prove that CSL inference rules are lemmas in SAGL with the specific assumption and guarantee. The soundness of SAGL is then proved following the syntactic approach to type soundness [19]. The proofs are formalized in the Coq proof assistant [17].

Our study is based on an assembly language with RISC-style instructions and built-in lock/unlock and memory allocation/free primitives. Instead of using the high-level parallel language proposed by Hoare [6], 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 [4] in Coq 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 [9]. CSL and the A-G method studied in this paper are all adapted to this low-level language. The relationship between the low-level CSL and the original logic by O'Hearn [11, 10] is discussed in Sect. 7.

$$
\begin{array}{rrcl}
(\mathit{Program}) & \mathbb{P} & ::= & (\mathbb{M}, [\mathbb{T}_1, \ldots, \mathbb{T}_n], \mathbb{L}) \\
(\mathit{Thread}) & \mathbb{T}_i & ::= & (\mathbb{C}, \mathbb{R}, \mathbb{I}, i) \\
(\mathit{CodeHeap}) & \mathbb{C} & \in & \mathit{Labels} \rightharpoonup \mathit{InstrSeq} \\
(\mathit{Memory}) & \mathbb{M} & \in & \mathit{Labels} \rightharpoonup \mathit{Word} \\
(\mathit{RegFile}) & \mathbb{R} & \in & \mathit{Register} \rightarrow \mathit{Word} \\
(\mathit{LockMap}) & \mathbb{L} & ::= & \mathit{Locks} \rightharpoonup \{1, \ldots, n\} \\
(\mathit{Register}) & \mathtt{r} & ::= & \mathtt{r}_0 \,|\, \ldots \,|\, \mathtt{r}_{31} \\
(\mathit{Labels}) & \mathtt{f}, \mathtt{l} & ::= & i \;\; (\mathit{nat\ nums}) \\
(\mathit{Locks}) & l & ::= & i \;\; (\mathit{nat\ nums}) \\
(\mathit{Word}) & \mathtt{w} & ::= & i \;\; (\mathit{nat\ nums}) \\
(\mathit{InstrSeq}) & \mathbb{I} & ::= & \mathsf{j}\ \mathtt{f} \,|\, \mathsf{jr}\ \mathtt{r}_s \,|\, \iota; \mathbb{I} \\
(\mathit{Instr}) & \iota & ::= & \mathsf{add}\ \mathtt{r}_d, \mathtt{r}_s, \mathtt{r}_t \mid \mathsf{addi}\ \mathtt{r}_d, \mathtt{r}_s, i \mid \mathsf{alloc}\ \mathtt{r}_d, \mathtt{r}_s \mid \mathsf{beq}\ \mathtt{r}_s, \mathtt{r}_t, \mathtt{f} \mid \mathsf{bgt}\ \mathtt{r}_s, \mathtt{r}_t, \mathtt{f} \\
 & & & \mid \mathsf{free}\ \mathtt{r}_s \mid \mathsf{lock}\ l \mid \mathsf{ld}\ \mathtt{r}_t, i(\mathtt{r}_s) \mid \mathsf{sub}\ \mathtt{r}_d, \mathtt{r}_s, \mathtt{r}_t \mid \mathsf{st}\ \mathtt{r}_t, i(\mathtt{r}_s) \mid \mathsf{unlock}\ l
\end{array}
$$

Fig. 1. The Abstract Machine

In the rest of this paper, we first present our low-level language in Sect. 2. We then present an A-G based logic (AGL) for this language in Sect. 3. We extend AGL with local reasoning and propose SAGL in Sect. 4. In Sect. 5, we adapt the original CSL to the low-level language and formalize the relationship between CSL and SAGL. We use two examples to illustrate the use of SAGL in Sect. 6. Finally, we discuss related work and conclude in Sect. 7.

2 The Language

Figure 1 defines the model of an abstract machine and the syntax of the assembly language. The whole program state $\mathbb{P}$ contains a shared memory $\mathbb{M}$, a lock mapping $\mathbb{L}$ which maps a lock to the id of its owner thread, and n threads $[\mathbb{T}_1, \ldots, \mathbb{T}_n]$. The memory is modeled as a finite partial mapping from memory locations $\mathtt{l}$ (natural numbers) to word values (natural numbers). Each thread $\mathbb{T}_i$ contains its own code heap $\mathbb{C}$, register file $\mathbb{R}$, the instruction sequence $\mathbb{I}$ that is currently being executed, and its thread id i. Here we allow each thread to have its own register file, which is consistent with most implementation of thread libraries where the register file is saved in the execution context when a thread is preempted.

The code heap $\mathbb{C}$ maps code labels to instruction sequences, which is a list of assembly instructions ending with a jump instruction. The set of instructions we present here are the commonly used subsets in RISC machines. We also use lock/unlock primitives to do synchronization, and use alloc/free to do dynamic memory allocation and free.

The step relation ($\longmapsto$) of program states ($\mathbb{P}$) is defined in Fig. 2. We use the auxiliary relation $(\mathbb{M}, \mathbb{T}, \mathbb{L}) \xmapsto{\mathsf{t}} (\mathbb{M}', \mathbb{T}', \mathbb{L}')$ to define the effects of the execution of the thread $\mathbb{T}$. Here we follow the preemptive thread model where execution of threads can be preempted at any program point, but execution of individual instructions is *atomic*.

$$(\mathbb{M}, [\mathbb{T}_1, \ldots, \mathbb{T}_n], \mathbb{L}) \longmapsto (\mathbb{M}', [\mathbb{T}_1, \ldots, \mathbb{T}_{k-1}, \mathbb{T}'_k, \mathbb{T}_{k+1}, \ldots, \mathbb{T}_n], \mathbb{L}')$$
$$\text{if } (\mathbb{M}, \mathbb{T}_k, \mathbb{L}) \overset{\mathsf{t}}{\longmapsto} (\mathbb{M}', \mathbb{T}'_k, \mathbb{L}') \text{ for any } k;$$

where

$(\mathbb{M}, (\mathbb{C}, \mathbb{R}, \mathbb{I}, k), \mathbb{L}) \overset{\mathsf{t}}{\longmapsto} (\mathbb{M}', \mathbb{T}', \mathbb{L}')$		
if $\mathbb{I} =$	then $(\mathbb{M}', \mathbb{T}', \mathbb{L}') =$	
j $\mathtt{f}$	$(\mathbb{M}, (\mathbb{C}, \mathbb{R}, \mathbb{I}', k), \mathbb{L})$	where $\mathbb{I}' = \mathbb{C}(\mathtt{f})$
jr $\mathtt{r}_s$	$(\mathbb{M}, (\mathbb{C}, \mathbb{R}, \mathbb{I}', k), \mathbb{L})$	where $\mathbb{I}' = \mathbb{C}(\mathbb{R}(\mathtt{r}_s))$
bgt $\mathtt{r}_s, \mathtt{r}_t, \mathtt{f}; \mathbb{I}'$	$(\mathbb{M}, (\mathbb{C}, \mathbb{R}, \mathbb{I}', k), \mathbb{L})$ $(\mathbb{M}, (\mathbb{C}, \mathbb{R}, \mathbb{I}'', k), \mathbb{L})$	if $\mathbb{R}(\mathtt{r}_s) \leq \mathbb{R}(\mathtt{r}_t)$ if $\mathbb{R}(\mathtt{r}_s) > \mathbb{R}(\mathtt{r}_t)$ and $\mathbb{I}'' = \mathbb{C}(\mathtt{f})$
beq $\mathtt{r}_s, \mathtt{r}_t, \mathtt{f}; \mathbb{I}'$	$(\mathbb{M}, (\mathbb{C}, \mathbb{R}, \mathbb{I}', k), \mathbb{L})$ $(\mathbb{M}, (\mathbb{C}, \mathbb{R}, \mathbb{I}'', k), \mathbb{L})$	if $\mathbb{R}(\mathtt{r}_s) \neq \mathbb{R}(\mathtt{r}_t)$ if $\mathbb{R}(\mathtt{r}_s) = \mathbb{R}(\mathtt{r}_t)$ and $\mathbb{I}'' = \mathbb{C}(\mathtt{f})$
lock $l; \mathbb{I}'$	$(\mathbb{M}, (\mathbb{C}, \mathbb{R}, \mathbb{I}', k), \mathbb{L}\{l \rightsquigarrow k\})$ $(\mathbb{M}, (\mathbb{C}, \mathbb{R}, \mathbb{I}, k), \mathbb{L})$	if $l \notin dom(\mathbb{L})$ if $l \in dom(\mathbb{L})$
unlock $l; \mathbb{I}'$	$(\mathbb{M}, (\mathbb{C}, \mathbb{R}, \mathbb{I}', k), \mathbb{L} \setminus \{l\})$	if $\mathbb{L}(l) = k$
$\iota; \mathbb{I}'$ for other ι	$(\mathbb{M}', (\mathbb{C}, \mathbb{R}', \mathbb{I}', k), \mathbb{L})$	where $(\mathbb{M}', \mathbb{R}') = \mathsf{Next}_\iota\ (\mathbb{M}, \mathbb{R})$

and

if $\iota =$	then $\mathsf{Next}_\iota\ (\mathbb{M}, \mathbb{R}) =$	
add $\mathtt{r}_d, \mathtt{r}_s, \mathtt{r}_t$	$(\mathbb{M}, \mathbb{R}\{\mathtt{r}_d \rightsquigarrow \mathbb{R}(\mathtt{r}_s) + \mathbb{R}(\mathtt{r}_t)\})$	
addi $\mathtt{r}_d, \mathtt{r}_s, i$	$(\mathbb{M}, \mathbb{R}\{\mathtt{r}_d \rightsquigarrow \mathbb{R}(\mathtt{r}_s) + i\})$	
ld $\mathtt{r}_t, i(\mathtt{r}_s)$	$(\mathbb{M}, \mathbb{R}\{\mathtt{r}_t \rightsquigarrow \mathbb{M}(\mathbb{R}(\mathtt{r}_s) + i)\})$	when $\mathbb{R}(\mathtt{r}_s) + i \in dom(\mathbb{M})$
sub $\mathtt{r}_d, \mathtt{r}_s, \mathtt{r}_t$	$(\mathbb{M}, \mathbb{R}\{\mathtt{r}_d \rightsquigarrow \mathbb{R}(\mathtt{r}_s) - \mathbb{R}(\mathtt{r}_t)\})$	
st $\mathtt{r}_t, i(\mathtt{r}_s)$	$(\mathbb{M}\{\mathbb{R}(\mathtt{r}_s) + i \rightsquigarrow \mathbb{R}(\mathtt{r}_t)\}, \mathbb{R})$	when $\mathbb{R}(\mathtt{r}_s) + i \in dom(\mathbb{M})$
alloc $\mathtt{r}_d, \mathtt{r}_s$	$(\mathbb{M}\{\mathtt{l}, \ldots, \mathtt{l} + \mathbb{R}(\mathtt{r}_s) - 1 \rightsquigarrow _\}, \mathbb{R}\{\mathtt{r}_d \rightsquigarrow \mathtt{l}\})$ where $\mathtt{l}, \ldots, \mathtt{l} + \mathbb{R}(\mathtt{r}_s) - 1 \notin dom(\mathbb{M})$	
free $\mathtt{r}_s$	$(\mathbb{M} \setminus \{\mathbb{R}(\mathtt{r}_s)\}, \mathbb{R})$	when $\mathbb{R}(\mathtt{r}_s) \in dom(\mathbb{M})$

Fig. 2. Operational Semantics of the Machine

Operational semantics of most instructions are standard. Note that we do not support reentrant-locks. If the lock l has been acquired, execution of the "lock l" instruction will be blocked even if the lock is owned by the current thread. The relation Next_ι defines the effects of the sequential instruction ι over memory and register files.

Note the way we distinguish "blocking" states from "stuck" states caused by unsafe operations, *e.g.,* freeing dangling pointers. If an unsafe operation is made, there is no resulting state satisfying the step relation ($\overset{\mathsf{t}}{\longmapsto}$) for the current thread. If a thread tries to acquire a lock which has been taken, it stutters: the resulting state will be the same as the current one (therefore the lock instruction will be executed again).

