Rely-Guarantee-Based Simulation for Compositional Verification of Concurrent Program Transformations

HONGJIN LIANG, XINYU FENG, and MING FU, University of Science and Technology of China

Verifying program transformations usually requires proving that the resulting program (the target) refines or is equivalent to the original one (the source). However, the refinement relation between individual sequential threads cannot be preserved in general with the presence of parallel compositions, due to instruction reordering and the different granularities of atomic operations at the source and the target. On the other hand, the refinement relation defined based on fully abstract semantics of concurrent programs assumes arbitrary parallel environments, which is too strong and cannot be satisfied by many well-known transformations.

In this article, we propose a *Rely-Guarantee-based Simulation* (RGSim) to verify concurrent program transformations. The relation is parametrized with constraints of the environments that the source and the target programs may compose with. It considers the interference between threads and their environments, thus is less permissive than relations over sequential programs. It is compositional with respect to parallel compositions as long as the constraints are satisfied. Also, RGSim does not require semantics preservation under all environments, and can incorporate the assumptions about environments made by specific program transformations in the form of rely/guarantee conditions. We use RGSim to reason about optimizations and prove atomicity of concurrent objects. We also propose a general garbage collector verification framework based on RGSim, and verify the Boehm et al. concurrent mark-sweep GC.

Categories and Subject Descriptors: D.2.4 [Software Engineering]: Software/Program Verification— *Correctness proofs, Formal methods*; F.3.1 [Logics and Meanings of Programs]: Specifying and Verifying and Reasoning about Programs

General Terms: Theory, Verification

Additional Key Words and Phrases: Concurrency, program transformation, rely-guarantee reasoning, simulation

ACM Reference Format:

Liang, H., Feng, X., and Fu, M. 2014. Rely-guarantee-based simulation for compositional verification of concurrent program transformations. *ACM Trans. Program. Lang. Syst.* 36, 1, Article 3 (March 2014), 55 pages. DOI:http://dx.doi.org/10.1145/2576235

1. INTRODUCTION

Many verification problems can be reduced to verifying program transformations, that is, proving the target program of the transformation has no more observable

© 2014 ACM 0164-0925/2014/03-ART3 \$15.00

DOI:http://dx.doi.org/10.1145/2576235

This work is supported in part by grants from National Natural Science Foundation of China (NSFC) under grant nos. 61379039, 61229201, 61103023, and 91318301, Program for New Century Excellent Talents in Universities (grant no. NCET-2010-0984), and the Fundamental Research Funds for the Central Universities (grant no. WK0110000031).

Authors' address: H. Liang, X. Feng (corresponding author), and M. Fu, School of Computer Science and Technology, University of Science and Technology of China, Hefei, Anhui 230026, China and Suzhou Institute for Advanced Study, University of Science and Technology of China, Suzhou, China 215123; email: xyfeng@ustc.edu.cn.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

behaviors than the source. Next we give some typical examples in concurrent settings.

- Correctness of compilation and optimizations of concurrent programs. In this most natural program transformation verification problem, every compilation phase does a program transformation **T**, which needs to preserve the semantics of the inputs.
- —*Atomicity of concurrent objects.* A concurrent object or library provides a set of methods that allow clients to manipulate the shared data structure with abstract atomic behaviors [Herlihy and Shavit 2008]. Their correctness can be reduced to the correctness of the transformation from abstract atomic operations to concrete and executable programs in a concurrent context.
- Verifying implementations of Software Transactional Memory (STM). Many languages supporting STM provide a high-level atomic block **atomic**{ \mathbb{C} }, so that programmers can assume the atomicity of the execution of \mathbb{C} . Atomic blocks are implemented using some STM protocol (e.g., TL2 [Dice et al. 2006]) that allows very fine-grained interleavings. Verifying that the fine-grained program respects the semantics of atomic blocks gives us the correctness of the STM implementation.
- *Correctness of concurrent Garbage Collectors (GCs).* High-level garbage-collected languages (e.g., Java) allow programmers to work at an abstract level without knowledge of the underlying GC algorithm. However, the concrete and executable low-level program involves interactions between the mutators and the collector. If we view the GC implementation as a transformation from high-level mutators to low-level ones with a concrete GC thread, the GC safety can be reduced naturally to the semantics preservation of the transformation.

To verify the correctness of a program transformation \mathbf{T} , we follow Leroy's approach [Leroy 2009] and define a refinement relation \sqsubseteq between the target and the source programs, which says the target has no more observable behaviors than the source. Then we can formalize the correctness of the transformation as follows.

$$\mathsf{Correct}(\mathbf{T}) \triangleq \forall C, \mathbb{C}. \ C = \mathbf{T}(\mathbb{C}) \implies C \sqsubseteq \mathbb{C}.$$
(1.1)

That is, for any source program \mathbb{C} acceptable by $\mathbf{T}, \mathbf{T}(\mathbb{C})$ is a refinement of \mathbb{C} . When the source and the target are shared-state concurrent programs, the refinement \sqsubseteq needs to satisfy the following requirements to support effective proof of $Correct(\mathbf{T})$.

- Since the target $\mathbf{T}(\mathbb{C})$ may be in a different language from the source, the refinement should be general and independent of the language details.
- To verify fine-grained implementations of abstract operations, the refinement should support different views of program states and different granularities of state accesses at the source and the target levels.
- When **T** is syntax-directed (and it is usually the case for parallel compositions, i.e., $\mathbf{T}(\mathbb{C} \parallel \mathbb{C}') = \mathbf{T}(\mathbb{C}) \parallel \mathbf{T}(\mathbb{C}')$), a *compositional* refinement is of particular importance for modular verification of **T**.

However, existing refinement (or equivalence) relations cannot satisfy all these requirements at the same time. Contextual equivalence, the canonical notion for comparing program behaviors, fails to handle different languages since the contexts of the source and the target will be different. Simulations and logical relations have been used to verify compilation [Benton and Hur 2009; Hur and Dreyer 2011; Leroy 2009; Lochbihler 2010], but they are usually designed for sequential programs (except Lochbihler [2010] and Ševčík et al. [2011], which we will discuss in Section 8). Since the refinement or equivalence relation between sequential threads cannot be preserved in general with parallel compositions, we cannot simply adapt existing work on sequential programs to verify transformations of concurrent programs. Refinement relations based on fully abstract semantics of concurrent programs are compositional, but they assume arbitrary program contexts, which is too strong for many practical transformations. We will explain the challenges in detail in Section 2.

In this article, we propose a *Rely-Guarantee-based Simulation* (RGSim) for compositional verification of concurrent transformations. By addressing the preceding problems, we make the following contributions.

- RGSim parametrizes the simulation between concurrent programs with rely/ guarantee conditions [Jones 1983] which specify the interactions between the programs and their environments. This makes the corresponding refinement relation compositional with respect to parallel compositions, allowing us to decompose refinement proofs for multithreaded programs into proofs for individual threads. On the other hand, the rely/guarantee conditions can incorporate the assumptions about environments made by specific program transformations, so RGSim can be applied to verify many practical transformations.
- Based on the simulation technique, RGSim focuses on comparing externally observable behaviors (e.g., I/O events) only, which gives us considerable leeway in the implementations of related programs. The relation is mostly independent of the language details. It can be used to relate programs in different languages with different views of program states and different granularities of atomic state accesses.
- RGSim makes relational reasoning about optimizations possible in parallel contexts. We present a set of relational reasoning rules to characterize and justify common optimizations in a concurrent setting, including hoisting loop invariants, strength reduction and induction variable elimination, dead code elimination, redundancy introduction, etc.
- RGSim gives us a refinement-based proof method to verify fine-grained implementations of abstract algorithms and concurrent objects. We successfully apply RGSim to verify concurrent counters, the concurrent GCD algorithm, Treiber's nonblocking stack, and the lock-coupling list.
- We reduce the problem of verifying concurrent garbage collectors to verifying transformations, and present a general GC verification framework which combines unary rely-guarantee-based verification [Jones 1983] with relational proofs based on RGSim.
- We verify the Boehm et al. concurrent garbage collection algorithm [Boehm et al. 1991] using our framework. As far as we know, it is the first time to formally prove the correctness of this algorithm.
- We give a mechanized formulation of RGSim, and prove its soundness and compositionality in the Coq proof assistant [2010]. Both the manual and mechanized proofs are available online¹.

This article extends the conference paper in POPL 2012 [Liang et al. 2012]. First, we add more examples, including strength reduction and induction variable elimination, the nonblocking concurrent counter, Treiber's stack algorithm, and the concurrent GCD algorithm. Second, we significantly expand the details for the concurrent GC verification, demonstrating that RGSim is a powerful proof technique for verifying program transformations which involve concurrent runtime systems.

In the rest of this article, we first analyze the challenges for compositional verification of concurrent program transformations, and explain our approach informally in Section 2. Then we give the basic technical settings in Section 3 and present the formal definition of RGSim in Section 4. We show the use of RGSim to reason about

¹http://kyhcs.ustcsz.edu.cn/relconcur/rgsim

ACM Transactions on Programming Languages and Systems, Vol. 36, No. 1, Article 3, Publication date: March 2014.

local r1; local r2; x := 1; y := 1; r1 := y; || r2 := x; if (r1 = 0) then if (r2 = 0) then critical region critical region

(a) Dekker's mutual exclusion algorithm

(b) different granularities of atomic operations

Fig. 1. Equivalence lost after parallel composition.

optimizations in Section 5, verify fine-grained algorithms and atomicity of concurrent objects in Section 6, and prove the correctness of concurrent GCs in Section 7. Finally we discuss related work and conclude in Section 8.

2. CHALLENGES AND OUR APPROACH

The major challenge we face is to have a compositional refinement relation \sqsubseteq between concurrent programs, that is, we should be able to know $\mathbf{T}(\mathbb{C}_1) \| \mathbf{T}(\mathbb{C}_2) \sqsubseteq \mathbb{C}_1 \| \mathbb{C}_2$ if we have $\mathbf{T}(\mathbb{C}_1) \sqsubseteq \mathbb{C}_1$ and $\mathbf{T}(\mathbb{C}_2) \sqsubseteq \mathbb{C}_2$.