3 AGL: an A-G Based Program Logic

In this section we present an A-G based program logic (AGL) for our assembly language. AGL is a variation of the CCAP logic [20] which applies the A-G method for

$$
\begin{array}{rrcl}
(\mathit{XState}) & \mathbb{X} & ::= & (\mathbb{M}, (\mathbb{R}, i), \mathbb{L}) \\
(\mathit{ProgSpec}) & \Phi & ::= & ([\Psi_1, \ldots, \Psi_n], [(\mathsf{A}_1, \mathsf{G}_1), \ldots, (\mathsf{A}_n, \mathsf{G}_n)]) \\
(\mathit{CdHpSpec}) & \Psi & ::= & \{\mathtt{f} \rightsquigarrow \mathtt{a}\}^* \\
(\mathit{Assertion}) & \mathtt{a} & \in & \mathit{XState} \rightarrow \mathit{Prop} \\
(\mathit{Assume}) & \mathsf{A} & \in & \mathit{XState} \rightarrow \mathit{XState} \rightarrow \mathit{Prop} \\
(\mathit{Guarantee}) & \mathsf{G} & \in & \mathit{XState} \rightarrow \mathit{XState} \rightarrow \mathit{Prop}
\end{array}
$$

Fig. 3. Specification Constructs for AGL

assembly code verification. Different from CCAP, AGL works for the preemptive thread model instead of the non-preemptive model.

Figure 3 shows the specification constructs for AGL. For each thread in the program, its specification contains three parts: the specification Ψ for the code heap, the assumption A and the guarantee G. The specification Φ of the whole program just groups specifications for each thread. We use CiC, our *meta-logic* mechanized by Coq [17], as the assertion language for assertions and program specifications. CiC corresponds to the higher-order predicate logic with inductive definitions via Curry-Howard isomorphism.

Assumptions and guarantees are meta-logic predicates over a pair of extended thread states $\mathbb{X}$, which contains the shared memory $\mathbb{M}$, the thread's register file $\mathbb{R}$ and id k, and the global lock mapping $\mathbb{L}$. The assumption A for a thread specifies the expected invariant of state transitions made by the environment. The arguments it takes are states before and after a transition, respectively. The guarantee G of a thread specifies the invariant of state transitions made by the thread.

The code heap specification Ψ assigns a precondition $\mathtt{a}$ to each instruction sequence in the code heap $\mathbb{C}$.The assertion $\mathtt{a}$ is a meta-logic predicate over the extended thread state $\mathbb{X}$. It ensures the safe execution of the corresponding instruction sequence. We do not assign postconditions to instruction sequences. Since each instruction sequence ends with a jump instruction, we use the assertion at the target address as the postcondition.

Inference rules. Inference rules of AGL are presented in Figs. 4 and 5. The PROG rule defines the well-formedness of the program $\mathbb{P}$ with respect to the program specification Φ and the set of preconditions $([\mathtt{a}_1, \ldots, \mathtt{a}_n])$ for the instruction sequences that are currently executed by all the threads. Checking the well-formedness of $\mathbb{P}$ involves two steps. First we check the compatibility of assumptions and guarantees for all the threads. The predicate NI is defined as follows:

$$
\begin{array}{l}
\mathsf{NI}([(\mathsf{A}_1, \mathsf{G}_1), \ldots, (\mathsf{A}_n, \mathsf{G}_n)]) \stackrel{\text{def}}{=} \forall i, j, \mathbb{M}, \mathbb{M}', \mathbb{R}_i, \mathbb{R}'_i, \mathbb{R}_j, \mathbb{L}, \mathbb{L}'. \\
\qquad i \neq j \rightarrow \mathsf{G}_i\ (\mathbb{M}, (\mathbb{R}_i, i), \mathbb{L})\ (\mathbb{M}', (\mathbb{R}'_i, i), \mathbb{L}') \rightarrow \mathsf{A}_j\ (\mathbb{M}, (\mathbb{R}_j, j), \mathbb{L})\ (\mathbb{M}', (\mathbb{R}_j, j), \mathbb{L}'),
\end{array} \tag{1}
$$

which simply says that the guarantee of each thread should satisfy assumptions of all other threads. Then we apply the THRD rule to check that implementation of each thread actually satisfies the specification. Each thread $\mathbb{T}_i$ is verified separately. therefore thread modularity is supported.

In the THRD rule, we require that the precondition $\mathtt{a}$ be satisfied by the current extended thread state $(\mathbb{M}, (\mathbb{R}, k), \mathbb{L})$; that the thread code heap satisfy its specification Ψ,

$\boxed{\Phi, [\mathsf{a}_1, \ldots, \mathsf{a}_n] \vdash \mathbb{P}}$ **(*Well-formed program*)**

$$\frac{\begin{array}{l}\Phi = ([\Psi_1, \ldots, \Psi_n], [(\mathsf{A}_1, \mathsf{G}_1), \ldots, (\mathsf{A}_n, \mathsf{G}_n)]) \\ \mathsf{NI}([(\mathsf{A}_1, \mathsf{G}_1), \ldots, (\mathsf{A}_n, \mathsf{G}_n)]) \qquad \Psi_k, \mathsf{A}_k, \mathsf{G}_k \vdash \{\mathsf{a}_k\}\,(\mathbb{M}, \mathbb{T}_k, \mathbb{L}) \ \text{ for all } k\end{array}}{\Phi, [\mathsf{a}_1, \ldots, \mathsf{a}_n] \vdash (\mathbb{M}, [\mathbb{T}_1, \ldots, \mathbb{T}_n], \mathbb{L})} \ (\text{PROG})$$

$\boxed{\Psi, \mathsf{A}, \mathsf{G} \vdash \{\mathsf{a}\}\,(\mathbb{M}, \mathbb{T}, \mathbb{L})}$ **(*Well-formed thread*)**

$$\frac{\mathsf{a}\,(\mathbb{M}, (\mathbb{R}, k), \mathbb{L}) \quad \Psi, \mathsf{A}, \mathsf{G} \vdash \mathbb{C} : \Psi \quad \Psi, \mathsf{A}, \mathsf{G} \vdash \{\mathsf{a}\}\,\mathbb{I}}{\Psi, \mathsf{A}, \mathsf{G} \vdash \{\mathsf{a}\}\,(\mathbb{M}, (\mathbb{C}, \mathbb{R}, \mathbb{I}, k), \mathbb{L})} \ (\text{THRD})$$

$\boxed{\Psi, \mathsf{A}, \mathsf{G} \vdash \mathbb{C} : \Psi'}$ **(*Well-formed code heap*)**

$$\frac{\forall \mathtt{f} \in dom(\Psi') : \quad \Psi, \mathsf{A}, \mathsf{G} \vdash \{\Psi'(\mathtt{f})\}\,\mathbb{C}(\mathtt{f})}{\Psi, \mathsf{A}, \mathsf{G} \vdash \mathbb{C} : \Psi'} \ (\text{CDHP})$$

Fig. 4. AGL Inference Rules

A and G; and that it be safe to execute the current instruction sequence $\mathbb{I}$ under the precondition a and the thread specification.

The CDHP rule checks the well-formedness of thread code heaps. It requires that each instruction sequence specified in Ψ' be well-formed with respect to the imported interfaces specified in Ψ, the assumption A and the guarantee G.

The SEQ rule and the JR rule ensure that it is safe to execute the instruction sequence if the precondition is satisfied. If the instruction sequence starts with a normal sequential instruction ι, we need to come up with an assertion a' which serves both as the postcondition of ι and as the precondition of the remaining instruction sequence. Also we need to ensure that, if the current thread is preempted at a state satisfying a, a must be preserved by any state transitions (by other threads) satisfying the assumption A. This is enforced by $(\mathsf{a} \circ \mathsf{A}) \Rightarrow \mathsf{a}$:

$$(\mathsf{a} \circ \mathsf{A}) \Rightarrow \mathsf{a} \ \stackrel{\text{def}}{=} \ \forall \mathbb{X}, \mathbb{X}'.\ \mathsf{a}\,\mathbb{X} \wedge \mathsf{A}\,\mathbb{X}\,\mathbb{X}' \rightarrow \mathsf{a}\,\mathbb{X}'.$$

If we reach the last jump instruction of the instruction sequence, the JR rule requires that the assertion assigned to the target address in Ψ be satisfied after the jump. It also requires that a be preserved by state transitions satisfying A. Here we use the syntactic sugar $\forall X @ (x_1, \ldots, x_n).\ P(X, x_1, \ldots, x_n)$ to mean that, for all tuple X containing elements $x_1, \ldots, x_n$, the predicate P holds. It is formally defined as:

$$\forall X, x_1, \ldots, x_n. (X = (x_1, \ldots, x_n)) \rightarrow P(X, x_1, \ldots, x_n).$$

The notation $\lambda X @ (x_1, \ldots, x_n).\ f(X, x_1, \ldots, x_n)$ that we use later is defined similarly. The rule for direct jumps (j $\mathtt{f}$) is similar to the JR rule and requires no further explanation.

Instruction rules require that the precondition ensure the safe execution of the instruction; and that the resulting state satisfy the postcondition. Also, if shared states ($\mathbb{M}$ and $\mathbb{L}$) are updated by the instruction, we need to ensure that the update satisfies the

$\boxed{\Psi, \mathsf{A}, \mathsf{G} \vdash \{\mathtt{a}\}\,\mathbb{I}}$ (***Well-formed instr. sequences***)

$$\frac{\Psi, \mathsf{A}, \mathsf{G} \vdash \{\mathtt{a}\}\,\iota\,\{\mathtt{a}'\} \quad \Psi, \mathsf{A}, \mathsf{G} \vdash \{\mathtt{a}'\}\,\mathbb{I} \quad (\mathtt{a} \circ \mathsf{A}) \Rightarrow \mathtt{a}}{\Psi, \mathsf{A}, \mathsf{G} \vdash \{\mathtt{a}\}\,\iota;\mathbb{I}} \text{ (SEQ)}$$

$$\frac{\forall \mathbb{X}@(\mathbb{M}, (\mathbb{R}, k), \mathbb{L}).\ \mathtt{a}\ \mathbb{X} \rightarrow \Psi(\mathbb{R}(\mathtt{r}_s))\ \mathbb{X} \quad (\mathtt{a} \circ \mathsf{A}) \Rightarrow \mathtt{a}}{\Psi, \mathsf{A}, \mathsf{G} \vdash \{\mathtt{a}\}\,\mathsf{jr}\ \mathtt{r}_s} \text{ (JR)}$$

$$\frac{\forall \mathbb{X}@(\mathbb{M}, (\mathbb{R}, k), \mathbb{L}).\ \mathtt{a}\ \mathbb{X} \rightarrow \Psi(\mathtt{f})\ \mathbb{X} \quad (\mathtt{a} \circ \mathsf{A}) \Rightarrow \mathtt{a}}{\Psi, \mathsf{A}, \mathsf{G} \vdash \{\mathtt{a}\}\,\mathsf{j}\ \mathtt{f}} \text{ (J)}$$

$\boxed{\Psi, \mathsf{A}, \mathsf{G} \vdash \{\mathtt{a}\}\,\iota\,\{\mathtt{a}'\}}$ (***Well-formed instructions***)

$$\frac{\begin{array}{r}\forall \mathbb{X}@(\mathbb{M}, (\mathbb{R}, k), \mathbb{L}).\ \mathtt{a}\ \mathbb{X} \wedge l \notin dom(\mathbb{L}) \rightarrow \mathtt{a}'\ \mathbb{X}' \wedge \mathsf{G}\ \mathbb{X}\ \mathbb{X}' \\ \text{where } \mathbb{X}' = (\mathbb{M}, (\mathbb{R}, k), \mathbb{L}\{l \rightsquigarrow k\}).\end{array}}{\Psi, \mathsf{A}, \mathsf{G} \vdash \{\mathtt{a}\}\,\mathsf{lock}\ l\,\{\mathtt{a}'\}} \text{ (LOCK)}$$

$$\frac{\begin{array}{r}\forall \mathbb{X}@(\mathbb{M}, (\mathbb{R}, k), \mathbb{L}).\ \mathtt{a}\ \mathbb{X} \rightarrow \mathbb{L}(l) = k \wedge \mathtt{a}'\ \mathbb{X}' \wedge \mathsf{G}\ \mathbb{X}\ \mathbb{X}' \\ \text{where } \mathbb{X}' = (\mathbb{M}, (\mathbb{R}, k), \mathbb{L} \setminus \{l\}).\end{array}}{\Psi, \mathsf{A}, \mathsf{G} \vdash \{\mathtt{a}\}\,\mathsf{unlock}\ l\,\{\mathtt{a}'\}} \text{ (UNLOCK)}$$