2.1. Sequential Refinement Loses Parallel Compositionality

Observable behaviors of sequential imperative programs usually refer to their control effects (e.g., termination and exceptions) and final program states. However, refinement relations defined correspondingly cannot be preserved after parallel compositions. It has been a well-known fact in the compiler community that sound optimizations for sequential programs may change the behaviors of multithreaded programs [Boehm 2005]. The Dekker's algorithm shown in Figure 1(a) has been widely used to demonstrate the problem. Reordering the first two assignment statements of the thread on the left preserves its sequential behaviors, but the whole program can no longer ensure exclusive access to the critical region.

In addition to instruction reordering, the different granularities of atomic operations between the source and the target programs can also break the compositionality of program equivalence in a concurrent setting. In Figure 1(b), the target program at the bottom behaves differently from the source at the top (assuming each statement is executed atomically), although the individual threads at the target and the source have the same behaviors.

2.2. Assuming Arbitrary Environments is Too Strong

The problem with the refinement for sequential programs is that it does not consider the effects of threads' intermediate state accesses on their parallel environments. People have given fully abstract semantics to concurrent programs (e.g., [Abadi and Plotkin 2009; Brookes 1996]). The semantics of a program is modeled as a set of execution traces. Each trace is an interleaving of state transitions made by the program itself and *arbitrary* transitions made by the environment. Then the refinement between programs can be defined as the subset relation between the corresponding trace sets. Since it considers all possible environments, the refinement relation has very nice compositionality, but unfortunately is too strong to formulate the correctness of many well-known transformations, including the four classes of transformations mentioned before.

- Many concurrent languages (e.g., C++ [Boehm and Adve 2008]) do not give semantics to programs with data races (like the examples shown in Figure 1). Therefore the compilers only need to guarantee the semantics preservation of data-race-free programs.
- When we prove that a fine-grained implementation of a concurrent object is a refinement of an abstract atomic object, we can assume that all accesses to the object are made through the object's methods only, for example, a stack object can only be accessed through push and pop methods, and its internal data cannot be arbitrarily updated.
- Usually the implementation of STM (e.g., TL2 [Dice et al. 2006]) ensures the atomicity of a transaction **atomic**{ \mathbb{C} } only when there are no data races. Therefore, the correctness of the transformation from high-level atomic blocks to fine-grained concurrent code assumes data-race freedom in the source.
- Many garbage-collected languages are type-safe and prohibit operations such as pointer arithmetic. Therefore the garbage collector could make corresponding assumptions about the mutators that run in parallel.

In all these cases, the transformations of individual threads are allowed to make various assumptions about the environments. They do not have to ensure semantics preservation within all contexts.

2.3. Languages at Source and Target May Be Different

The use of different languages at the source and the target levels makes the formulation of the transformation correctness more difficult. If the source and the target languages have different views of program states and different atomic primitives, we cannot directly compare the state transitions made by the source and the target programs. This is another reason that makes the aforementioned subset relation between sets of program traces in fully abstract semantics infeasible. For the same reason, many existing techniques for proving refinement or equivalence of programs in the same language cannot be applied either.

2.4. Different Observers Make Different Observations

Concurrency introduces tensions between two kinds of observers: human beings (as external observers) and the parallel program contexts. External observers do not care about the implementation details of the source and the target programs. For them, intermediate state accesses (such as memory reads and writes) are silent steps (unobservable), and only external events (such as I/O operations) are observable. On the other hand, state accesses have effects on the parallel program contexts, and are not silent to them.

If the refinement relation relates externally observable event traces only, it cannot have parallel compositionality, as we explained in Section 2.1. On the other hand, relating all state accesses of programs is too strong. Any reordering of state accesses or change of atomicity would fail the refinement.

2.5. Our Approach

In this article we propose a *Rely-Guarantee-based Simulation* (RGSim) \leq between the target and the source programs. It establishes a weak simulation, ensuring that for every externally observable event made by the target program there is a corresponding one in the source. We choose to view intermediate state accesses as silent steps, thus we can relate programs with different implementation details. This also makes our simulation independent of language details.

To support parallel compositionality, our relation takes into account explicitly the expected interference between threads and their parallel environments. Inspired by the rely-guarantee (R-G) verification method [Jones 1983], we specify the interference using rely/guarantee conditions. In rely-guarantee reasoning, the rely condition R of a thread specifies the permitted state transitions that its environment may have, and its guarantee G specifies the possible transitions made by the thread itself. To ensure parallel threads can collaborate, we need to check the interference constraint, that is, the guarantee of each thread is permitted in the rely of every other. Then we can verify their parallel composition by separately verifying each thread, showing its behaviors under the rely condition indeed satisfy its guarantee. After parallel composition, the threads should be executed under their common environment (i.e., the intersection of their relies) and guarantee all the possible transitions made by them (i.e., the union of their guarantees).

Parametrized with rely/guarantee conditions for the two levels, our relation $(C, \mathcal{R}, \mathcal{G}) \leq (\mathbb{C}, \mathbb{R}, \mathbb{G})$ talks about not only the target C and the source \mathbb{C} , but also the interference \mathcal{R} and \mathcal{G} between C and its target-level environment, and \mathbb{R} and \mathbb{G} between \mathbb{C} and its environment at the source level. Informally, $(C, \mathcal{R}, \mathcal{G}) \leq (\mathbb{C}, \mathbb{R}, \mathbb{G})$ says the executions of C under the environment \mathcal{R} do not exhibit more observable behaviors than the executions of \mathbb{C} under the environment \mathbb{R} , and the state transitions of C and \mathbb{C} satisfy \mathcal{G} and \mathbb{G} respectively. RGSim is now compositional, as long as the threads are composed with well-behaved environments only. The parallel compositionality lemma is in the following form. If we know $(C_1, \mathcal{R}_1, \mathcal{G}_1) \leq (\mathbb{C}_1, \mathbb{R}_1, \mathbb{G}_1)$ and $(C_2, \mathcal{R}_2, \mathcal{G}_2) \leq (\mathbb{C}_2, \mathbb{R}_2, \mathbb{G}_2)$, and also the interference constraints are satisfied, that is, $\mathcal{G}_2 \subseteq \mathcal{R}_1$, $\mathcal{G}_1 \subseteq \mathcal{R}_2$, $\mathbb{G}_2 \subseteq \mathbb{R}_1$ and $\mathbb{G}_1 \subseteq \mathbb{R}_2$, we could get

$$(C_1 || C_2, \mathcal{R}_1 \cap \mathcal{R}_2, \mathcal{G}_1 \cup \mathcal{G}_2) \preceq (\mathbb{C}_1 || \mathbb{C}_2, \mathbb{R}_1 \cap \mathbb{R}_2, \mathbb{G}_1 \cup \mathbb{G}_2).$$

The compositionality of RGSim gives us a proof theory for concurrent program transformations.

Also different from fully abstract semantics for threads, which assumes arbitrary behaviors of environments, RGSim allows us to instantiate the interference \mathcal{R} , \mathcal{G} , \mathbb{R} and \mathbb{G} differently for different assumptions about environments, therefore it can be used to verify the aforementioned four classes of transformations. For instance, if we want to prove that a transformation preserves the behaviors of data-race-free programs, we can specify the data-race freedom in \mathbb{R} and \mathbb{G} . Then we are no longer concerned with the examples in Figure 1, both of which have data races.

Example. Next we give an example of loop invariant hoisting to illustrate how RGSim works. The formal proofs are shown in Section 5.2.1.

Target Code (C_1)	Source Code (C)
local t;	local t;
t := x + 1;	while(i < n) {
while(i < n) {	← t := x + 1;
i := i + t;	i := i + t;
}	}

(Events)
$$e ::= \dots$$
 (Labels) $o ::= e \mid \tau$

(a) events and transition labels

 $\begin{array}{rcl} (LState) & \sigma & ::= & \dots \\ (LExpr) & E & \in & LState \rightarrow Int_{\perp} \\ (LBExp) & B & \in & LState \rightarrow \{ \mathbf{true}, \mathbf{false} \}_{\perp} \\ (LInstr) & c & \in & LState \rightarrow \mathcal{P}((Labels \times LState) \cup \{ \mathbf{abort} \}) \\ (LStmt) & C & ::= & \mathbf{skip} \mid c \mid C_1; C_2 \mid \mathbf{if} \ (B) \ C_1 \ \mathbf{else} \ C_2 \mid \mathbf{while} \ (B) \ C \mid C_1 \parallel C_2 \\ (LStep) \longrightarrow_L \in & \mathcal{P}((LStmt \setminus \{ \mathbf{skip} \} \times LState) \times Labels \times ((LStmt \times LState) \cup \{ \mathbf{abort} \})) \end{array}$

(b) the low-level language

(c) the high-level language

Fig. 2. Generic languages at target and source levels.

Benton [2004] has proved that the optimized code C_1 preserves the *sequential* behaviors of the source C. In a concurrent setting, this optimization is incorrect within arbitrary environments. For instance, if other threads may update x, the final values of i might be different at the two levels. In fact, this optimization works only when the environments \mathcal{R} at both levels do not update x nor t. The guarantees \mathcal{G} of both C_1 and C can be specified as arbitrary transitions. Then we can prove the RGSim relation $(C_1, \mathcal{R}, \mathcal{G}) \leq (C, \mathcal{R}, \mathcal{G})$ and conclude the correctness of the transformation.

3. BASIC TECHNICAL SETTINGS

In this section, we present the source and the target programming languages. Then we define a basic refinement \sqsubseteq , which naturally says the target has no more externally observable event traces than the source. We use \sqsubseteq as an intuitive formulation of the correctness of transformations. Our RGSim relation, which will be formally defined in Section 4, is proposed as a proof technique for \sqsubseteq .