$$\frac{\begin{array}{r}\forall \mathbb{X}@(\mathbb{M}, (\mathbb{R}, k), \mathbb{L}).\forall \mathtt{l}.\ \mathtt{a}\ \mathbb{X} \wedge (\{\mathtt{l}, \ldots, \mathtt{l}{+}\mathbb{R}(\mathtt{r}_s){-}1\} \cap dom(\mathbb{M}) = \emptyset) \rightarrow \\ \mathbb{R}(\mathtt{r}_s) > 0 \wedge \mathtt{a}'\ \mathbb{X}' \wedge \mathsf{G}\ \mathbb{X}\ \mathbb{X}' \\ \text{where } \mathbb{X}' = (\mathbb{M}\{\mathtt{l}, \ldots, \mathtt{l}{+}\mathbb{R}(\mathtt{r}_s){-}1 \rightsquigarrow _\}, (\mathbb{R}\{\mathtt{r}_d \rightsquigarrow \mathtt{l}\}, k), \mathbb{L})\end{array}}{\Psi, \mathsf{A}, \mathsf{G} \vdash \{\mathtt{a}\}\,\mathsf{alloc}\ \mathtt{r}_d, \mathtt{r}_s\,\{\mathtt{a}'\}} \text{ (ALLOC)}$$

$$\frac{\begin{array}{r}\forall \mathbb{X}@(\mathbb{M}, (\mathbb{R}, k), \mathbb{L}).\ \mathtt{a}\ \mathbb{X} \rightarrow \mathbb{R}(\mathtt{r}_s) \in dom(\mathbb{M}) \wedge \mathtt{a}'\ \mathbb{X}' \wedge \mathsf{G}\ \mathbb{X}\ \mathbb{X}' \\ \text{where } \mathbb{X}' = (\mathbb{M} \setminus \{\mathbb{R}(\mathtt{r}_s)\}, (\mathbb{R}, k), \mathbb{L})\end{array}}{\Psi, \mathsf{A}, \mathsf{G} \vdash \{\mathtt{a}\}\,\mathsf{free}\ \mathtt{r}_s\,\{\mathtt{a}'\}} \text{ (FREE)}$$

$$\frac{\begin{array}{r}\forall \mathbb{X}@(\mathbb{M}, (\mathbb{R}, k), \mathbb{L}).\ \mathtt{a}\ \mathbb{X} \rightarrow (\mathbb{R}(\mathtt{r}_s){+}i) \in dom(\mathbb{M}) \wedge \mathtt{a}'\ \mathbb{X}' \wedge \mathsf{G}\ \mathbb{X}\ \mathbb{X}' \\ \text{where } \mathbb{X}' = (\mathbb{M}\{\mathbb{R}(\mathtt{r}_s){+}i \rightsquigarrow \mathbb{R}(\mathtt{r}_t)\}, (\mathbb{R}, k), \mathbb{L})\end{array}}{\Psi, \mathsf{A}, \mathsf{G} \vdash \{\mathtt{a}\}\,\mathsf{st}\ \mathtt{r}_t, i(\mathtt{r}_s)\,\{\mathtt{a}'\}} \text{ (SW)}$$

$$\frac{\begin{array}{r}\forall \mathbb{X}@(\mathbb{M}, (\mathbb{R}, k), \mathbb{L}).\ \mathtt{a}\ \mathbb{X} \rightarrow (\mathbb{R}(\mathtt{r}_s){+}i) \in dom(\mathbb{M}) \wedge \mathtt{a}'\ \mathbb{X}' \\ \text{where } \mathbb{X}' = (\mathbb{M}, (\mathbb{R}\{\mathtt{r}_d \rightsquigarrow \mathbb{M}(\mathbb{R}(\mathtt{r}_s){+}i)\}, k), \mathbb{L})\end{array}}{\Psi, \mathsf{A}, \mathsf{G} \vdash \{\mathtt{a}\}\,\mathsf{ld}\ \mathtt{r}_d, i(\mathtt{r}_s)\,\{\mathtt{a}'\}} \text{ (LW)}$$

$$\frac{\forall \mathbb{X}@(\mathbb{M}, (\mathbb{R}, k), \mathbb{L}).\ \mathtt{a}\ \mathbb{X} \rightarrow ((\mathbb{R}(\mathtt{r}_s) \leq \mathbb{R}(\mathtt{r}_t) \rightarrow \mathtt{a}'\ \mathbb{X}) \wedge (\mathbb{R}(\mathtt{r}_s) > \mathbb{R}(\mathtt{r}_t) \rightarrow \Psi(\mathtt{f})\ \mathbb{X}))}{\Psi, \mathsf{A}, \mathsf{G} \vdash \{\mathtt{a}\}\,\mathsf{bgt}\ \mathtt{r}_s, \mathtt{r}_t, \mathtt{f}\,\{\mathtt{a}'\}} \text{ (BGT)}$$

$$\frac{\forall \mathbb{X}@(\mathbb{M}, (\mathbb{R}, k), \mathbb{L}).\ \mathtt{a}\ \mathbb{X} \rightarrow ((\mathbb{R}(\mathtt{r}_s) \neq \mathbb{R}(\mathtt{r}_t) \rightarrow \mathtt{a}'\ \mathbb{X}) \wedge (\mathbb{R}(\mathtt{r}_s) = \mathbb{R}(\mathtt{r}_t) \rightarrow \Psi(\mathtt{f})\ \mathbb{X}))}{\Psi, \mathsf{A}, \mathsf{G} \vdash \{\mathtt{a}\}\,\mathsf{beq}\ \mathtt{r}_s, \mathtt{r}_t, \mathtt{f}\,\{\mathtt{a}'\}} \text{ (BEQ)}$$

Fig. 5. AGL Inference Rules (cont'd)

guarantee G. For the lock instruction, if the control falls through, we know that the lock is not held by any thread. This extra knowledge can be used together with the precondition a to show the postcondition is satisfied by the resulting state. The rest of instruction rules are straightforward and will not be explained here.

Soundness. The soundness of AGL shows that the PROG rule enforces the non-interference and the partial correctness of programs with respect to the specifications. It is proved following the syntactic approach [19] to proving soundness of type systems. We first prove the following progress and preservation lemma.

Lemma 1 (AGL-Progress). *For any program* $\mathbb{P} = (\mathbb{M}, [\mathbb{T}_1, \ldots, \mathbb{T}_n], \mathbb{L})$*, if* $\Phi, [\mathsf{a}_1, \ldots, \mathsf{a}_n] \vdash \mathbb{P}$*, then, for any thread* $\mathbb{T}_k$*, there exist* $\mathbb{M}'$*,* $\mathbb{T}'_k$ *and* $\mathbb{L}'$ *such that* $(\mathbb{M}, \mathbb{T}_k, \mathbb{L}) \xmapsto{\mathsf{t}} (\mathbb{M}', \mathbb{T}'_k, \mathbb{L}')$*.*

Lemma 2 (AGL-Preservation). *If* $\Phi, [\mathsf{a}_1, \ldots, \mathsf{a}_n] \vdash \mathbb{P}$ *and* $(\mathbb{P} \longmapsto \mathbb{P}')$*, then there exist* $\mathsf{a}'_1, \ldots, \mathsf{a}'_n$ *such that* $\Phi, [\mathsf{a}'_1, \ldots, \mathsf{a}'_n] \vdash \mathbb{P}'$*.*

The soundness theorem follows the progress and preservation lemmas.

Theorem 3 (AGL-Soundness). *For any program* $\mathbb{P}$ *with specification* $\Phi = [\Psi_1, \ldots, \Psi_n], [(\mathsf{A}_1, \mathsf{G}_1), \ldots, (\mathsf{A}_n, \mathsf{G}_n)]$*, if* $\Phi, [\mathsf{a}_1, \ldots, \mathsf{a}_n] \vdash \mathbb{P}$*, then,*

- *for any m, there exists a* $\mathbb{P}'$ *such that* $(\mathbb{P} \longmapsto^m \mathbb{P}')$*;*
- *for any m and* $\mathbb{P}' = (\mathbb{M}', [\mathbb{T}'_1, \ldots, \mathbb{T}'_n], \mathbb{L}')$*, if* $(\mathbb{P} \longmapsto^m \mathbb{P}')$*, then,*
  - $\Phi, [\mathsf{a}'_1, \ldots, \mathsf{a}'_n] \vdash \mathbb{P}'$ *for some* $\mathsf{a}'_1, \ldots, \mathsf{a}'_n$*;*
  - *for any k, there exist* $\mathbb{M}''$*,* $\mathbb{T}''_k$ *and* $\mathbb{L}''$ *such that* $(\mathbb{M}', \mathbb{T}'_k, \mathbb{L}') \xmapsto{\mathsf{t}} (\mathbb{M}'', \mathbb{T}''_k, \mathbb{L}'')$*;*
  - *for any k, if* $\mathbb{T}'_k = (\mathbb{C}_k, \mathbb{R}'_k, \mathsf{j}\ f, k)$*, then* $\Psi_k(f)\ (\mathbb{M}', (\mathbb{R}'_k, k), \mathbb{L}')$ *holds;*
  - *for any k, if* $\mathbb{T}'_k = (\mathbb{C}_k, \mathbb{R}'_k, \mathsf{jr}\ r_s, k)$*, then* $\Psi_k(\mathbb{R}'_k(r_s))\ (\mathbb{M}', (\mathbb{R}'_k, k), \mathbb{L}')$ *holds;*
  - *for any k, if* $\mathbb{T}'_k = (\mathbb{C}_k, \mathbb{R}'_k, \mathsf{beq}\ r_s, r_t, f; \mathbb{I}, k)$ *and* $\mathbb{R}'_k(r_s) = \mathbb{R}'_k(r_t)$*, then* $\Psi_k(f)\ (\mathbb{M}', (\mathbb{R}'_k, k), \mathbb{L}')$ *holds;*
  - *for any k, if* $\mathbb{T}'_k = (\mathbb{C}_k, \mathbb{R}'_k, \mathsf{bgt}\ r_s, r_t, f; \mathbb{I}, k)$ *and* $\mathbb{R}'_k(r_s) > \mathbb{R}'_k(r_t)$*, then* $\Psi_k(f)\ (\mathbb{M}', (\mathbb{R}'_k, k), \mathbb{L}')$ *holds.*

Note that our AGL does not guarantee deadlock-freedom, which can be easily achieved by enforcing a partial order of lock acquiring in the LOCK rule.

4 SAGL: Separated A-G Logic

AGL is a general program logic supporting thread modular verification of concurrent code. However, because it treats all memory as shared resources, it does not have good memory modularity, and assumptions and guarantees are hard to define and use. For instance, suppose we partition the memory into n blocks and each block is assigned to one thread. Each thread simply works on its own part of memory and never accesses other parts. This scenario is not unusual in parallel programs and non-interference is trivially satisfied. However, to certify the code in AGL, the assumption for each thread has to be like "my private part of memory is never updated by others", and the guarantee is like "I will not touch other threads' private memory". During program verification, we have to prove for each individual instruction that the guarantee is not broken, even if

$$
\begin{array}{rrcl}
(\mathit{CdHpSpec}) & \Psi & ::= & \{\mathtt{f} \rightsquigarrow (\mathsf{a}, \mathsf{v})\}^* \\
(\mathit{Assertion}) & \mathsf{a}, \mathsf{v} & \in & \mathit{XState} \rightarrow \mathit{Prop}
\end{array}
$$

Fig. 6. Extension of AGL Specification Constructs in SAGL

$\boxed{\Phi, [(\mathsf{a}_1, \mathsf{v}_1), \ldots, (\mathsf{a}_n, \mathsf{v}_n)] \vdash \mathbb{P}}$ (***Well-formed program***)

$$
\frac{\begin{array}{l}\Phi = ([\Psi_1, \ldots, \Psi_n], [(\mathsf{A}_1, \mathsf{G}_1), \ldots, (\mathsf{A}_n, \mathsf{G}_n)]) \quad \mathsf{NI}([(\mathsf{A}_1, \mathsf{G}_1), \ldots, (\mathsf{A}_n, \mathsf{G}_n)]) \\ \mathbb{M}_s \uplus \mathbb{M}_1 \uplus \cdots \uplus \mathbb{M}_n = \mathbb{M} \quad \Psi_k, \mathsf{A}_k, \mathsf{G}_k \vdash \{(\mathsf{a}_k, \mathsf{v}_k)\}\,(\mathbb{M}_s, \mathbb{M}_k, \mathbb{T}_k, \mathbb{L}) \ \text{ for all } k\end{array}}{\Phi, [(\mathsf{a}_1, \mathsf{v}_1), \ldots, (\mathsf{a}_n, \mathsf{v}_n)] \vdash (\mathbb{M}, [\mathbb{T}_1, \ldots, \mathbb{T}_n], \mathbb{L})} \ (\text{PROG})
$$

$\boxed{\Psi, \mathsf{A}, \mathsf{G} \vdash \{(\mathsf{a}, \mathsf{v})\}\,(\mathbb{M}_s, \mathbb{M}_v, \mathbb{T}, \mathbb{L})}$ (***Well-formed thread***)

$$
\frac{\mathsf{a}\,(\mathbb{M}_s, (\mathbb{R}, k), \mathbb{L}) \quad \mathsf{v}\,(\mathbb{M}_v, (\mathbb{R}, k), \mathbb{L}|_k) \quad \Psi, \mathsf{A}, \mathsf{G} \vdash \mathbb{C} : \Psi \quad \Psi, \mathsf{A}, \mathsf{G} \vdash \{(\mathsf{a}, \mathsf{v})\}\,\mathbb{I}}{\Psi, \mathsf{A}, \mathsf{G} \vdash \{(\mathsf{a}, \mathsf{v})\}\,(\mathbb{M}_s, \mathbb{M}_v, (\mathbb{C}, \mathbb{R}, \mathbb{I}, k), \mathbb{L})} \ (\text{THRD})
$$

$\boxed{\Psi, \mathsf{A}, \mathsf{G} \vdash \mathbb{C} : \Psi'}$ (***Well-formed code heap***)

$$
\frac{\forall \mathtt{f} \in \mathit{dom}(\Psi') : \quad \Psi, \mathsf{A}, \mathsf{G} \vdash \{\Psi'(\mathtt{f})\}\,\mathbb{C}(\mathtt{f})}{\Psi, \mathsf{A}, \mathsf{G} \vdash \mathbb{C} : \Psi'} \ (\text{CDHP})
$$

Fig. 7. SAGL Inference Rules

it is trivially true. To make things worse, if each thread dynamically allocates memory and uses the allocated memory as private resources, as shown in the example in Sect. 6, the domain of private memory becomes dynamic, which makes it very hard to define the assumption and guarantee.