3.1. The Languages

Following standard simulation techniques, we model the semantics of target and source programs as labeled transition systems. Before showing the languages, we first define events and labels in Figure 2(a). We leave the set of events unspecified here. It

$$\frac{(o, \Sigma') \in c \Sigma}{(c, \Sigma) \xrightarrow{o} (skip, \Sigma')} \qquad \frac{abort \in c \Sigma}{(c, \Sigma) \rightarrow abort} \qquad \frac{\Sigma \notin dom(c)}{(c, \Sigma) \rightarrow (c, \Sigma)}$$

$$\frac{(\mathbb{C}_1, \Sigma) \xrightarrow{o} (Skip, \Sigma) \rightarrow (skip, \Sigma)}{(\mathbb{C}_1 ||| \mathbb{C}_2, \Sigma) \xrightarrow{o} (\mathbb{C}'_1, \Sigma')} \qquad \frac{(\mathbb{C}_1, \Sigma) \xrightarrow{o} (\mathbb{C}'_1, \Sigma')}{(\mathbb{C}_1 ||| \mathbb{C}_2, \Sigma) \xrightarrow{o} (\mathbb{C}'_1 ||| \mathbb{C}_2, \Sigma')}$$

$$\frac{(\mathbb{C}_2, \Sigma) \xrightarrow{o} (\mathbb{C}'_2, \Sigma')}{(\mathbb{C}_1 ||| \mathbb{C}'_2, \Sigma')} \qquad \frac{(\mathbb{C}_1, \Sigma) \rightarrow abort \text{ or } (\mathbb{C}_2, \Sigma) \rightarrow abort}{(\mathbb{C}_1 ||| \mathbb{C}_2, \Sigma) \rightarrow abort}$$

Fig. 3. Selected operational semantics rules of the high-level language.

can be instantiated by program verifiers, depending on their interest (e.g., input/output events). A label that will be associated with a state transition is either an event or τ , which means the corresponding transition does not generate any event (i.e., a silent step).

The target language, which we also call the low-level language, is shown in Figure 2(b). We abstract away the forms of states, expressions, and primitive instructions in the language. An arithmetic expression E is modeled as a function from states to integers lifted with an undefined value \bot . Boolean expressions Bs are modeled similarly. An instruction c is a partial function from states to sets of label and state pairs, describing the state transitions and the events it generates. We use $\mathcal{P}(_)$ to denote the power set. Unsafe executions lead to **abort**. Note that the semantics of an instruction could be nondeterministic. Moreover, it might be undefined on some states, making it possible to model blocking operations such as acquiring a lock.

Statements are either primitive instructions or compositions of them. **skip** is a special statement used as a flag to show the end of executions. When it is sequentially composed with other statements, it has no computational effects. A single-step execution of statements is modeled as a labeled transition $_ \xrightarrow{-}_{L_{-}}$, which is a triple of an initial program configuration (a pair of statement and state), a label and a resulting configuration. It is undefined when the initial statement is **skip**. The step aborts if an unsafe instruction is executed.

The high-level language (source language) is defined similarly in Figure 2(c), but it is important to note that its states and primitive instructions may be different from those in the low-level language. The compound statements are almost the same as their low-level counterparts. \mathbb{C}_1 ; \mathbb{C}_2 and $\mathbb{C}_1 ||| \mathbb{C}_2$ are sequential and parallel compositions of \mathbb{C}_1 and \mathbb{C}_2 respectively. Note that we choose to use the same set of compound statements in the two languages for simplicity only. This is not required by our simulation relation, although the analogous program constructs of the two languages (e.g., parallel compositions $C_1 \parallel C_2$ and $\mathbb{C}_1 \parallel \mathbb{C}_2$) make it convenient for us to discuss the compositionality later.

Figure 3 shows part of the definition of $_ \xrightarrow{} H_{-}$, which gives the high-level operational semantics of statements. We often omit the subscript H (or L) in $_ \xrightarrow{} H_{-}$ (or $_ \xrightarrow{} L_{-}$) and the label on top of the arrow when it is τ . The semantics is mostly standard. We only show the rules for primitive instructions and parallel compositions here. Note that when a primitive instruction c is blocked at state Σ (i.e., $\Sigma \notin dom(c)$), we let the program configuration reduce to itself. For example, the instruction lock(1) would be blocked when l is not 0, making it be repeated until l becomes 0; whereas unlock(1) simply sets l to 0 at any time and would never be blocked. Primitive instructions in the high-level and low-level languages are *atomic* in the interleaving semantics. Shortly

we use $_ \longrightarrow *_$ for zero or multiple-step transitions with no events generated, and $_ \stackrel{e}{\longrightarrow} *_$ for multiple-step transitions with *only one* event *e* generated.

3.2. The Event Trace Refinement

Now we can formally define the refinement relation \sqsubseteq that relates the set of externally observable event traces generated by the target and the source programs. A trace is a sequence of events *e*, and may end with a termination marker **done** or a fault marker **abort**.

(*EvtTrace*) $\mathcal{E} ::= \epsilon \mid \mathbf{done} \mid \mathbf{abort} \mid e :: \mathcal{E}$

Definition 3.1 (Event Trace Set). $ETrSet_n(C,\sigma)$ represents a set of external event traces produced by C in n steps from the state σ .

- (1) $ETrSet_0(C,\sigma) \triangleq \{\epsilon\};$
- (2) $ETrSet_{n+1}(C,\sigma) \triangleq$ $\{\mathcal{E} \mid (C,\sigma) \longrightarrow (C',\sigma') \land \mathcal{E} \in ETrSet_n(C',\sigma')$ $\lor (C,\sigma) \xrightarrow{e} (C',\sigma') \land \mathcal{E}' \in ETrSet_n(C',\sigma') \land \mathcal{E} = e :: \mathcal{E}'$ $\lor (C,\sigma) \longrightarrow abort \land \mathcal{E} = abort$ $\lor C = skip \land \mathcal{E} = done \}.$

We define $ETrSet(C, \sigma)$ as $\bigcup_n ETrSet_n(C, \sigma)$.

We overload the notation and use $ETrSet(\mathbb{C}, \Sigma)$ for the high-level language. Note that we treat **abort** as a specific behavior instead of undefined arbitrary behaviors. The choices should depend on applications. The ideas in the article should also apply for the latter setting, though we need to change our refinement and simulation relations defined shortly.

Then we define an event trace refinement as the subset relation between event trace sets, which is similar to Leroy's refinement property [Leroy 2009].

Definition 3.2 (Event Trace Refinement). We say (C, σ) is an *e*-trace refinement of (\mathbb{C}, Σ) , that is, $(C, \sigma) \subseteq (\mathbb{C}, \Sigma)$, if and only if

$$ETrSet(C, \sigma) \subseteq ETrSet(\mathbb{C}, \Sigma)$$

The refinement is defined for program configurations instead of for code only because the initial states may affect the behaviors of programs. In this case, the transformation **T** should translate states as well as code. We overload the notation and use $\mathbf{T}(\Sigma)$ to represent the state transformation, and use $C \sqsubseteq_{\mathbf{T}} \mathbb{C}$ for

$$\forall \sigma, \Sigma. \ \sigma = \mathbf{T}(\Sigma) \Longrightarrow (C, \sigma) \sqsubseteq (\mathbb{C}, \Sigma),$$

then Correct(T) defined in Eq. (1.1) can be reformulated as

$$\mathsf{Correct}(\mathbf{T}) \stackrel{\scriptscriptstyle \Delta}{=} \forall C, \mathbb{C}, C = \mathbf{T}(\mathbb{C}) \implies C \sqsubseteq_{\mathbf{T}} \mathbb{C}. \tag{3.1}$$

From the aforesaid e-trace refinement definition, we can derive an *e-trace equivalence* relation by requiring both directions hold

$$(C,\sigma) \approx (\mathbb{C},\Sigma) \triangleq (C,\sigma) \sqsubseteq (\mathbb{C},\Sigma) \land (\mathbb{C},\Sigma) \sqsubseteq (C,\sigma),$$

and use $C \approx_{\mathbf{T}} \mathbb{C}$ for $\forall \sigma, \Sigma, \sigma = \mathbf{T}(\Sigma) \Longrightarrow (C, \sigma) \approx (\mathbb{C}, \Sigma)$.

4. THE RGSIM RELATION

The e-trace refinement is defined directly over the externally observable behaviors of programs. It is intuitive, and also abstract in that it is independent of language details.



(a) α -related transitions

(b) the side condition of TRANS

Fig. 4. Related transitions.

However, as we explained before, it is *not* compositional with respect to parallel compositions. In this section we propose RGSim, which can be viewed as a compositional proof technique that allows us to derive the simple e-trace refinement and then verify the corresponding transformation \mathbf{T} .

4.1. The Definition

Our co-inductively defined RGSim relation is in the form of $(C, \sigma, \mathcal{R}, \mathcal{G}) \leq_{\alpha;\gamma} (\mathbb{C}, \Sigma, \mathbb{R}, \mathbb{G})$, which is a simulation between program configurations (C, σ) and (\mathbb{C}, Σ) . It is parametrized with the rely and guarantee conditions at the low level and the high level, which are binary relations over states.

$$\mathcal{R}, \mathcal{G} \in \mathcal{P}(LState \times LState), \quad \mathbb{R}, \mathbb{G} \in \mathcal{P}(HState \times HState).$$

The simulation also takes two additional parameters: the *step invariant* α and the *postcondition* γ , which are both relations between the low-level and the high-level states.

$$\alpha, \gamma \in \mathcal{P}(LState \times HState)$$
.

Before we formally define RGSim in Definition 4.2, we first introduce the α -related transitions as follows.

Definition 4.1 (α -Related Transitions). $\langle \mathcal{R}, \mathbb{R} \rangle_{\alpha} \triangleq \{((\sigma, \sigma'), (\Sigma, \Sigma')) \mid (\sigma, \sigma') \in \mathcal{R} \land (\Sigma, \Sigma') \in \mathbb{R} \land (\sigma, \Sigma) \in \alpha \land (\sigma', \Sigma') \in \alpha\}.$

 $\langle \mathcal{R}, \mathbb{R} \rangle_{\alpha}$ represents a set of the α -related transitions in \mathcal{R} and \mathbb{R} , putting together the corresponding transitions in \mathcal{R} and \mathbb{R} that can be related by α , as illustrated in Figure 4(a). $\langle \mathcal{G}, \mathbb{G} \rangle_{\alpha}$ is defined in the same way.

Definition 4.2 (*RGSim*). Whenever $(C, \sigma, \mathcal{R}, \mathcal{G}) \leq_{\alpha; \gamma} (\mathbb{C}, \Sigma, \mathbb{R}, \mathbb{G})$, then $(\sigma, \Sigma) \in \alpha$ and the following are true.

- (1) If $(C, \sigma) \longrightarrow (C', \sigma')$, then there exist \mathbb{C}' and Σ' such that $(\mathbb{C}, \Sigma) \longrightarrow^* (\mathbb{C}', \Sigma'), ((\sigma, \sigma'), (\Sigma, \Sigma')) \in \langle \mathcal{G}, \mathbb{G}^* \rangle_{\alpha}$ and $(C', \sigma', \mathcal{R}, \mathcal{G}) \preceq_{\alpha; \gamma} (\mathbb{C}', \Sigma', \mathbb{R}, \mathbb{G}).$
- (2) If $(C, \sigma) \xrightarrow{e} (C', \sigma')$, then there exist \mathbb{C}' and Σ' such that $(\mathbb{C}, \Sigma) \xrightarrow{e} (\mathbb{C}', \Sigma'), ((\sigma, \sigma'), (\Sigma, \Sigma')) \in \langle \mathcal{G}, \mathbb{G}^* \rangle_{\alpha}$ and $(C', \sigma', \mathcal{R}, \mathcal{G}) \preceq_{\alpha; \gamma} (\mathbb{C}', \Sigma', \mathbb{R}, \mathbb{G}).$
- (3) If $C = \mathbf{skip}$, then there exists Σ' such that $(\mathbb{C}, \Sigma) \longrightarrow^* (\mathbf{skip}, \Sigma'), ((\sigma, \sigma), (\Sigma, \Sigma')) \in \langle \mathcal{G}, \mathbb{G}^* \rangle_{\alpha}, (\sigma, \Sigma') \in \gamma \text{ and } \gamma \subseteq \alpha.$
- (4) If $(C, \sigma) \longrightarrow \text{abort}$, then $(\mathbb{C}, \Sigma) \longrightarrow {}^*\text{abort}$.
- (5) If $((\sigma, \sigma'), (\Sigma, \Sigma')) \in \langle \mathcal{R}, \mathbb{R}^* \rangle_{\alpha}$, then $(C, \sigma', \mathcal{R}, \mathcal{G}) \preceq_{\alpha; \gamma} (\mathbb{C}, \Sigma', \mathbb{R}, \mathbb{G})$.



Fig. 5. Simulation diagrams of RGSim.

Then, $(C, \mathcal{R}, \mathcal{G}) \preceq_{\alpha; \zeta} {}_{\gamma} (\mathbb{C}, \mathbb{R}, \mathbb{G})$ iff for all σ and Σ , if $(\sigma, \Sigma) \in \zeta$, then $(C, \sigma, \mathcal{R}, \mathcal{G}) \preceq_{\alpha; \gamma} (\mathbb{C}, \Sigma, \mathbb{R}, \mathbb{G})$. Here the *precondition* $\zeta \in \mathcal{P}(LState \times HState)$ is used to relate the initial states σ and Σ .

Informally, $(C, \sigma, \mathcal{R}, \mathcal{G}) \preceq_{\alpha; \gamma} (\mathbb{C}, \Sigma, \mathbb{R}, \mathbb{G})$ says the low-level configuration (C, σ) is simulated by the high-level configuration (\mathbb{C}, Σ) with behaviors \mathcal{G} and \mathbb{G} respectively, no matter how their environments \mathcal{R} and \mathbb{R} interfere with them. It requires the following hold for every execution of C.

- Starting from α -related states, each step of *C* corresponds to zero or multiple steps of \mathbb{C} , and the resulting states are α -related too. If an external event is produced in the step of *C*, the same event should be produced by \mathbb{C} . We show the simulation diagram with events generated by the program steps in Figure 5(a), where solid lines denote hypotheses and dashed lines denote conclusions, following Leroy's notations [Leroy 2009].
- The α relation reflects the abstractions from the low-level machine model to the high-level one, and is preserved by the related transitions at the two levels (so it is an *invariant*). For instance, when verifying a fine-grained implementation of sets, the α relation may relate a concrete representation in memory (e.g., a linked-list) at the low level to the corresponding abstract mathematical set at the high level.
- The corresponding transitions of C and \mathbb{C} need to be in $\langle \mathcal{G}, \mathbb{G}^* \rangle_{\alpha}$. That is, for each step of C, its state transition should satisfy the guarantee \mathcal{G} , and the corresponding transition made by the multiple steps of \mathbb{C} should be in the transitive closure of \mathbb{G} . The guarantees are abstractions of the programs' behaviors. As we will show later in the PAR rule in Figure 7, they will serve as the rely conditions of the sibling threads at the time of parallel compositions. Note that we do not need each step of \mathbb{C} to be in \mathbb{G} , although we could do so. This is because we only care about the coarse-grained behaviors (with mumbling) of the source that are used to simulate the target. We will explain more by the example (4.1) in Section 4.2.
- If *C* terminates, then \mathbb{C} terminates as well, and the final states should be related by the postcondition γ . We require $\gamma \subseteq \alpha$, that is, the final state relation is not weaker than the step invariant.
- -C is not safe only if \mathbb{C} is not safe either. This means the transformation should not make a safe high-level program unsafe at the low level.
- Whatever the low-level environment \mathcal{R} and the high-level one \mathbb{R} do, as long as the state transitions are α -related, they should not affect the simulation between C and \mathbb{C} , as shown in Figure 5(b). Here a step in \mathcal{R} may correspond to zero or multiple steps of \mathbb{R} . Note that different from the program steps, some steps of \mathcal{R} may not correspond to steps of \mathbb{R} . On the other hand, only requiring that \mathcal{R} be simulated by \mathbb{R} (see (4.2) in Section 4.2) is not sufficient for parallel compositionality, which we will explain later in Section 4.2.

$$\begin{split} & \mathsf{InitRel}_{\mathbf{T}}(\zeta) \, \triangleq \, \forall \sigma, \Sigma, \sigma = \mathbf{T}(\Sigma) \implies (\sigma, \Sigma) \in \zeta \\ & \mathcal{B} \Leftrightarrow \mathbb{B} \, \triangleq \, \{(\sigma, \Sigma) \mid \mathcal{B} \, \sigma = \mathbb{B} \, \Sigma\} \qquad \mathcal{B} \land \mathbb{B} \, \triangleq \, \{(\sigma, \Sigma) \mid \mathcal{B} \, \sigma \land \mathbb{B} \, \Sigma\} \\ & \mathsf{Intuit}(\alpha) \, \triangleq \, \forall \sigma, \Sigma, \sigma', \Sigma', (\sigma, \Sigma) \in \alpha \land \sigma \subseteq \sigma' \land \Sigma \subseteq \Sigma' \implies (\sigma', \Sigma') \in \alpha \\ & \alpha \uplus \beta \, \triangleq \, \{(\sigma_1 \uplus \sigma_2, \Sigma_1 \uplus \Sigma_2) \mid (\sigma_1, \Sigma_1) \in \alpha \land (\sigma_2, \Sigma_2) \in \beta\} \qquad \eta \ \# \, \alpha \, \triangleq \, (\eta \cap \alpha) \subseteq (\eta \uplus \alpha) \\ & \beta \circ \alpha \, \triangleq \, \{(\sigma, \Sigma) \mid \exists \theta, (\sigma, \theta) \in \alpha \land (\theta, \Sigma) \in \beta\} \qquad \alpha^{-1} \, \triangleq \, \{(\Sigma, \sigma) \mid (\sigma, \Sigma) \in \alpha\} \\ & \mathsf{Id} \, \triangleq \, \{(\sigma, \sigma) \mid \sigma \in LState\} \qquad \mathsf{True} \, \triangleq \, \{(\sigma, \sigma') \mid \sigma, \sigma' \in LState\} \\ & \mathsf{R}_{\mathsf{M}} \text{ isMidOf } (\alpha, \beta; \mathcal{R}, \mathbb{R}) \, \triangleq \, \forall \sigma, \sigma', \Sigma, \Sigma', \, ((\sigma, \sigma'), (\Sigma, \Sigma')) \in \langle \mathcal{R}, \mathbb{R} \rangle_{\beta \circ \alpha} \\ \qquad \qquad \qquad \Rightarrow \forall \theta, (\sigma, \theta) \in \alpha \land (\theta, \Sigma) \in \beta \\ & \implies \exists \theta', ((\sigma, \sigma'), (\theta, \theta')) \in \langle \mathcal{R}, \mathsf{R}_{\mathsf{M}} \rangle_{\alpha} \land ((\theta, \theta'), (\Sigma, \Sigma')) \in \langle \mathsf{R}_{\mathsf{M}}, \mathbb{R} \rangle_{\beta} \end{split}$$

Fig. 6. Auxiliary definitions for RGSim.

Then based on the simulation, we hide the states by the precondition ζ and define the RGSim relation between programs only. By the definition we know $\zeta \subseteq \alpha$ if $(C, \mathcal{R}, \mathcal{G}) \preceq_{\alpha; \zeta} \gamma$ ($\mathbb{C}, \mathbb{R}, \mathbb{G}$), that is, the precondition needs to be no weaker than the step invariant. Usually in practice α is very weak and naturally implied by the preand postconditions ζ and γ , for example, ζ and γ are the same as α in examples in Section 6.