In this section, we extend AGL with explicit partition of private resources and shared resources. The extended logic, which we call Separated A-G Logic (SAGL), has much better support of memory modularity than AGL without sacrificing any expressiveness. Borrowing the local-reasoning idea in separation logic, private resources of one thread are not visible to other threads, therefore will not be touched by others. Assumptions and guarantees in SAGL only specifies shared resources, therefore in the scenario above the definition of SAGL assumption and guarantee becomes completely trivial since there is no shared resources. The dynamic domain of private memory caused by dynamic memory allocation/free is no longer a challenge to define assumptions and guarantees because private memory does not have to be specified.

Figure 6 shows our extensions of AGL specifications for SAGL. In the specification Ψ of each thread code heap, the precondition assigned to each code label now becomes a pair of assertions (a, v). The assertion a plays the same role as in AGL. It specifies the shared resources (all memory are treated as shared in AGL). The assertion v specifies the private resources of the thread. Other threads' private resources are not specified.

$\boxed{\Psi, \mathsf{A}, \mathsf{G} \vdash \{(\mathsf{a}, \mathsf{v})\}\,\mathbb{I}}$ **(*Well-formed instr. sequences*)**

$$\frac{\Psi, \mathsf{A}, \mathsf{G} \vdash \{(\mathsf{a}, \mathsf{v})\}\,\iota\,\{(\mathsf{a}', \mathsf{v}')\} \quad \Psi, \mathsf{A}, \mathsf{G} \vdash \{(\mathsf{a}', \mathsf{v}')\}\,\mathbb{I} \quad (\mathsf{a} \circ \mathsf{A}) \Rightarrow \mathsf{a} \quad \mathsf{Precise}(\mathsf{a})}{\Psi, \mathsf{A}, \mathsf{G} \vdash \{(\mathsf{a}, \mathsf{v})\}\,\iota; \mathbb{I}} \ (\textsc{seq})$$

$$\frac{\mathsf{Precise}(\mathsf{a}) \quad (\mathsf{a} \circ \mathsf{A}) \Rightarrow \mathsf{a} \quad \forall \mathbb{X} @ (\mathbb{M}, (\mathbb{R}, k), \mathbb{L}).\ (\mathsf{a} * \mathsf{v})\ \mathbb{X} \rightarrow (\mathsf{a}' * \mathsf{v}')\ \mathbb{X} \wedge (\lfloor \mathsf{G} \rfloor_{(\mathsf{a}, \mathsf{a}')}\ \mathbb{X}\ \mathbb{X}) \text{ where } (\mathsf{a}', \mathsf{v}') = \Psi(\mathbb{R}(\mathtt{r}_s))}{\Psi, \mathsf{A}, \mathsf{G} \vdash \{(\mathsf{a}, \mathsf{v})\}\,\mathsf{jr}\ \mathtt{r}_s} \ (\textsc{jr})$$

$\boxed{\Psi, \mathsf{A}, \mathsf{G} \vdash \{(\mathsf{a}, \mathsf{v})\}\,\iota\,\{(\mathsf{a}', \mathsf{v}')\}}$ **(*Well-formed instructions*)**

$$\frac{\forall \mathbb{X} @ (\mathbb{M}, (\mathbb{R}, k), \mathbb{L}).\ (\mathsf{a} * \mathsf{v})\ \mathbb{X} \wedge l \notin dom(\mathbb{L}) \rightarrow (\mathsf{a}' * \mathsf{v}')\ \mathbb{X}' \wedge (\lfloor \mathsf{G} \rfloor_{(\mathsf{a}, \mathsf{a}')}\ \mathbb{X}\ \mathbb{X}') \text{ where } \mathbb{X}' = (\mathbb{M}, (\mathbb{R}, k), \mathbb{L}\{l \rightsquigarrow k\})}{\Psi, \mathsf{A}, \mathsf{G} \vdash \{(\mathsf{a}, \mathsf{v})\}\,\mathsf{lock}\ l\,\{(\mathsf{a}', \mathsf{v}')\}} \ (\textsc{lock})$$

$$\frac{\forall \mathbb{X} @ (\mathbb{M}, (\mathbb{R}, k), \mathbb{L}).\ (\mathsf{a} * \mathsf{v})\ \mathbb{X} \rightarrow \mathbb{L}(l) = k \wedge (\mathsf{a}' * \mathsf{v}')\ \mathbb{X}' \wedge (\lfloor \mathsf{G} \rfloor_{(\mathsf{a}, \mathsf{a}')}\ \mathbb{X}\ \mathbb{X}') \text{ where } \mathbb{X}' = (\mathbb{M}, (\mathbb{R}, k), \mathbb{L} \setminus \{l\})}{\Psi, \mathsf{A}, \mathsf{G} \vdash \{(\mathsf{a}, \mathsf{v})\}\,\mathsf{unlock}\ l\,\{(\mathsf{a}', \mathsf{v}')\}} \ (\textsc{unlock})$$

Fig. 8. SAGL Inference Rules (cont'd)

Inference rules. The inference rules of SAGL are shown in Figs. 7 and 8. They look very similar to AGL rules. In the PROG rule, as in AGL, we check the compatibility of assumptions and guarantees, and check the well-formedness of each thread. However, here we require that there be a partition of memory into $n+1$ parts: one part $\mathbb{M}_s$ is shared and other parts $\mathbb{M}_1, \ldots, \mathbb{M}_n$ are privately owned by the threads $\mathbb{T}_1, \ldots, \mathbb{T}_n$, respectively. When we check the well-formedness of thread $\mathbb{T}_k$, the memory in the extended thread state is not the global memory. It just contains $\mathbb{M}_s$ and $\mathbb{M}_k$.

The THRD rule in SAGL is similar to the one in AGL, except that the memory visible by each thread is separated into two parts: the shared $\mathbb{M}_s$ and the private $\mathbb{M}_v$. We require that assertions a and v hold over $\mathbb{M}_s$ and $\mathbb{M}_v$ respectively. Since v only specifies the private resource, we use the "filter" operator $\mathbb{L}|_k$ to prevent v from having access to the ownership information of locks not owned by the current thread:

$$(\mathbb{L}|_k)(l) \stackrel{\text{def}}{=} \begin{cases} k & \mathbb{L}(l) = k \\ \textit{undefined} & \text{otherwise} \end{cases} \tag{2}$$

i.e., $\mathbb{L}|_k$ is a subset of $\mathbb{L}$ which maps locks to k.

Instruction rules are shown in Fig. 8. In the SEQ rule, we use (a, v) as the precondition. However, to ensure that the precondition is preserved by state transitions satisfying A, we only check a (*i.e.,* we check $(\mathsf{a} \circ \mathsf{A}) \Rightarrow \mathsf{a}$) because A only specifies shared resources. We know that the private resources will not be touched by the environment. We require a to be precise to enforce the unique boundary between shared and private

resources. Following the definition in CSL [11], an assertion a is precise if and only if for any memory $\mathbb{M}$, there is at most one subset $\mathbb{M}'$ that satisfies a, *i.e.,*

$$\begin{aligned}\mathsf{Precise}(\mathsf{a}) \stackrel{\text{def}}{=} \forall \mathbb{M}, \mathbb{R}, k, \mathbb{L}, \mathbb{M}_1, \mathbb{M}_2.\ (\mathbb{M}_1 \subseteq \mathbb{M}) \wedge (\mathbb{M}_2 \subseteq \mathbb{M}) \wedge \\ \mathsf{a}\ (\mathbb{M}_1, (\mathbb{R}, k), \mathbb{L}) \wedge \mathsf{a}\ (\mathbb{M}_2, (\mathbb{R}, k), \mathbb{L}) \rightarrow \mathbb{M}_1 = \mathbb{M}_2\,.\end{aligned} \tag{3}$$

The JR rule requires a be precise and it be preserved by state transitions satisfying the assumption. Also, the specification assigned to the target address needs to be satisfied by the resulting state of the jump, and the identity state transition made by the jump satisfies the guarantee G. We use the separating conjunction of the shared and private predicates as the pre- and post-condition. We define $\mathsf{a} * \mathsf{v}$ as:

$$\begin{aligned}\mathsf{a} * \mathsf{v} \stackrel{\text{def}}{=} \lambda(\mathbb{M}, (\mathbb{R}, k), \mathbb{L}). \\ \exists \mathbb{M}_1, \mathbb{M}_2.(\mathbb{M}_1 \uplus \mathbb{M}_2 = \mathbb{M}) \wedge \mathsf{a}\ (\mathbb{M}_1, (\mathbb{R}, k), \mathbb{L}) \wedge \mathsf{v}\ (\mathbb{M}_2, (\mathbb{R}, k), \mathbb{L}|_k)\,.\end{aligned} \tag{4}$$

Again, the use of $\mathbb{L}|_k$ prevents v from having access to the ownership information of locks not owned by the current thread. We use $f_1 \uplus f_2$ to represent the union of finite partial mappings with disjoint domains.

To ensure G is satisfied over shared resources, we lift G to $\lfloor \mathsf{G} \rfloor_{(\mathsf{a},\mathsf{a}')}$:

$$\begin{aligned}\lfloor \mathsf{G} \rfloor_{(\mathsf{a},\mathsf{a}')} \stackrel{\text{def}}{=} \lambda \mathbb{X}@(\mathbb{M}, (\mathbb{R}, k), \mathbb{L}), \mathbb{X}'@(\mathbb{M}', (\mathbb{R}', k'), \mathbb{L}'). \\ \exists \mathbb{M}_1, \mathbb{M}_2, \mathbb{M}_1', \mathbb{M}_2'.\ (\mathbb{M}_1 \uplus \mathbb{M}_2 = \mathbb{M}) \wedge (\mathbb{M}_1' \uplus \mathbb{M}_2' = \mathbb{M}') \\ \wedge\ \mathsf{a}\ (\mathbb{M}_1, (\mathbb{R}, k), \mathbb{L}) \wedge \mathsf{a}'\ (\mathbb{M}_1', (\mathbb{R}', k'), \mathbb{L}') \\ \wedge\ \mathsf{G}\ (\mathbb{M}_1, (\mathbb{R}, k), \mathbb{L})\ (\mathbb{M}_1', (\mathbb{R}', k'), \mathbb{L}')\,,\end{aligned} \tag{5}$$

Here we use precise predicates a and a' to enforce the unique boundary between shared and private resources.

As expected, the SAGL rule for each individual instruction is almost the same as its counterpart in AGL, except that we always use the separating conjunction of predicates for shared and private resources. Each instruction rule requires that memory in states before and after the transition can be partitioned to private and shared; private parts satisfy private predicates and shared parts satisfy shared predicates and G.

It is important that we always combine shared predicates with private predicates instead of checking separately the relationship between a and a' and between v and v'. This gives us the ability to support *dynamic redistribution* of private and shared memory. Instead of enforcing static partition, we allow that part of private memory becomes shared under certain conditions and vice versa. As we will show in the next section, this ability makes our SAGL very expressive and is the enabling feature that makes the embedding of CSL into SAGL possible.