RGSim is sound with respect to the e-trace refinement (Definition 3.2). That is, $(C, \sigma, \mathcal{R}, \mathcal{G}) \preceq_{\alpha; \gamma} (\mathbb{C}, \Sigma, \mathbb{R}, \mathbb{G})$ ensures that (C, σ) does not have more observable behaviors than (\mathbb{C}, Σ) .

THEOREM 4.3 (SOUNDNESS/ADEQUACY). If there exist \mathcal{R} , \mathcal{G} , \mathbb{R} , \mathbb{G} , α and γ such that $(C, \sigma, \mathcal{R}, \mathcal{G}) \leq_{\alpha:\gamma} (\mathbb{C}, \Sigma, \mathbb{R}, \mathbb{G})$, then $(C, \sigma) \subseteq (\mathbb{C}, \Sigma)$.

The soundness theorem shows that RGSim is a proof technique for the simple and natural refinement \sqsubseteq , which is what we ultimately care about. The theorem can be proved by first strengthening the relies to the identity transitions and weakening the guarantees to the universal relations. Then we prove that the resulting simulation under identity environments implies the e-trace refinement. The mechanized proof in the Coq proof assistant [2010] is available online.

For program transformations, since the initial state for the target program is transformed from the initial state for the source, we use $\text{InitRel}_{\mathbf{T}}(\zeta)$ (defined in Figure 6) to say the transformation \mathbf{T} over states ensures the binary precondition ζ .

COROLLARY 4.4. If there exist \mathcal{R} , \mathcal{G} , \mathbb{R} , \mathbb{G} , α , ζ and γ such that $\mathsf{InitRel}_{\mathbf{T}}(\zeta)$ and $(C, \mathcal{R}, \mathcal{G}) \leq_{\alpha' \zeta' - \gamma} (\mathbb{C}, \mathbb{R}, \mathbb{G})$, then $C \sqsubseteq_{\mathbf{T}} \mathbb{C}$.

4.2. Compositionality Rules

RGSim is compositional with respect to various program constructs, including parallel compositions. We present the compositionality rules in Figure 7, which gives us a relational proof method for concurrent program transformations.

As in the R-G logic [Jones 1983], we require that the pre- and postconditions be *stable* under the interference from the environments. Here we introduce the concept of stability of a relation ζ with respect to a set of transition pairs $\Lambda \in \mathcal{P}((LState \times LState) \times (HState \times HState))$.

Definition 4.5 (*Stability*). Sta(ζ , Λ) holds iff for all σ , σ' , Σ and Σ' , if (σ , Σ) $\in \zeta$ and ((σ , σ'), (Σ , Σ')) $\in \Lambda$, then (σ' , Σ') $\in \zeta$.

$$\frac{\zeta \subseteq \alpha}{(\mathbf{skip}, \mathcal{R}, \mathbf{ld}) \leq_{\alpha;\zeta - \gamma} (\mathbf{c}_{1}, \mathbb{R}, \mathbb{G})} (\mathbf{c}_{2}, \mathcal{R}, \mathcal{G}) \leq_{\alpha;\gamma - \eta} (\mathbb{C}_{2}, \mathbb{R}, \mathbb{G})} (\mathbf{c}_{1}, \mathcal{R}, \mathcal{G}) \leq_{\alpha;\zeta - \gamma} (\mathbb{C}_{1}, \mathbb{R}, \mathbb{G})} (\mathbb{C}_{2}, \mathcal{R}, \mathcal{G}) \leq_{\alpha;\gamma - \eta} (\mathbb{C}_{2}, \mathbb{R}, \mathbb{G})} (\mathbf{c}_{2}, \mathcal{R}, \mathcal{G}) = (\mathbb{C}_{1}, \mathcal{R}, \mathcal{G}) \leq_{\alpha;\zeta - \gamma} (\mathbb{C}_{1}, \mathbb{R}, \mathbb{G})} (\mathbb{C}_{2}, \mathcal{R}, \mathcal{G}) \leq_{\alpha;\zeta - \gamma} (\mathbb{C}_{2}, \mathbb{R}, \mathbb{G})} (\mathbb{C}_{1}, \mathcal{R}, \mathcal{G}) = (\mathbb{C}_{1}, \mathcal{R}, \mathcal{G}) = (\mathbb{C}_{2}, \mathcal{R}, \mathcal{G}) \leq_{\alpha;\zeta - \gamma} (\mathbb{C}_{2}, \mathbb{R}, \mathbb{G})} (\mathbb{C}_{1}, \mathbb{R}, \mathbb{G}) = (\mathbb{C}_{2}, \mathbb{R}, \mathbb{G}) = (\mathbb{C}_{1}, \mathbb{C}, \mathbb{R}, \mathbb{G}) = (\mathbb{C}, \mathbb{R}, \mathbb{G}, \mathbb{G}) = (\mathbb{C}, \mathbb{R}, \mathbb{G}, \mathbb{G}, \mathbb{G}) = (\mathbb{C}, \mathbb{R}, \mathbb{G}, \mathbb{G}) = (\mathbb{C}, \mathbb{R}, \mathbb{G}, \mathbb{G}) = (\mathbb{C}, \mathbb{R}, \mathbb{G}, \mathbb{G}) = ($$

Fig. 7. Compositionality rules for RGSim. At each proof rule, we implicitly assume that the pre- and postconditions are stable under the environments' interference (Definition 4.5), and the relies and guarantees are closed over identity transitions.

Usually we need $\operatorname{Sta}(\zeta, \langle \mathcal{R}, \mathbb{R}^* \rangle_{\alpha})$, which says whenever ζ holds initially and \mathcal{R} and \mathbb{R}^* perform related actions, the resulting states still satisfy ζ . By unfolding $\langle \mathcal{R}, \mathbb{R}^* \rangle_{\alpha}$, we could see that α itself is stable with respect to any α -related transitions, that is,

Sta(α , $(\mathcal{R}, \mathbb{R}^*)_{\alpha}$). Another simple example is given next, where both environments could increment x and the unary stable assertion $x \ge 0$ is lifted to the relation ζ .

$$\begin{split} \zeta &\triangleq \{(\sigma, \Sigma) \mid \sigma(\mathbf{x}) = \Sigma(\mathbf{x}) \land \sigma(\mathbf{x}) \geq 0\} \\ \mathcal{R} &\triangleq \{(\sigma, \sigma') \mid \sigma' = \sigma\{\mathbf{x} \quad \sigma(\mathbf{x}) + 1\}\} \end{split} \qquad \begin{aligned} \alpha &\triangleq \{(\sigma, \Sigma) \mid \sigma(\mathbf{x}) = \Sigma(\mathbf{x})\} \\ \mathbb{R} &\triangleq \{(\Sigma, \Sigma') \mid \Sigma' = \Sigma\{\mathbf{x} \quad \Sigma(\mathbf{x}) + 1\}\} \end{aligned}$$

We can prove $\operatorname{Sta}(\zeta, \langle \mathcal{R}, \mathbb{R}^* \rangle_{\alpha})$. Stability of the pre- and postconditions under the environments' interference is assumed as an implicit side condition at every proof rule in Figure 7, for example, we assume $\operatorname{Sta}(\zeta, \langle \mathcal{R}, \mathbb{R}^* \rangle_{\alpha})$ in the SKIP rule. We also require implicitly that the relies and guarantees are closed over identity transitions, since stuttering steps will not affect observable event traces.

In Figure 7, the rules SKIP, SEQ, IF and WHILE reveal a high degree of similarity to the corresponding inference rules in Hoare logic. In the SEQ rule, γ serves as the postcondition of C_1 and \mathbb{C}_1 and the precondition of C_2 and \mathbb{C}_2 at the same time. The IF rule requires the boolean conditions of both sides to be evaluated to the same value under the precondition ζ . The definitions of the sets $B \Leftrightarrow \mathbb{B}$ and $B \land \mathbb{B}$ are given in Figure 6. The rule also requires the precondition ζ to imply the step invariant α . In the WHILE rule, the γ relation is viewed as a loop invariant preserved at the loop entry point, and needs to ensure $B \Leftrightarrow \mathbb{B}$.

Parallel compositionality. The PAR rule shows parallel compositionality of RGSim. The interference constraints say that two threads can be composed in parallel if one thread's guarantee implies the rely of the other. After parallel composition, they are expected to run in the common environment and their guaranteed behaviors contain each single thread's behaviors.

Note that, although RGSim does not require every step of the high-level program to be in its guarantee (see the first two conditions in Definition 4.2), this relaxation does not affect the parallel compositionality. This is because the target could have less behaviors than the source. To let $\mathbb{C}_1 || \mathbb{C}_2$ simulate $C_1 || C_2$, we only need a subset of the interleavings of \mathbb{C}_1 and \mathbb{C}_2 to simulate those of C_1 and C_2 . Thus the high-level relies and guarantees need to ensure the existence of those interleavings only. Next we give a simple example to explain this subtle issue. We can prove

$$(\mathbf{x}:=\mathbf{x}+2,\mathcal{R},\mathcal{G}) \leq_{\alpha:\tau} (\mathbf{x}:=\mathbf{x}+1;\mathbf{x}:=\mathbf{x}+1,\mathbb{R},\mathbb{G}),$$

$$(4.1)$$

where the relies and the guarantees say x can be increased by 2 and α , ζ and γ relate x of the two sides.

$$\mathcal{R} = \mathcal{G} \triangleq \{(\sigma, \sigma') \mid \sigma' = \sigma \lor \sigma' = \sigma \{\mathbf{x} \quad \sigma(\mathbf{x}) + 2\}\};\\ \mathbb{R} = \mathbb{G} \triangleq \{(\Sigma, \Sigma') \mid \Sigma' = \Sigma \lor \Sigma' = \Sigma \{\mathbf{x} \quad \Sigma(\mathbf{x}) + 2\}\};\\ \alpha = \zeta = \gamma \triangleq \{(\sigma, \Sigma) \mid \sigma(\mathbf{x}) = \Sigma(\mathbf{x})\}.$$

Note that the high-level program is actually finer grained than its guarantee, but to prove Eq. (4.1) we only need the execution in which it goes two steps to the end without interference from its environment. Also we can prove $(print(x), \mathcal{R}, \mathcal{G}) \leq_{\alpha; \zeta} \gamma$ $(print(x), \mathbb{R}, \mathbb{G})$. Here we use the instruction **print**(*E*) to observe the value of x, which will produce an external event **out**(*n*) if *E* evaluates to *n*. Then by the PAR rule, we get

$$(\mathbf{x}:=\mathbf{x}+2\|\operatorname{print}(\mathbf{x}),\mathcal{R},\mathcal{G}) \leq_{\alpha:\mathcal{L}} ((\mathbf{x}:=\mathbf{x}+1;\mathbf{x}:=\mathbf{x}+1)\|\|\operatorname{print}(\mathbf{x}),\mathbb{R},\mathbb{G}),$$

which does not violate the natural meaning of refinements. That is, all the possible external events produced by the low-level side can also be produced by the highlevel side, although the latter could have more external behaviors due to its finer granularity.

ACM Transactions on Programming Languages and Systems, Vol. 36, No. 1, Article 3, Publication date: March 2014.

Another subtlety in the RGSim definition is with the fifth condition over the environments, which is crucial for parallel compositionality. One may think a more natural alternative to this condition is to require that \mathcal{R} be simulated by \mathbb{R} .

If
$$(\sigma, \sigma') \in \mathcal{R}$$
, then there exists Σ' such that
 $(\Sigma, \Sigma') \in \mathbb{R}^*$ and $(C, \sigma', \mathcal{R}, \mathcal{G}) \preceq'_{\alpha;\gamma} (\mathbb{C}, \Sigma', \mathbb{R}, \mathbb{G})$.
$$(4.2)$$

We refer to this modified simulation definition as \leq' . Unfortunately, \leq' does not have parallel compositionality. As a counter-example, if the invariant α says the left side x is not greater than the right side x, that is,

$$\alpha \triangleq \{(\sigma, \Sigma) \mid \sigma(\mathbf{x}) \leq \Sigma(\mathbf{x})\},\$$

we could prove the following.

$$(x:=x+1, \text{ Id, True}) \leq_{\alpha;\alpha}' \alpha \quad (x:=x+2, \text{ Id, True});$$

$$(4.3)$$

$$(x:=0;print(x), True, Id) \leq_{\alpha'\alpha'}' (x:=0;print(x), True, Id).$$
 (4.4)

Here we use Id and True (defined in Figure 6) for the sets of identity transitions and arbitrary transitions respectively, and overload the notations at the low level to the high level. However, the following refinement does *not* hold after parallel composition.

$$(x:=x+1||(x:=0;print(x)), Id, True) \leq_{\alpha:\alpha} \alpha (x:=x+2|||(x:=0;print(x)), Id, True)$$

This is because the rely \mathcal{R} (or \mathbb{R}) is an abstraction of all the permitted behaviors in the environment of a thread t. Any thread t' whose behaviors are allowed in \mathcal{R} (or \mathbb{R}) can run in parallel with t. Thus to obtain parallel compositionality, we have to ensure that the simulation is preserved with *any* possible sibling thread t'. With *our* definition \leq , the refinement of Eq. (4.4) is not provable, because after some α -related transitions of environments, the target may print a value smaller than the one printed by the source.

Other rules. We also develop some other useful rules about RGSim. For example, the STREN- α rule allows us to replace the invariant α by a stronger invariant α' . We need to check that α' is indeed an invariant preserved by the related program steps, that is, $Sta(\alpha', \langle \mathcal{G}, \mathbb{G}^* \rangle_{\alpha})$ holds. Symmetrically, the WEAKEN- α rule requires α to be preserved by environment steps related by the weaker invariant α' . As usual, the pre- and postconditions, the relies and the guarantees can be strengthened or weakened by the CONSEQ rule.

The FRAME rule allows us to use local specifications [Reynolds 2002]. When verifying the simulation between C and \mathbb{C} , we need to only talk about the locally used resource in α , ζ and γ , and the local relies and guarantees \mathcal{R} , \mathcal{G} , \mathbb{R} and \mathbb{G} . Then the proof can be reused in contexts where some extra resource η is used, and the accesses of it respect the invariant β and \mathcal{R}_1 , \mathcal{G}_1 , \mathbb{R}_1 and \mathbb{G}_1 . We give the auxiliary definitions in Figure 6. The disjoint union \uplus between states is lifted to state pairs. A state relation α is intuitionistic, denoted by Intuit(α), if it is monotone with respect to the extension of states. The disjoint state pairs satisfying η and α respectively. For example, let $\eta \triangleq \{(\sigma, \Sigma) \mid \sigma(y) = \Sigma(y)\}$ and $\alpha \triangleq \{(\sigma, \Sigma) \mid \sigma(x) = \Sigma(x)\}$ where x and y are two distinct variables, then both η and α are intuitionistic and $\eta \# \alpha$ holds. We also require η to be stable under interference from the programs (i.e., the programs do not change the extra resource) and the extra environments. We use $\eta \# \{\zeta, \gamma, \alpha\}$ as a shorthand for $(\eta \# \zeta) \land (\eta \# \gamma) \land (\eta \# \alpha)$. Similar representations are used in this rule.

Finally, the transitivity rule TRANS allows us to verify a transformation by using an intermediate level as a bridge. The intermediate environment R_M should be chosen

with caution so that the $(\beta \circ \alpha)$ -related transitions can be decomposed into β -related and α -related transitions, as illustrated in Figure 4(b). Here \circ defines the composition of two relations and isMidOf defines the side condition over the environments, as shown in Figure 6. We use θ for a middle-level state.

Soundness. All the rules in Figure 7 are sound, that is, for each rule the premises imply the conclusion. We prove their soundness by co-induction, directly following the definition of RGSim. The proofs are checked in the Coq proof assistant [2010].

Instantiations of relies and guarantees. We can derive the sequential refinement and the fully-abstract-semantics-based refinement by instantiating the rely conditions in RGSim. For example, the refinement of Eq. (4.5) over closed programs assumes identity environments, making the interference constraints in the PAR rule unsatisfiable. This confirms the observation in Section 2.1 that the sequential refinement loses parallel compositionality.

$$(C, \mathsf{Id}, \mathsf{True}) \preceq_{\alpha; \ell} (\mathbb{C}, \mathsf{Id}, \mathsf{True})$$

$$(4.5)$$

The refinement of Eq. (4.6) assumes arbitrary environments, which makes the interference constraints in the PAR rule trivially true. But this assumption is too strong: usually (4.6) cannot be satisfied in practice.

$$(C, \mathsf{True}, \mathsf{True}) \preceq_{\alpha; \zeta} _{\nu} (\mathbb{C}, \mathsf{True}, \mathsf{True})$$

$$(4.6)$$

4.3. A Simple Example

Shortly we give a simple example to illustrate the use of RGSim and its parallel compositionality in verifying concurrent program transformations. The high-level program $\mathbb{C}_1 || \mathbb{C}_2$ is transformed to $C_1 || C_2$, using a lock 1 to synchronize the accesses of the shared variable x. We aim to prove $C_1 || C_2 \sqsubseteq_{\mathbf{T}} \mathbb{C}_1 || \mathbb{C}_2$. That is, although x:=x+2 is implemented by two steps of incrementing x in C_2 , the parallel observer C_1 will not print unexpected values. Here we view output events as externally observable behaviors.

To facilitate the proof, we introduce an auxiliary shared variable X at the low level to record the value of x at the time when releasing the lock. It specifies the value of x outside every critical section, thus should match the value of the high-level x after every corresponding action. Here $\langle C \rangle$ means C is executed atomically. Its semantics follows RGSep [Vafeiadis 2008] (or see Section 6.2). The auxiliary variable is writeonly and would not affect the external behaviors of the program [Abadi and Lamport 1991]. Thus in what follows we can focus on the instrumented target program with the auxiliary code.

By the soundness and compositionality of RGSim, we only need to prove simulations over individual threads, providing appropriate relies and guarantees. We first define the invariant α , which only cares about the value of x when the lock is free.

$$\alpha \triangleq \{(\sigma, \Sigma) \mid \sigma(X) = \Sigma(x) \land (\sigma(1) = 0 \Longrightarrow \sigma(x) = \sigma(X))\}.$$

We let the pre- and postconditions be α as well.

ACM Transactions on Programming Languages and Systems, Vol. 36, No. 1, Article 3, Publication date: March 2014.

The high-level threads can be executed in arbitrary environments with arbitrary guarantees: $\mathbb{R} = \mathbb{G} \triangleq$ True. The transformation uses the lock to protect every access of x, thus the low-level relies and guarantees are not arbitrary.

$$\mathcal{R} \triangleq \{(\sigma, \sigma') \mid \sigma(1) = \operatorname{cid} \Longrightarrow \\ \sigma(\mathbf{x}) = \sigma'(\mathbf{x}) \land \sigma(\mathbf{X}) = \sigma'(\mathbf{X}) \land \sigma(1) = \sigma'(1)\}; \\ \mathcal{G} \triangleq \{(\sigma, \sigma') \mid \sigma' = \sigma \lor \sigma(1) = \mathbf{0} \land \sigma' = \sigma\{\mathbf{1} \quad \operatorname{cid}\} \\ \lor \sigma(1) = \operatorname{cid} \land \sigma' = \sigma\{\mathbf{x} \quad _\} \\ \lor \sigma(1) = \operatorname{cid} \land \sigma' = \sigma\{\mathbf{1} \quad \mathbf{0}, \mathbf{X} \quad _\}\}.$$

Every low-level thread guarantees that it updates x only when the lock is acquired. Its environment cannot update x or l if the current thread holds the lock. Here cid is the identifier of the current thread. When acquired, the lock holds the identifier of the owner thread.