AGL can be viewed as a specialized version of SAGL where all the v's are set to emp (emp is an assertion which can only be satisfied by memory with empty domain).

Soundness. Soundness of SAGL is proved following the syntactic approach [19] to proving soundness of type systems.

Lemma 4. *If* $\Psi, \mathsf{A}, \mathsf{G} \vdash \mathbb{C} : \Psi$*, then* $dom(\Psi) \subseteq dom(\mathbb{C})$.

Lemma 5 (Thread-Progress). *If* $\Psi, \mathsf{A}, \mathsf{G} \vdash \{(\mathscr{a}, \mathsf{v})\}\,(\mathbb{M}_s, \mathbb{M}_v, \mathbb{T}_k, \mathbb{L})$, $dom(\mathbb{M}_s) \cap dom(\mathbb{M}_v) = \emptyset$, *and* $\mathbb{T}_k = (\mathbb{C}, \mathbb{R}, \mathbb{I}, k)$, *then there exist* $\mathbb{M}'$, $\mathbb{R}'$, $\mathbb{I}'$ *and* $\mathbb{L}'$ *such that* $(\mathbb{M}_s \cup \mathbb{M}_v, \mathbb{T}_k, \mathbb{L}) \xmapsto{\mathsf{t}} (\mathbb{M}', \mathbb{T}'_k, \mathbb{L}')$, *where* $\mathbb{T}'_k = (\mathbb{C}, \mathbb{R}', \mathbb{I}', k)$.

Proof sketch: From $\Psi, \mathsf{A}, \mathsf{G} \vdash \{(\mathsf{a}, \mathsf{v})\}\,(\mathbb{M}_s, \mathbb{M}_v, \mathbb{T}_k, \mathbb{L})$ and $dom(\mathbb{M}_s) \cap dom(\mathbb{M}_v) = \emptyset$, we know (1) $\Psi, \mathsf{A}, \mathsf{G} \vdash \{(\mathsf{a}, \mathsf{v})\}\,\mathbb{I}$; (2) $\Psi, \mathsf{A}, \mathsf{G} \vdash \mathbb{C} : \Psi$; and (3) $(\mathsf{a} * \mathsf{v})\ (\mathbb{M}_s \cup \mathbb{M}_v, (\mathbb{R}, k), \mathbb{L})$. If $\mathbb{I} = \mathsf{j}\ \mathtt{f}$ ($\mathbb{I} = \mathsf{jr}\ \mathtt{r}_s$), by (1) and the J (JR) rule, we know $\mathtt{f} \in dom(\Psi)$ ($\mathbb{R}(\mathtt{r}_s) \in dom(\Psi)$). By Lemma 4, we know the target address is a valid code label in $\mathbb{C}$, therefore it is safe to execute the jump instruction. If $\mathbb{I} = \iota; \mathbb{I}'$, we know by (1) that there are a' and v' such that $\Psi, \mathsf{A}, \mathsf{G} \vdash \{(\mathsf{a}, \mathsf{v})\}\,\iota\,\{(\mathsf{a}', \mathsf{v}')\}$. By inspection of instruction rules, we know it is always safe to execute the instruction ι as long as the state satisfy $\mathsf{a} * \mathsf{v}$. Since we have (3), the thread can progress by executing ι. □

Lemma 6 (Thread-Progress-Monotone). *If* $(\mathbb{M}, \mathbb{T}_k, \mathbb{L}) \xmapsto{\mathsf{t}} (\mathbb{M}', \mathbb{T}'_k, \mathbb{L}')$, *where* $\mathbb{T}_k = (\mathbb{C}, \mathbb{R}, \mathbb{I}, k)$ *and* $\mathbb{T}'_k = (\mathbb{C}, \mathbb{R}', \mathbb{I}', k)$, *then, for any* $\mathbb{M}_0$ *such that* $dom(\mathbb{M}) \cap dom(\mathbb{M})_0 = \emptyset$, *there exists* $\mathbb{M}''$ *and* $\mathbb{R}''$ *such that* $(\mathbb{M} \cup \mathbb{M}_0, \mathbb{T}_k, \mathbb{L}) \xmapsto{\mathsf{t}} (\mathbb{M}'', \mathbb{T}''_k, \mathbb{L}')$, *where* $\mathbb{T}''_k = (\mathbb{C}, \mathbb{R}'', \mathbb{I}', k)$.

Proof sketch: The proof trivially follows the definition of the operational semantics. Note that we model memory $\mathbb{M}$ as a *finite* partial mapping, so the monotonicity also holds for alloc. □

Lemma 7 (SAGL-Progress). *For any program* $\mathbb{P} = (\mathbb{M}, [\mathbb{T}_1, \ldots, \mathbb{T}_n], \mathbb{L})$, *if* $\Phi, [(\mathscr{a}_1, \mathsf{v}_1), \ldots, (\mathscr{a}_n, \mathsf{v}_n)] \vdash \mathbb{P}$, *then, for any thread* $\mathbb{T}_k$, *there exist* $\mathbb{M}'$, $\mathbb{T}'_k$ *and* $\mathbb{L}'$ *such that* $(\mathbb{M}, \mathbb{T}_k, \mathbb{L}) \xmapsto{\mathsf{t}} (\mathbb{M}', \mathbb{T}'_k, \mathbb{L}')$.

Proof sketch: By inversion of the PROG rule, Lemma 4 and Lemma 6. □

Lemma 8. *If* $\Psi, \mathsf{A}, \mathsf{G} \vdash \{(\mathscr{a}, \mathsf{v})\}\,(\mathbb{M}_s, \mathbb{M}_v, \mathbb{T}_k, \mathbb{L})$, *then* $(\mathscr{a} \circ \mathsf{A}) \Rightarrow \mathscr{a}$ *and* $\mathsf{Precise}(\mathscr{a})$.

Lemma 9 (Thread-Preservation). *If* $\Psi, \mathsf{A}, \mathsf{G} \vdash \{(\mathscr{a}, \mathsf{v})\}\,(\mathbb{M}_s, \mathbb{M}_v, \mathbb{T}_k, \mathbb{L})$, $dom(\mathbb{M}_s) \cap dom(\mathbb{M}_v) = \emptyset$, *and* $(\mathbb{M}_s \cup \mathbb{M}_v, \mathbb{T}_k, \mathbb{L}) \xmapsto{\mathsf{t}} (\mathbb{M}', \mathbb{T}'_k, \mathbb{L}')$, *where* $\mathbb{T}_k = (\mathbb{C}, \mathbb{R}, \mathbb{I}, k)$ *and* $\mathbb{T}'_k = (\mathbb{C}, \mathbb{R}', \mathbb{I}', k)$, *then there exist* $\mathscr{a}'$, v', $\mathbb{M}'_s$ *and* $\mathbb{M}'_v$ *such that* $\mathbb{M}'_s \uplus \mathbb{M}'_v = \mathbb{M}'$, $\mathsf{G}\ (\mathbb{M}_s, (\mathbb{R}, k), \mathbb{L})\ (\mathbb{M}'_s, (\mathbb{R}', k), \mathbb{L}')$, *and* $\Psi, \mathsf{A}, \mathsf{G} \vdash \{(\mathscr{a}', \mathsf{v}')\}\,(\mathbb{M}'_s, \mathbb{M}'_v, \mathbb{T}'_k, \mathbb{L}')$.

Proof sketch: By $\Psi, \mathsf{A}, \mathsf{G} \vdash \{(\mathsf{a}, \mathsf{v})\}\,(\mathbb{M}_s, \mathbb{M}_v, \mathbb{T}_k, \mathbb{L})$ and inversion of the THRD rule, we know $\Psi, \mathsf{A}, \mathsf{G} \vdash \{(\mathsf{a}, \mathsf{v})\}\,\mathbb{I}$. Then we prove the lemma by inspection of the structure of $\mathbb{I}$ and the inversion of the corresponding instruction sequence rules and instruction rules. □

Lemma 10 (Thread-Inv). *If* $\Psi, \mathsf{A}, \mathsf{G} \vdash \{(\mathscr{a}, \mathsf{v})\}\,(\mathbb{M}_s, \mathbb{M}_v, \mathbb{T}_k, \mathbb{L})$, $dom(\mathbb{M}_s) \cap dom(\mathbb{M}_v) = \emptyset$, *and* $(\mathbb{M}_s \cup \mathbb{M}_v, \mathbb{T}_k, \mathbb{L}) \xmapsto{\mathsf{t}} (\mathbb{M}', \mathbb{T}'_k, \mathbb{L}')$, *where* $\mathbb{T}_k = (\mathbb{C}, \mathbb{R}, \mathbb{I}, k)$ *and* $\mathbb{T}'_k = (\mathbb{C}, \mathbb{R}', \mathbb{I}', k)$, *then*

- *if* $\mathbb{I} = \mathsf{j}\ f$, *then* $(\mathscr{a}' * \mathsf{v}')\ (\mathbb{M}', (\mathbb{R}', k), \mathbb{L}')$ *holds, where* $(\mathscr{a}', \mathsf{v}') = \Psi(f)$;
- *if* $\mathbb{I} = \mathsf{jr}\ r_s$, *then* $(\mathscr{a}' * \mathsf{v}')\ (\mathbb{M}', (\mathbb{R}', k), \mathbb{L}')$ *holds, where* $(\mathscr{a}', \mathsf{v}') = \Psi(\mathbb{R}(r_s))$;
- *if* $\mathbb{I} = \mathsf{beq}\ r_s, r_t, f; \mathbb{I}'$ *and* $\mathbb{R}(r_s) = \mathbb{R}(r_t)$, *then* $(\mathscr{a}' * \mathsf{v}')\ (\mathbb{M}', (\mathbb{R}', k), \mathbb{L}')$ *holds, where* $(\mathscr{a}', \mathsf{v}') = \Psi(f)$;

– *if* $\mathbb{I} = \mathsf{bgt}\ r_s, r_t, f; \mathbb{I}'$ *and* $\mathbb{R}(r_s) > \mathbb{R}(r_t)$*, then* $(\alpha' * \nu')\ (\mathbb{M}', (\mathbb{R}', k), \mathbb{L}')$ *holds, where* $(\alpha', \nu') = \Psi(f)$*;*

Proof sketch: By $\Psi, \mathsf{A}, \mathsf{G} \vdash \{(\mathsf{a}, \nu)\}\,(\mathbb{M}_s, \mathbb{M}_v, \mathbb{T}_k, \mathbb{L})$ and inversion of the THRD rule, we know $\Psi, \mathsf{A}, \mathsf{G} \vdash \{(\mathsf{a}, \nu)\}\,\mathbb{I}$. The lemma trivially follows inversion of the J rule, the JR rule, the BEQ rule and the BGT rule. (The J rule is similar to the JR rule. The BEQ rule and the BGT rule are similar to their counterparts in AGL. These rules are not shown in Fig. 8.) □

Lemma 11 (Thread-Frame). *If* $(\mathbb{M}, \mathbb{T}_k, \mathbb{L}) \xmapsto{\mathsf{t}} (\mathbb{M}', \mathbb{T}'_k, \mathbb{L}')$*,* $\mathbb{M} = \mathbb{M}_1 \uplus \mathbb{M}_2$*, and there exists* $\mathbb{M}''_1$ *and* $\mathbb{T}''_k$ *such that* $(\mathbb{M}_1, \mathbb{T}_k, \mathbb{L}) \xmapsto{\mathsf{t}} (\mathbb{M}''_1, \mathbb{T}''_k, \mathbb{L}')$*, then there exists* $\mathbb{M}'_1$ *such that* $\mathbb{M}' = \mathbb{M}'_1 \uplus \mathbb{M}_2$*, and* $(\mathbb{M}_1, \mathbb{T}_k, \mathbb{L}) \xmapsto{\mathsf{t}} (\mathbb{M}'_1, \mathbb{T}'_k, \mathbb{L}')$*.*

Proof sketch: By inspection of the operational semantics of instructions. □

Lemma 12 (SAGL-Preservation). *If* $\Phi, [(\alpha_1, \nu_1) \ldots, (\alpha_n, \nu_n)] \vdash \mathbb{P}$ *and* $(\mathbb{P} \longmapsto \mathbb{P}')$*, where* $\Phi = ([\Psi_1, \ldots, \Psi_n], [(\mathsf{A}_1, \mathsf{G}_1), \ldots, (\mathsf{A}_n, \mathsf{G}_n)])$*, then there exist* α'_1*,* ν'_1*, ...,* α'_n*,* ν'_n *such that* $\Phi, [(\alpha'_1, \nu'_1) \ldots, (\alpha'_n, \nu'_n)] \vdash \mathbb{P}'$*.*

Proof sketch: If the thread k executes its instruction in the step $\mathbb{P} \longmapsto \mathbb{P}'$, we know:

– the private memory of other threads will not be touched, following Lemma 11; and
– other threads' assertions a_i for shared resources are preserved, following Lemma 8, Lemma 9 and the non-interference of specifications, *i.e.,* $\mathsf{NI}([(\mathsf{A}_1, \mathsf{G}_1), \ldots, (\mathsf{A}_n, \mathsf{G}_n)])$.

Then, by the PROG rule, it is easy to prove $\Phi, [(\mathsf{a}'_1, \nu'_1) \ldots, (\mathsf{a}'_n, \nu'_n)] \vdash \mathbb{P}'$ for some a'_1, ν'_1, ..., a'_n, ν'_n. □

Finally, the soundness of SAGL is formulated in Theorem 13. In addition to the safety of well-formed programs, it also characterizes partial correctness: assertions assigned to labels in Ψ will hold whenever the labels are reached.

Theorem 13 (SAGL-Soundness). *For any program* $\mathbb{P}$ *with specification* $\Phi = ([\Psi_1, \ldots, \Psi_n], [(\mathsf{A}_1, \mathsf{G}_1), \ldots, (\mathsf{A}_n, \mathsf{G}_n)])$*, if* $\Phi, [(\alpha_1, \nu_1) \ldots, (\alpha_n, \nu_n)] \vdash \mathbb{P}$*, then,*

– *for any natural number m, there exists* $\mathbb{P}'$ *such that* $(\mathbb{P} \longmapsto^m \mathbb{P}')$*;*
– *for any m and* $\mathbb{P}' = (\mathbb{M}', [\mathbb{T}'_1, \ldots, \mathbb{T}'_n], \mathbb{L}')$*, if* $(\mathbb{P} \longmapsto^m \mathbb{P}')$*, then,*
  - $\Phi, [(\alpha'_1, \nu'_1), \ldots, (\alpha'_n, \nu'_n)] \vdash \mathbb{P}'$ *for some* $\alpha'_1, \ldots, \alpha'_n$ *and* $\nu'_1, \ldots, \nu'_n$*;*
  - *for any k, there exist* $\mathbb{M}''$*,* $\mathbb{T}''_k$ *and* $\mathbb{L}''$ *such that* $(\mathbb{M}', \mathbb{T}'_k, \mathbb{L}') \xmapsto{\mathsf{t}} (\mathbb{M}'', \mathbb{T}''_k, \mathbb{L}'')$*;*
  - *for any k, if* $\mathbb{T}'_k = (\mathbb{C}_k, \mathbb{R}'_k, \mathsf{j}\ f, k)$*, then* $(\alpha''_k * \nu''_k * \mathsf{True})\ (\mathbb{M}', (\mathbb{R}'_k, k), \mathbb{L}')$ *holds, where* $(\alpha''_k, \nu''_k) = \Psi_k(f)$*;*
  - *for any k, if* $\mathbb{T}'_k = (\mathbb{C}_k, \mathbb{R}'_k, \mathsf{jr}\ r_s, k)$*, then* $(\alpha''_k * \nu''_k * \mathsf{True})\ (\mathbb{M}', (\mathbb{R}'_k, k), \mathbb{L}')$ *holds, where* $(\alpha''_k, \nu''_k) = \Psi_k(\mathbb{R}'_k(r_s))$*;*
  - *for any k, if* $\mathbb{T}'_k = (\mathbb{C}_k, \mathbb{R}'_k, \mathsf{beq}\ r_s, r_t, f; \mathbb{I}, k)$ *and* $\mathbb{R}'_k(r_s) = \mathbb{R}'_k(r_t)$*, then* $(\alpha''_k * \nu''_k * \mathsf{True})\ (\mathbb{M}', (\mathbb{R}'_k, k), \mathbb{L}')$ *holds, where* $(\alpha''_k, \nu''_k) = \Psi_k(f)$*;*
  - *for any k, if* $\mathbb{T}'_k = (\mathbb{C}_k, \mathbb{R}'_k, \mathsf{bgt}\ r_s, r_t, f; \mathbb{I}, k)$ *and* $\mathbb{R}'_k(r_s) > \mathbb{R}'_k(r_t)$*, then* $(\alpha''_k * \nu''_k * \mathsf{True})\ (\mathbb{M}', (\mathbb{R}'_k, k), \mathbb{L}')$ *holds, where* $(\alpha''_k, \nu''_k) = \Psi_k(f)$*;*

Proof sketch: By Lemma 7, 12 and 10. □

$$
\begin{array}{rcl}
(\mathit{ProgSpec})\ \phi & ::= & ([\psi_1,\ldots,\psi_n],\Gamma) \\
(\mathit{CdHpSpec})\ \psi & ::= & \{\mathtt{f} \rightsquigarrow \nu\}^* \\
(\mathit{ResourceINV})\ \Gamma & \in & \mathit{Locks} \rightharpoonup \mathit{MemPred} \\
(\mathit{MemPred})\ \mathtt{m} & \in & \mathit{Memory} \rightarrow \mathit{Prop}
\end{array}
$$

Fig. 9. Specification Constructs for CSL

5 Concurrent Separation Logic (CSL)

Both AGL and SAGL treat lock/unlock primitives as normal instructions. They do not require that shared memory be protected by locks. This shows the generality of the A-G method, which makes no assumption about language constructs for synchronizations. Any ad-hoc synchronizations can be verified using the A-G method.

If we focus on a special class of programs following Hoare [6] where accesses of shared resources are protected by critical regions (implemented by locks in our language), we can further simplify our SAGL logic and derive a variation of CSL (CSL adapted to our assembly language).

5.1 CSL Specifications and Rules

In CSL, shared memory is partitioned and each part is protected by a unique lock. For each part of the partition, an invariant is assigned to specify its well-formedness. A thread cannot access shared memory unless it has acquired the corresponding lock. After the lock is acquired, the thread takes advantage of mutual-exclusion provided by locks and treats the part of memory as private. When the thread releases the lock, it must ensure that the part of memory is well-formed with regard to the corresponding invariant. In this way the following global invariant is enforced:

Shared resources are well-formed outside critical regions.

Figure 9 shows the specification constructs for CSL. The program specification ϕ contains a collection of code heap specifications for each thread and the specification Γ for lock-protected memory. Code heap specification ψ maps a code label to an assertion ν as the precondition of the corresponding instruction sequence. Here ν plays similar role of the private predicate in SAGL. Since each thread privately owns the lock protected memory if it owns the lock, all memory accessible by a thread is viewed as private memory. Therefore we do not need an assertion $\mathtt{a}$ to specify the shared memory as we did in SAGL. This also explains why we do not need assumptions and guarantees in CSL. The specification Γ of lock-protected memory maps a lock to an invariant $\mathtt{m}$, which specifies the corresponding part of memory. The invariant $\mathtt{m}$ is simply a predicate over memory because the register file is private to each thread.

Inference rules. The inference rules for CSL are presented in Fig. 11. The PROG rule requires that there be a partition of the global memory into $n+1$ parts. Each $\mathbb{M}_k$ is privately owned by thread $\mathbb{T}_k$. The well-formedness of $\mathbb{T}_k$ is checked by applying the

$$\mathtt{m} * \mathtt{m}' \stackrel{\text{def}}{=} \lambda\mathbb{M}.\exists\mathbb{M}_1,\mathbb{M}_2.\ (\mathbb{M}_1 \uplus \mathbb{M}_2 = \mathbb{M}) \wedge \mathtt{m}\ \mathbb{M}_1 \wedge \mathtt{m}'\ \mathbb{M}_2$$

$$\nu * \mathtt{m} \stackrel{\text{def}}{=} \lambda\mathbb{X}@(\mathbb{M},(\mathbb{R},k),\mathbb{L}).\exists\mathbb{M}_1,\mathbb{M}_2.\ (\mathbb{M}_1 \uplus \mathbb{M}_2 = \mathbb{M}) \wedge \nu\ (\mathbb{M}_1,(\mathbb{R},k),\mathbb{L}) \wedge \mathtt{m}\ \mathbb{M}_2$$

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

$$\mathsf{Precise}(\mathtt{m}) \stackrel{\text{def}}{=} \forall\mathbb{M},\mathbb{M}_1,\mathbb{M}_2.\ \mathbb{M}_1 \subseteq \mathbb{M} \wedge \mathbb{M}_2 \subseteq \mathbb{M} \wedge \mathtt{m}\ \mathbb{M}_1 \wedge \mathtt{m}\ \mathbb{M}_2 \rightarrow \mathbb{M}_1 = \mathbb{M}_2$$

$$\mathsf{Precise}(\Gamma) \stackrel{\text{def}}{=} \forall l \in dom(\Gamma).\ \mathsf{Precise}(\Gamma(l))$$

$$\nu \Rightarrow \nu' \stackrel{\text{def}}{=} \forall\mathbb{X}.\ \nu\ \mathbb{X} \rightarrow \nu'\ \mathbb{X}$$

$$\mathsf{acq}\ l\ \nu \stackrel{\text{def}}{=} \lambda(\mathbb{M},(\mathbb{R},k),\mathbb{L}).\ \nu\ (\mathbb{M},(\mathbb{R},k),\mathbb{L}\{l \rightsquigarrow k\})$$

$$\mathsf{rel}\ l\ \nu \stackrel{\text{def}}{=} \lambda(\mathbb{M},(\mathbb{R},k),\mathbb{L}).\ \mathbb{L}(l) = k \wedge \nu\ (\mathbb{M},(\mathbb{R},k),\mathbb{L} \setminus \{l\})$$

Fig. 10. Definitions of Notations in CSL

THRD rule. $\mathbb{M}_s$ is the part of memory protected by free locks (locks not owned by any threads). It must satisfy the invariants specified in Γ. Here a_Γ is the separating conjunction of invariants assigned to free locks in Γ, which is defined as:

$$\mathsf{a}_\Gamma \stackrel{\text{def}}{=} \lambda(\mathbb{M},(\mathbb{R},k),\mathbb{L}).\ (\forall_* l \in (dom(\Gamma) - dom(\mathbb{L})).\ \Gamma(l))\ \mathbb{M}, \tag{6}$$

that is, *shared resources are well-formed outside of critical regions*. Here $\forall_*$ is an indexed, finitely iterated separating conjunction, which is formalized in Fig. 10. Separating conjunctions with memory predicates ($\nu * \mathtt{m}$ and $\mathtt{m} * \mathtt{m}'$) are also defined in Fig. 10. As in O'Hearn's original work on CSL [11], we also require invariants specified in Γ to be precise (*i.e.,* $\mathsf{Precise}(\Gamma)$, as defined in Fig. 10), therefore we know a_Γ is also precise.

The THRD rule checks the well-formedness of threads. It requires that the current extended thread state satisfies the precondition ν. Since ν only cares about the resource privately owned by the thread, it takes $\mathbb{L}|_k$ instead of complete $\mathbb{L}$ as argument. Recall that $\mathbb{L}|_k$ is defined in (2) in Section 4 to represent the subset of $\mathbb{L}$ which maps locks to k. The CDHP rule and rules for instruction sequences are similar to their counterparts in AGL and SAGL and require no more explanation.

In the LOCK rule, we use "$\mathsf{acq}\ l\ \nu'$" to represent the weakest precondition of ν'; and "$\nu \Rightarrow \nu'$" for logical implication lifted for state predicates. They are formalized in Fig. 10. If the $\mathsf{lock}\ l$ instruction successfully acquires the lock l, we know by our global invariant that the part of memory protected by l satisfies the invariant $\Gamma(l)$ (*i.e.,* $\mathtt{m}$), because l is a free lock before $\mathsf{lock}\ l$ is executed. Therefore, we can carry the knowledge $\mathtt{m}$ in the postcondition ν'. Also, carrying $\mathtt{m}$ in ν' allows subsequent instructions to access that part of memory, since separation logic predicates capture ownerships of memory.

In the UNLOCK rule, "$\mathsf{rel}\ l\ \nu'$" is the weakest precondition for ν' (see Fig. 10). At the time the lock l is released, the memory protected by l must be well formed with respect to $\mathtt{m} = \Gamma(l)$. The separating conjunction here ensures that ν' does not specify this part of memory. Therefore the following instructions cannot use the part of memory unless the lock is acquired again. Rules for other instructions are straightforward and are not shown here.