Following the definition, we can prove $(C_1, \mathcal{R}, \mathcal{G}) \preceq_{\alpha;\alpha \ \alpha} (\mathbb{C}_1, \mathbb{R}, \mathbb{G})$ and $(C_2, \mathcal{R}, \mathcal{G}) \preceq_{\alpha;\alpha \ \alpha} (\mathbb{C}_2, \mathbb{R}, \mathbb{G})$. By applying the PAR rule and from the soundness of RGSim (Corollary 4.4), we know $C_1 || C_2 \sqsubseteq_{\mathbf{T}} \mathbb{C}_1 || \mathbb{C}_2$ holds for any **T** that respects α . Perhaps interestingly, if we omit the lock and unlock operations in C_1 , then $C_1 || C_2$

Perhaps interestingly, if we omit the lock and unlock operations in C_1 , then $C_1 || C_2$ would have more externally observable behaviors than $\mathbb{C}_1 ||| \mathbb{C}_2$. This does *not* indicate the unsoundness of our PAR rule (which is sound!). The reason is that x might have different values on the two levels after the environments' α -related transitions, so that we cannot have $(\text{print}(x), \mathcal{R}, \mathcal{G}) \leq_{\alpha;\alpha} \alpha$ (print(x), \mathbb{R}, \mathbb{G}) with the current definitions of α , \mathcal{R} and \mathcal{G} , even though the code of the two sides is syntactically identical.

The use of the auxiliary variable. The auxiliary variable X helps us define the invariant α and do the proof. It is difficult to prove the refinement without this auxiliary variable. One may wish to prove

$$(C_1, \mathcal{R}', \mathcal{G}') \leq_{\alpha':\alpha'=\alpha'} (\mathbb{C}_1, \mathbb{R}, \mathbb{G}), \tag{4.7}$$

where α' , \mathcal{R}' and \mathcal{G}' are defined as follows by eliminating X from α , \mathcal{R} and \mathcal{G} .

$$\begin{array}{l} \alpha' \triangleq \{(\sigma, \Sigma) \mid \sigma(1) = 0 \Longrightarrow \sigma(\mathbf{x}) = \Sigma(\mathbf{x})\}; \\ \mathcal{R}' \triangleq \{(\sigma, \sigma') \mid \sigma(1) = \operatorname{cid} \Longrightarrow \sigma(\mathbf{x}) = \sigma'(\mathbf{x}) \land \sigma(1) = \sigma'(1)\}; \\ \mathcal{G}' \triangleq \{(\sigma, \sigma') \mid \sigma' = \sigma \lor \sigma(1) = 0 \land \sigma' = \sigma\{1 \quad \operatorname{cid}\} \\ \lor \sigma(1) = \operatorname{cid} \land \sigma' = \sigma\{\mathbf{x} \quad .\} \\ \lor \sigma(1) = \operatorname{cid} \land \sigma' = \sigma\{1 \quad 0\}\}. \end{array}$$

But Eq. (4.7) does not hold because $\langle \mathcal{R}', \mathbb{R}^* \rangle_{\alpha'}$ (which is used in Definition 4.2(5)) permits unexpected transitions. For instance, we allow $((\sigma, \sigma'), (\Sigma, \Sigma')) \in \langle \mathcal{R}', \mathbb{R}^* \rangle_{\alpha'}$ for the following σ, σ', Σ and Σ' .

$$\sigma = \sigma' \triangleq \{ \mathbf{x} \quad 0, 1 \quad \text{cid} \}; \qquad \Sigma \triangleq \{ \mathbf{x} \quad 0 \}; \qquad \Sigma' \triangleq \{ \mathbf{x} \quad 1 \}$$

The high-level environment is allowed to change x even if the thread holds the lock at the low level. Then the left thread may print out different values at the two levels, breaking the simulation (4.7).

It is possible to define the RGSim relation in another way that allows us to get rid of the auxiliary variable for this example. Instead of defining separate rely/guarantee relations at the two levels and using α to relate them, we can directly define "relational rely/guarantee" relations $r, g \in \mathcal{P}((LState \times LState) \times (HState \times HState))$ The new simulation is in the form of $C \preceq_{\alpha;\zeta} \underset{\gamma;r;g}{\cong} \mathbb{C}$ and defined by substituting r and g for $\langle \mathcal{R}, \mathbb{R}^* \rangle_{\alpha}$

and $\langle \mathcal{G}, \mathbb{G}^* \rangle_{\alpha}$ in Definition 4.2. It has all the nice properties of our current RGSim relation (including parallel compositionality) and we no longer need auxiliary variables to prove the simple example. We can prove the new simulations $C_1 \leq_{\alpha';\alpha'} \alpha';r;g \ \mathbb{C}_1$ and $C'_2 \leq_{\alpha';\alpha'} \alpha';r;g \ \mathbb{C}_2$. Here C'_2 results from removing X from C_2 , α' is defined as given before, and r and g are as follows.

$$r \triangleq \{ ((\sigma, \sigma'), (\Sigma, \Sigma')) \mid \sigma(1) = \operatorname{cid} \Longrightarrow \sigma(x) = \sigma'(x) \land \sigma(1) = \sigma'(1) \land \Sigma(x) = \Sigma'(x) \}; \\ g \triangleq \{ ((\sigma, \sigma'), (\Sigma, \Sigma')) \mid \sigma' = \sigma \land \Sigma' = \Sigma \lor \sigma(1) = 0 \land \sigma' = \sigma\{1 \quad \operatorname{cid}\} \land \Sigma' = \Sigma \\ \lor \sigma(1) = \operatorname{cid} \land \sigma' = \sigma\{x \quad _{-}\} \land \Sigma' = \Sigma \\ \lor \sigma(1) = \operatorname{cid} \land \sigma' = \sigma\{1 \quad 0\} \land \Sigma' = \Sigma\{x \quad \sigma(x)\} \}.$$

We can see that if the thread holds the lock at the low level, neither the high-level or the low-level environment can change x. This relational r does not permit the unexpected transitions discussed before. It is more expressive than $\langle \mathcal{R}', \mathbb{R}^* \rangle_{\alpha'}$, but is also much heavier. We choose to present the current RGSim relation because in practice it is usually easier to define separate rely/guarantee conditions at the two levels.

More discussions. RGSim ensures that the target program preserves safety properties (including the partial correctness) of the source, but allows a terminating source program to be transformed to a target having infinite silent steps. In the previous example, this allows the low-level programs to be blocked forever (e.g., at the time when the lock is held but never released by some other thread). Proving the preservation of the termination behavior would require liveness proofs in a concurrent setting (e.g., proving the absence of deadlock), which we leave as future work.

In the next three sections, we show more serious examples to demonstrate the applicability of RGSim.

5. RELATIONAL REASONING ABOUT OPTIMIZATIONS

As a general correctness notion of concurrent program transformations, RGSim establishes a relational approach to justify compiler optimizations on concurrent programs. In the following we adapt Benton's work [2004] on sequential optimizations to the concurrent setting.

5.1. Optimization Rules

Usually optimizations depend on particular contexts, for example, the assignment x := E can be eliminated only in the context that the value of x is never used after the assignment. In a shared-state concurrent setting, we should also consider the parallel context for an optimization. RGSim enables us to specify various sophisticated requirements for the parallel contexts by rely/guarantee conditions. Based on RGSim, we provide a set of inference rules to characterize and justify common optimizations (e.g., dead code elimination) with information of both the sequential and the parallel contexts. Note in this section the target and the source programs are in the same language.

Sequential Unit Laws

$$\frac{(C_1, \mathcal{R}_1, \mathcal{G}_1) \preceq_{\alpha; \zeta - \gamma} (C_2, \mathcal{R}_2, \mathcal{G}_2)}{(\mathbf{skip}; C_1, \mathcal{R}_1, \mathcal{G}_1) \preceq_{\alpha; \zeta - \gamma} (C_2, \mathcal{R}_2, \mathcal{G}_2)} \qquad \qquad \frac{(C_1, \mathcal{R}_1, \mathcal{G}_1) \preceq_{\alpha; \zeta - \gamma} (C_2, \mathcal{R}_2, \mathcal{G}_2)}{(C_1, \mathcal{R}_1, \mathcal{G}_1) \preceq_{\alpha; \zeta - \gamma} (\mathbf{skip}; C_2, \mathcal{R}_2, \mathcal{G}_2)}$$

Plus the variants with **skip** after the code C_1 or C_2 . That is, **skip**s could be arbitrarily introduced and eliminated.

ACM Transactions on Programming Languages and Systems, Vol. 36, No. 1, Article 3, Publication date: March 2014.

Common Branch

$$\begin{array}{c} \forall \sigma_1, \sigma_2. \ (\sigma_1, \sigma_2) \in \zeta \implies B \ \sigma_2 \neq \bot \\ (C, \mathcal{R}, \mathcal{G}) \preceq_{\alpha; \zeta_1 \quad \gamma} (C_1, \mathcal{R}', \mathcal{G}') \qquad \zeta_1 = (\zeta \cap (\mathbf{true} \land B)) \\ (C, \mathcal{R}, \mathcal{G}) \preceq_{\alpha; \zeta_2 \quad \gamma} (C_2, \mathcal{R}', \mathcal{G}') \qquad \zeta_2 = (\zeta \cap (\mathbf{true} \land \neg B)) \\ \hline (C, \mathcal{R}, \mathcal{G}) \preceq_{\alpha; \zeta_2 \quad \gamma} (\mathbf{if} \ (B) \ C_1 \ \mathbf{else} \ C_2, \mathcal{R}', \mathcal{G}') \end{array}$$

This rule says that, when the if-condition can be evaluated and both branches can be optimized to the same code C, we can transform the whole if-statement to C without introducing new behaviors.