Figure 12 shows admissible rules for CSL, including the frame rules and conjunction rules, which can be proved as lemmas in our meta-logic based on the rules shown

$\boxed{\phi, [\nu_1, \ldots, \nu_n] \vdash \mathbb{P}}$ (***Well-formed program***)

$$\frac{\begin{array}{ll}\phi = ([\psi_1, \ldots, \psi_n], \Gamma) & \mathbb{M}_s \uplus \mathbb{M}_1 \uplus \cdots \uplus \mathbb{M}_n = \mathbb{M} \\ \mathsf{a}_\Gamma\,(\mathbb{M}_s, _, \mathbb{L}) \quad \mathsf{Precise}(\Gamma) & \psi_k, \Gamma \vdash \{\nu_k\}\,(\mathbb{M}_k, \mathbb{T}_k, \mathbb{L}) \ \text{ for all } k\end{array}}{\phi, [\nu_1, \ldots, \nu_n] \vdash (\mathbb{M}, [\mathbb{T}_1, \ldots, \mathbb{T}_n], \mathbb{L})} \ (\text{PROG})$$

$\boxed{\psi, \Gamma \vdash \{\nu\}\,(\mathbb{M}, \mathbb{T}, \mathbb{L})}$ (***Well-formed thread***)

$$\frac{\nu\,(\mathbb{M}, (\mathbb{R}, k), \mathbb{L}|_k) \quad \psi, \Gamma \vdash \mathbb{C} : \psi \quad \psi, \Gamma \vdash \{\nu\}\,\mathbb{I}}{\psi, \Gamma \vdash \{\nu\}\,(\mathbb{M}, (\mathbb{C}, \mathbb{R}, \mathbb{I}, k), \mathbb{L})} \ (\text{THRD})$$

$\boxed{\psi, \Gamma \vdash \mathbb{C} : \psi'}$ (***Well-formed code heap***)

$$\frac{\forall \mathtt{f} \in dom(\psi') : \quad \psi, \Gamma \vdash \{\psi'(\mathtt{f})\}\,\mathbb{C}(\mathtt{f})}{\psi, \Gamma \vdash \mathbb{C} : \psi'} \ (\text{CDHP})$$

$\boxed{\psi, \Gamma \vdash \{\nu\}\,\mathbb{I}}$ (***Well-formed instr. sequences***)

$$\frac{\psi, \Gamma \vdash \{\nu\}\,\iota\,\{\nu'\} \quad \psi, \Gamma \vdash \{\nu'\}\,\mathbb{I}}{\psi, \Gamma \vdash \{\nu\}\,\iota; \mathbb{I}} \ (\text{SEQ}) \qquad \frac{\forall \mathbb{X} @ (\mathbb{M}, (\mathbb{R}, k), \mathbb{L}).\ \nu\,\mathbb{X} \rightarrow \psi(\mathbb{R}(\mathtt{r}_s))\,\mathbb{X}}{\psi, \Gamma \vdash \{\nu\}\,\mathsf{jr}\ \mathtt{r}_s} \ (\text{JR})$$

$\boxed{\psi, \Gamma \vdash \{\nu\}\,\iota\,\{\nu'\}}$ (***Well-formed instructions***)

$$\frac{\nu * \mathtt{m} \Rightarrow \mathsf{acq}\ l\ \nu'}{\psi, \Gamma\{l \rightsquigarrow \mathtt{m}\} \vdash \{\nu\}\,\mathsf{lock}\ l\,\{\nu'\}} \ (\text{LOCK}) \qquad \frac{\nu \Rightarrow (\mathsf{rel}\ l\ \nu') * \mathtt{m}}{\psi, \Gamma\{l \rightsquigarrow \mathtt{m}\} \vdash \{\nu\}\,\mathsf{unlock}\ l\,\{\nu'\}} \ (\text{UNLOCK})$$

Fig. 11. CSL Inference Rules

in Fig. 11. The frame rules (the FRAME-S rule and the FRAME-I rule) are very similar to the hypothetical frame rules presented in [12]. Interested readers can refer to previous papers on separation logic for their meanings.

5.2 Interpretation of CSL in SAGL

We prove the soundness of CSL by giving it an interpretation in SAGL, and proving CSL rules as derivable lemmas. This interpretation also formalizes the specialization made for CSL to achieve the simplicity.

From SAGL's point of view, each thread has two parts of memory: the private and the shared. In CSL, the private memory of a thread includes the memory protected by locks held by the thread and the memory that will never be shared. The shared memory are the parts protected by free locks. Therefore, we can use the following interpretation to translate a CSL specification to a SAGL specification:

$$[\![\nu]\!]_\Gamma \stackrel{\text{def}}{=} (\mathsf{a}_\Gamma, \nu) \tag{7}$$

$$[\![\psi]\!]_\Gamma \stackrel{\text{def}}{=} \lambda \mathtt{f}.[\![\psi(\mathtt{f})]\!]_\Gamma \ \text{ if } \mathtt{f} \in dom(\psi)\,, \tag{8}$$

$$\frac{\psi, \Gamma \vdash \{\nu\}\,\mathbb{I}}{\psi * \mathtt{m}, \Gamma \uplus \Gamma' \vdash \{\nu * \mathtt{m}\}\,\mathbb{I}}\ (\text{FRAME-S}) \qquad \frac{\psi, \Gamma \vdash \{\nu\}\,\iota\,\{\nu'\}}{\psi * \mathtt{m}, \Gamma \uplus \Gamma' \vdash \{\nu * \mathtt{m}\}\,\iota\,\{\nu' * \mathtt{m}\}}\ (\text{FRAME-I})$$

$$\text{where } \psi * \mathtt{m} \stackrel{\text{def}}{=} \lambda \mathtt{f}.\, \psi(\mathtt{f}) * \mathtt{m} \quad \text{if } \mathtt{f} \in dom(\psi)$$

$$\frac{\psi, \Gamma \vdash \{\nu\}\,\mathbb{I} \quad \psi, \Gamma \vdash \{\nu'\}\,\mathbb{I}}{\psi, \Gamma \vdash \{\nu \wedge \nu'\}\,\mathbb{I}}\ (\text{CONJ-S}) \qquad \frac{\psi, \Gamma \vdash \{\nu_1\}\,\iota\,\{\nu_1'\} \quad \psi, \Gamma \vdash \{\nu_2\}\,\iota\,\{\nu_2'\}}{\psi, \Gamma \vdash \{\nu_1 \wedge \nu_2\}\,\iota\,\{\nu_1' \wedge \nu_2'\}}\ (\text{CONJ-I})$$

Fig. 12. Admissible Rules for CSL

where a_Γ formalizes the CSL invariant and is defined by (6). We just reuse CSL specification ν as the specification of private memory, and use the separating conjunction a_Γ of invariants assigned to free locks as the specification for shared memory.

Since the assumption and guarantee in SAGL only specifies shared memory, we can define A_Γ and G_Γ for CSL threads:

$$\mathsf{A}_\Gamma \stackrel{\text{def}}{=} \lambda \mathbb{X}@(\mathbb{M}, (\mathbb{R}, k), \mathbb{L}), \mathbb{X}'@(\mathbb{M}', (\mathbb{R}', k'), \mathbb{L}').\mathbb{R} = \mathbb{R}' \wedge k = k' \wedge (\mathsf{a}_\Gamma\, \mathbb{X} \rightarrow \mathsf{a}_\Gamma\, \mathbb{X}') \quad (9)$$

$$\mathsf{G}_\Gamma \stackrel{\text{def}}{=} \lambda \mathbb{X}@(\mathbb{M}, (\mathbb{R}, k), \mathbb{L}), \mathbb{X}'@(\mathbb{M}', (\mathbb{R}', k'), \mathbb{L}').\, k = k' \wedge \mathsf{a}_\Gamma\, \mathbb{X} \wedge \mathsf{a}_\Gamma\, \mathbb{X}' \quad (10)$$

which enforces the invariant a_Γ of shared memory.

With above interpretations, we can prove the following soundness theorem.

Theorem 14 (CSL-Soundness).

1. *If* $\psi, \Gamma \vdash \{\nu\}\,\iota\,\{\nu'\}$ *in CSL, then* $[\![\psi]\!]_\Gamma, \mathsf{A}_\Gamma, \mathsf{G}_\Gamma \vdash \{[\![\nu]\!]_\Gamma\}\,\iota\,\{[\![\nu']\!]_\Gamma\}$ *in SAGL;*
2. *If* $\psi, \Gamma \vdash \{\nu\}\,\mathbb{I}$ *in CSL and* $\mathsf{Precise}(\Gamma)$*, then* $[\![\psi]\!]_\Gamma, \mathsf{A}_\Gamma, \mathsf{G}_\Gamma \vdash \{[\![\nu]\!]_\Gamma\}\,\mathbb{I}$ *in SAGL;*
3. *If* $\psi, \Gamma \vdash \mathbb{C} : \psi'$ *in CSL and* $\mathsf{Precise}(\Gamma)$*, then* $[\![\psi]\!]_\Gamma, \mathsf{A}_\Gamma, \mathsf{G}_\Gamma \vdash \mathbb{C} : [\![\psi']\!]_\Gamma$ *in SAGL;*
4. *If* $\psi, \Gamma \vdash \{\nu\}\,(\mathbb{M}_k, \mathbb{T}_k, \mathbb{L})$ *in CSL,* $\mathsf{Precise}(\Gamma)$*, and* $\mathsf{a}_\Gamma\,(\mathbb{M}_s, _, \mathbb{L})$*, then* $[\![\psi]\!]_\Gamma, \mathsf{A}_\Gamma, \mathsf{G}_\Gamma \vdash \{[\![\nu]\!]_\Gamma\}\,(\mathbb{M}_s, \mathbb{M}_k, \mathbb{T}_k, \mathbb{L})$ *in SAGL;*
5. *If* $([\psi_1, \ldots, \psi_n], \Gamma), [\nu_1, \ldots, \nu_n] \vdash \mathbb{P}$ *in CSL, then* $\Phi, [[\![\nu_1]\!]_\Gamma, \ldots, [\![\nu_n]\!]_\Gamma] \vdash \mathbb{P}$ *in SAGL, where* $\Phi = ([[\![\psi_1]\!]_\Gamma, \ldots, [\![\psi_n]\!]_\Gamma], [(\mathsf{A}_\Gamma, \mathsf{G}_\Gamma), \ldots, (\mathsf{A}_\Gamma, \mathsf{G}_\Gamma)])$.

6 SAGL Examples

We use two complementary examples to demonstrate how SAGL combines merits of AGL and CSL. Figure 13 shows a simple program, which allocates a fresh memory cell and then writes into and reads from it. Following the MIPS convention, we assume the register $\mathtt{r}_0$ always contains 0. The corresponding high-level pseudo code is given as comments (followed by "`;;`"). It is obvious that two threads executing the same code (but may use different *m*) will never interfere with each other, therefore the test in line (7) is always True and the program never reaches the unsafe branch.

It is trivial to certify the code in CSL since there is no memory-sharing at all. However, due to the nondeterministic operation of the alloc instruction, it is challenging to

```
(1)    start:  -{(emp, emp)}
(2)            addi  r1, r0, 1                  ;; local int x, y;
(3)            alloc r2, r1                     ;; x := alloc(1);
               -{(emp, r2 ↦ _)}
(4)            addi  r1, r0, m
(5)            st    r1, 0(r2)                  ;; [x] := m;
               -{(emp, (r2 ↦ m) ∧ r1 = m)}
(6)            ld    r3, 0(r2)                  ;; y := [x];
               -{(emp, (r2 ↦ m) ∧ r1 = m ∧ r3 = m)}
(7)            beq   r1, r3, safe               ;; while(y == m){}
(8)    unsafe: -{(emp, False)}
(9)            free  r0                         ;; free(0);  (* unsafe! *)
(10)   safe:   -{(emp, r2 ↦ _)}
(11)           j     safe
```

Fig. 13. Example 1: Memory Allocation

$(\mathtt{m} \mapsto \alpha) * (\mathtt{n} \mapsto \beta)$

```
local int x, y;                          local int x, y;
while(true){                             while(true){
   x := [m];                                x := [n];
   y := [n];                     ||         y := [m];
   if(x > y) {[m] := x-y;}                  if(x > y) {[n] := x-y;}
   if(x == y) { break;}                     if(x == y) { break;}
}                                        }
```

$(\mathtt{m} \mapsto gcd(\alpha,\beta)) * (\mathtt{n} \mapsto gcd(\alpha,\beta))$

Fig. 14. Example 2: Parallel GCD

certify the code in AGL because the specification of A and G requires global knowledge of memory. We certify the code in SAGL. Assertions are shown as annotations enclosed in "-{}". Recall that in SAGL the first assertion in the pair specifies shared resources and the second one specifies private resources. We treat all the resources as private, therefore the shared predicate is simply emp. The corresponding A and G are trivial. The whole verification is as simple as in CSL.

Our second example is adapted from Yu and Shao [20], which shows a parallel implementation of the Euclidean algorithm to compute the greatest common divisor (GCD) of α and β, stored at locations m and n initially. The high-level pseudo code is shown in Fig. 14. Memory cells at m and n are shared, but locks are not used for synchronization. To certify the code in CSL, we have to rewrite it by wrapping each memory-access command using "lock" and "unlock" commands and by introducing auxiliary variables. This time we use the "AGL part" of SAGL to certify the code. Figure 15 shows the assembly code of the first thread, with specifications as annotations. Private predicates are simply emp. The assumption and guarantee are defined below, where we

```
loop: -{(∃x,y. (m ↦ x) * (n ↦ y) ∧ gcd(x,y) = gcd(α,β), emp)}
      ld  r1, m(r0)                                  ;; r1 <- [m]
      -{(∃x,y. (m ↦ x) * (n ↦ y) ∧ gcd(x,y) = gcd(α,β) ∧ r1 = x, emp)}
      ld  r2, n(r0)                                  ;; r2 <- [n]
      -{(∃x,y. (m ↦ x) * (n ↦ y) ∧ gcd(x,y) = gcd(α,β) ∧ r1 = x
                                  ∧r2 ≥ y ∧ (x ≥ y → r2 = y), emp)}
      bgt r1, r2, calc                               ;; if (r1 > r2)  goto calc
      beq r1, r2, done                               ;; if (r1 == r2) goto done
      j   loop                                       ;; goto loop
calc: -{(∃x,y. (m ↦ x) * (n ↦ y) ∧ gcd(x,y) = gcd(α,β) ∧ r1 = x ∧ r2 = y ∧ x > y, emp)}
      sub r3, r1, r2                                 ;; r3 = r1 - r2
      st  r3, m(r0)                                  ;; [m] <- r3
      j   loop                                       ;; goto loop
done: -{(∃x. (m ↦ x) * (n ↦ x) ∧ x = gcd(α,β), emp)}
      j   done
```

Fig. 15. Parallel GCD–Assembly Code for The First Thread

use primed values (*e.g.,* $[\mathtt{m}]'$ and $[\mathtt{n}]'$) to represent memory values in the resulting state of each action.

$$\mathsf{A}_1 \stackrel{\text{def}}{=} ([\mathtt{m}] = [\mathtt{m}]') \wedge ([\mathtt{n}] \geq [\mathtt{n}]') \wedge ([\mathtt{m}] \geq [\mathtt{n}] \rightarrow [\mathtt{n}] = [\mathtt{n}]') \wedge (gcd([\mathtt{m}],[\mathtt{n}]) = gcd([\mathtt{m}]',[\mathtt{n}]'))$$
$$\mathsf{G}_1 \stackrel{\text{def}}{=} ([\mathtt{n}] = [\mathtt{n}]') \wedge ([\mathtt{m}] \geq [\mathtt{m}]') \wedge ([\mathtt{n}] \geq [\mathtt{m}] \rightarrow [\mathtt{m}] = [\mathtt{m}]') \wedge (gcd([\mathtt{m}],[\mathtt{n}]) = gcd([\mathtt{m}]',[\mathtt{n}]'))$$

The example shown in Fig. 16 is adapted from O'Hearn [10]. P-V primitives are firstly implemented using locks, then they are used for synchronization. This example illustrates the support of redistribution of shared and private memory in SAGL.

7 Related Work and Conclusion

O'Hearn [11] proposed CSL for a high-level parallel language following Hoare [6]. Synchronization in the language is achieved by the conditional critical region (CCR) in the form of "with r when b do c". Semantics of CCRs is as follows: the statement c can be executed only if the resource r has not been acquired by others and the Boolean expression b is true; otherwise the thread will be blocked. We adapt CSL to an assembly language. The CCR can be implemented using our lock/unlock primitives. Each lock in our language corresponds to a resource name at the high-level. Atomic instructions in our assembly language are very similar to actions in Brookes Semantics [2], where semantic functions are defined for statements and expressions. These semantic functions can be viewed as a translation from the high-level language to a low-level language similar to ours. Recently, Reynolds [16] and Brookes [3] have studied grainless semantics for concurrency. Brookes also gives a grainless semantics to CSL [3].

The PROG rule of our CSL corresponds to O'Hearn's parallel composition rule [11]. The number of threads in our machine is fixed, therefore the nested parallel composition statement supported by Brookes [2] is not supported in our language. We studied verification of assembly code with dynamic thread creation in an earlier paper [5].

$$
\begin{aligned}
\mathsf{I}(\mathtt{s},x) &\stackrel{\text{def}}{=} (\mathtt{s} \mapsto x) * ((x = 0 \wedge \mathsf{emp}) \vee (x = 1 \wedge 10 \mapsto _)) \\
\mathsf{I}(\mathtt{s}) &\stackrel{\text{def}}{=} \exists x.\mathsf{I}((\mathtt{s},x)) \\
\Gamma &\stackrel{\text{def}}{=} \{l_1 \rightsquigarrow \mathsf{I}(\mathtt{free}), l_2 \rightsquigarrow \mathsf{I}(\mathtt{busy})\} \\
\mathsf{own}(l) &\stackrel{\text{def}}{=} \mathbb{L}(l) = \mathsf{self}
\end{aligned}
$$

```
P_free: -{(a_Γ, emp)}                                  ;; while(true){
        lock   l_1                                     ;;   lock l_1
        -{(a_Γ, I(free))}
        ld     r1, free(r0)                            ;;
        -{(a_Γ, I(free,r_1) ∧ own(l_1))}
        bgt    r1, r0, dec_P                           ;;   if([free]>0) break;
        -{(a_Γ, r_1 = 0 ∧ I(free,r_1) ∧ own(l_1))}
        unlock l_1                                     ;;   unlock l_1
        -{(a_Γ, emp)}
        j      P_free                                  ;; }
dec_P:  -{(a_Γ, r_1 = 1 ∧ I(free,r_1) ∧ own(l_1))}
        subi   r1, r1, 1                               ;; [free] <- [free]-1
        -{(a_Γ, r_1 = 0 ∧ I(free,1) ∧ own(l_1))}
        st     r1, free(r0)
        -{(a_Γ, (10 ↦ _) * I(free,0) ∧ own(l_1))}
        -{(a_Γ, (10 ↦ _) ∧ own(l_1))}
        unlock l_1                                     ;;  unlock l_1
        -{(a_Γ, 10 ↦ _)}
        j      body

body:   -{(a_Γ, 10 ↦ _)}
        addi   r2, r0, m
        st     r2, 10(r0)                              ;;  [10] <- m
        j      V_busy

V_busy: -{(a_Γ, 10 ↦ _)}
        lock   l_2                                     ;; lock l_2
        -{(a_Γ, (10 ↦ _) ∧ own(l_2))}
        -{(a_Γ, (10 ↦ _) * I(busy,0) ∧ own(l_2))}
        ld     r1, busy(r0)
        addi   r1, r0, 1                               ;;  [busy] <- [busy]+1
        st     r1, busy(r0)
        -{(a_Γ, I(busy,1) ∧ own(l_2))}
        unlock l_2                                     ;; unlock l_2
        -{(a_Γ, emp)}
        j      done
done:   -{(a_Γ, emp)}
        j      done
```

Fig. 16. SAGL Example: Synchronizations Based on P-V Operations

CSL is still evolving. Bornat *et al.* [1] proposed a refinement of CSL with fine-grained resource accounting. Parkinson *et al.* [14] applied CSL to verify a non-blocking implementation of stacks. As in the original CSL, these works also assume language constructs for synchronizations. We suspect that there exist reductions from these variations to SAGL-like logics. We leave this as our future work.

Concurrently with our work on SAGL, Vafeiadis and Parkinson [18] proposed another approach to combining rely/guarantee and separation logic, which we refer to here as RGSep. Both RGSep and SAGL partition memory into shared and private parts. However, shared memory cannot be accessed directly in RGSep. It has to be converted into private first to be accessed. Conversions can only occur at boundaries of critical regions, which is a built-in language construct required by RGSep to achieve atomicity. RGSep, in principle, does not assume smallest granularity of transitions. In SAGL, shared memory can be accessed directly, or be converted into private first and then accessed. Conversions can be made dynamically at any program point, instead of being coupled with critical regions. However, like A-G reasoning, SAGL assumes smallest granularity. We suspect that RGSep can be compiled into a specialized version of SAGL, following the way we translate CSL. On the other hand, if our instructions are wrapped using critical regions, SAGL might be derived from RGSep too.

We also use SAGL as the basis to formalize the relationship between CSL and A-G reasoning. We encode the CSL invariant as an assumption and guarantee in SAGL, and prove that CSL rules are derivable from corresponding SAGL rules with the specific assumption and guarantee. Soundness of SAGL is proved following the syntactic approach to type soundness. Our work has been formalized in Coq [4].

References

[1] R. Bornat, C. Calcagno, P. O'Hearn, and M. Parkinson. Permission accounting in separation logic. In *Proc. 32nd ACM Symp. on Principles of Prog. Lang.*, pages 259–270, 2005.

[2] S. Brookes. A semantics for concurrent separation logic. In *Proc. 15th International Conference on Concurrency Theory (CONCUR'04)*, volume 3170 of *LNCS*, pages 16–34, 2004.

[3] S. Brookes. A grainless semantics for parallel programs with shared mutable data. In *Proc. MFPS XXI*, volume 155 of *Electr. Notes Theor. Comput. Sci.*, pages 277–307, 2006.

[4] X. Feng, R. Ferreira, and Z. Shao. On the relationship between concurrent separation logic and assume-guarantee reasoning: Formalization in coq. http://flint.cs.yale.edu/publications/sagl.html, October 2006.

[5] X. Feng and Z. Shao. Modular verification of concurrent assembly code with dynamic thread creation and termination. In *Proc. ICFP'05*, pages 254–267, 2005.

[6] C. A. R. Hoare. Towards a theory of parallel programming. In C. A. R. Hoare and R. H. Perrott, editors, *Operating Systems Techniques*, pages 61–71. Academic Press, 1972.

[7] 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.*, pages 14–26, 2001.

[8] C. B. Jones. Tentative steps toward a development method for interfering programs. *ACM Trans. on Programming Languages and Systems*, 5(4):596–619, 1983.

[9] G. Necula. Proof-carrying code. In *Proc. 24th ACM Symp. on Principles of Prog. Lang.*, pages 106–119. ACM Press, Jan. 1997.

[10] P. W. O'Hearn. Resources, concurrency and local reasoning. *Theoretical Computer Science (to appear)*. Journal version of [11].

[11] P. W. O'Hearn. Resources, concurrency and local reasoning. In *Proc. 15th Int'l Conf. on Concurrency Theory (CONCUR'04)*, volume 3170 of *LNCS*, pages 49–67, 2004.

[12] P. W. O'Hearn, H. Yang, and J. C. Reynolds. Separation and information hiding. In *Proc. 31th ACM Symp. on Principles of Prog. Lang.*, pages 268–280, Venice, Italy, Jan. 2004.

[13] S. Owicki and D. Gries. Verifying properties of parallel programs: an axiomatic approach. *Commun. ACM*, 19(5):279–285, 1976.

[14] M. Parkinson, R. Bornat, and P. O'Hearn. Modular verification of a non-blocking stack. In *Proc. 34th ACM Symp. on Principles of Prog. Lang.*, page to appear. ACM Press, Jan. 2007.

[15] J. C. Reynolds. Separation logic: A logic for shared mutable data structures. In *Proc. LICS'02*, pages 55–74, July 2002.

[16] J. C. Reynolds. Toward a grainless semantics for shared-variable concurrency. In *Proc. FSTTCS'04*, volume 3328 of *LNCS*, pages 35–48, 2004.

[17] The Coq Development Team. The Coq proof assistant reference manual. The Coq release v8.0, Oct. 2004.

[18] V. Vafeiadis and M. Parkinson. A marriage of rely/guarantee and separation logic. Available at `http://www.cl.cam.ac.uk/~mjp41/RGSep.pdf`, 2007.

[19] A. K. Wright and M. Felleisen. A syntactic approach to type soundness. *Information and Computation*, 115(1):38–94, 1994.

[20] D. Yu and Z. Shao. Verification of safety properties for concurrent assembly code. In *Proc. 2004 ACM SIGPLAN Int'l Conf. on Functional Prog.*, pages 175–188, September 2004.