Known Branch

$$\frac{(C, \mathcal{R}, \mathcal{G}) \preceq_{\alpha; \zeta \quad \gamma} (C_1, \mathcal{R}', \mathcal{G}') \qquad \zeta = (\zeta \cap (\mathbf{true} \land B))}{(C, \mathcal{R}, \mathcal{G}) \preceq_{\alpha; \zeta \quad \gamma} (\mathbf{if} (B) C_1 \mathbf{else} C_2, \mathcal{R}', \mathcal{G}')}$$
$$\frac{(C, \mathcal{R}, \mathcal{G}) \preceq_{\alpha; \zeta \quad \gamma} (C_2, \mathcal{R}', \mathcal{G}') \qquad \zeta = (\zeta \cap (\mathbf{true} \land \neg B))}{(C, \mathcal{R}, \mathcal{G}) \preceq_{\alpha; \zeta \quad \gamma} (\mathbf{if} (B) C_1 \mathbf{else} C_2, \mathcal{R}', \mathcal{G}')}$$

Since the if-condition B is **true** (or **false**) initially, we can consider the then-branch (or the else-branch) only. These rules can be derived from the common-branch rule.

Dead While

$$\frac{\zeta = (\zeta \cap (\mathbf{true} \land \neg B))}{(\mathbf{skip}, \mathcal{R}_1, \mathsf{Id}) \preceq_{\alpha; \zeta = \zeta} (\mathbf{while} \ (B)\{C\}, \mathcal{R}_2, \mathsf{Id})}$$

We can eliminate the loop, if the loop condition is **false** (no matter how the environments update the states) at the loop entry point.

Loop Peeling

$$(\textbf{while} (B)\{C\}, \mathcal{R}_1, \mathcal{G}_1) \preceq_{\alpha; \zeta - \gamma} (\textbf{while} (B)\{C\}, \mathcal{R}_2, \mathcal{G}_2)$$
$$(\textbf{if} (B) \{C; \textbf{while} (B)\{C\}\} \textbf{else skip}, \mathcal{R}_1, \mathcal{G}_1) \preceq_{\alpha; \zeta - \gamma} (\textbf{while} (B)\{C\}, \mathcal{R}_2, \mathcal{G}_2)$$

Loop Unrolling

(while
$$(B)\{C\}, \mathcal{R}_1, \mathcal{G}_1) \preceq_{\alpha; \zeta} \psi$$
 (while $(B)\{C\}, \mathcal{R}_2, \mathcal{G}_2$)

 $(\textbf{while} (B)\{C; \textbf{if} (B) \ C \textbf{ else skip}\}, \mathcal{R}_1, \mathcal{G}_1) \preceq_{\alpha; \zeta = \gamma} (\textbf{while} (B)\{C\}, \mathcal{R}_2, \mathcal{G}_2)$

Dead Code Elimination

$$\frac{(\mathbf{skip},\mathsf{Id},\mathsf{Id}) \preceq_{\alpha;\zeta \quad \gamma} (C,\mathsf{Id},\mathcal{G}) \qquad \mathsf{Sta}(\{\zeta,\gamma\},\langle\mathcal{R}_1,\mathcal{R}_2^*\rangle_\alpha)}{(\mathbf{skip},\mathcal{R}_1,\mathsf{Id}) \preceq_{\alpha;\zeta \quad \gamma} (C,\mathcal{R}_2,\mathcal{G})}$$

Intuitively $(\mathbf{skip}, \mathsf{Id}, \mathsf{Id}) \preceq_{\alpha; \zeta} \gamma (C, \mathsf{Id}, \mathcal{G})$ says that the code *C* can be eliminated in a sequential context where the initial and the final states satisfy ζ and γ respectively. If both ζ and γ are stable with respect to the interference from the environments \mathcal{R}_1 and \mathcal{R}_2 , then the code *C* can be eliminated in such a parallel context as well.

Redundancy Introduction

$$\frac{(c,\mathsf{Id},\mathcal{G}) \preceq_{\alpha;\zeta} \gamma}{(c,\mathcal{R}_1,\mathcal{G}) \preceq_{\alpha;\zeta} \gamma} (\mathbf{skip},\mathsf{Id},\mathsf{Id}) \qquad \mathsf{Sta}(\{\zeta,\gamma\},\langle\mathcal{R}_1,\mathcal{R}_2^*\rangle_\alpha)}{(c,\mathcal{R}_1,\mathcal{G}) \preceq_{\alpha;\zeta} \gamma} (\mathbf{skip},\mathcal{R}_2,\mathsf{Id})$$

As we lifted sequential dead code elimination, we can also lift sequential redundant code introduction to the concurrent setting, so long as the pre- and postconditions are stable with respect to the environments. Note that here c is a single instruction, because we should consider the interference from the environments at every intermediate state when introducing a sequence of redundant instructions.

5.2. Examples

With these rules, we can prove the correctness of many traditional compiler optimizations performed on concurrent programs in appropriate contexts. In this section, we give some examples of hoisting loop invariants, strength reduction, and induction variable elimination.

5.2.1. Invariant Hoisting. We first formally prove the example in Section 2.5. As we discussed, safely hoisting the invariant code t:=x+1 requires that the environment should not update x nor t.

$$\mathcal{R} \triangleq \{(\sigma, \sigma') \mid \sigma(\mathbf{x}) = \sigma'(\mathbf{x}) \land \sigma(\mathbf{t}) = \sigma'(\mathbf{t})\}$$

The guarantee of the program can be specified as arbitrary transitions. Since we only care about the values of i, n and x, the invariant relation α can be defined as

$$\alpha \triangleq \{(\sigma_1, \sigma) \mid \sigma_1(i) = \sigma(i) \land \sigma_1(n) = \sigma(n) \land \sigma_1(x) = \sigma(x)\}.$$

We do not need special pre- and postconditions, thus the correctness of the optimization is formalized as follows.

$$(C_1, \mathcal{R}, \mathsf{True}) \preceq_{\alpha: \alpha = \alpha} (C, \mathcal{R}, \mathsf{True})$$
 (5.1)

We could prove Eq. (5.1) directly by the RGSim definition and the operational semantics of the code. But shortly we give a more convenient proof using the optimization rules and the compositionality rules instead. We first prove the following by the deadcode-elimination and redundancy-introduction rules.

$$\begin{array}{l} (\texttt{t:=x+1},\mathcal{R},\mathsf{True}) \leq_{\alpha;\alpha} \gamma (\texttt{skip},\mathcal{R},\mathsf{True});\\ (\texttt{skip},\mathcal{R},\mathsf{True}) \leq_{\alpha:\gamma} n (\texttt{t:=x+1},\mathcal{R},\mathsf{True}), \end{array}$$

Here γ and η specify the states at the specific program points.

$$\gamma \stackrel{\text{\tiny def}}{=} \alpha \cap \{(\sigma_1, \sigma) \mid \sigma_1(t) = \sigma_1(x) + 1\}; \\ \eta \stackrel{\text{\tiny def}}{=} \gamma \cap \{(\sigma_1, \sigma) \mid \sigma(t) = \sigma(x) + 1\}.$$

Then by the compositionality rules SEQ and WHILE, we can get $(C'_1, \mathcal{R}, \text{True}) \preceq_{\alpha;\alpha \quad \alpha} (C', \mathcal{R}, \text{True})$ where C'_1 and C' result from adding **skip**s to C_1 and C.

C'_1 :	C':
t := x + 1;	skip;
while(i < n) {	<pre>while(i < n) {</pre>
skip;	t := x + 1;
i := i + t;	i := i + t;
}	}

Besides, from sequential-unit laws and compositionality rules SEQ and WHILE, we can prove $(C_1, \mathcal{R}, \text{True}) \preceq_{\alpha;\alpha} \alpha (C'_1, \mathcal{R}, \text{True})$ and $(C', \mathcal{R}, \text{True}) \preceq_{\alpha;\alpha} \alpha (C, \mathcal{R}, \text{True})$. Finally, by the TRANS rule, we can conclude Eq. (5.1), that is, the correctness of the optimization in appropriate contexts. Since the rely conditions only prohibit updates of x and t, we

can execute C_1 and C concurrently with other threads which update i and n or read x, still ensuring semantics preservation.

5.2.2. Strength Reduction and Induction Variable Elimination

Torrat Loval C.		Middle-Level C_1		
Target-Level C2		local i, k;		Source-Level C
local k, r;		i := 0;		local i;
k := 0;		k := 0;		i := 0;
r := 6*n; while(k <r) td="" {<=""><td>while(i<n) td="" {<=""><td>⇐</td><td>while(i<n) td="" {<=""></n)></td></n)></td></r)>	while(i <n) td="" {<=""><td>⇐</td><td>while(i<n) td="" {<=""></n)></td></n)>	⇐	while(i <n) td="" {<=""></n)>	
		x := x+k;		x := x+6*i;
x := x+k;		i := i+1;		i := i+1:
k := k+6;		k := k+6;		}
ł		1		

The source program C is first transformed to C_1 by strength reduction which introduces a local variable k and replaces multiplication by addition. The original induction variable i and the introduced local variable k cannot be updated by the environments. Then C_1 is transformed to the target C_2 by eliminating i and using the new induction variable k in the while-condition. We assume n and r will not be updated by the target environment, so we can compute the new boundary outside the loop. Next, we give the environments \mathcal{R} , \mathcal{R}_1 and \mathcal{R}_2 at the source, intermediate, and target levels respectively.

$$\begin{aligned} \mathcal{R} &\triangleq \{(\sigma, \sigma') \mid \sigma(\mathbf{i}) = \sigma'(\mathbf{i})\} \\ \mathcal{R}_1 &\triangleq \{(\sigma_1, \sigma'_1) \mid \sigma_1(\mathbf{i}) = \sigma'_1(\mathbf{i}) \land \sigma_1(\mathbf{k}) = \sigma'_1(\mathbf{k})\} \\ \mathcal{R}_2 &\triangleq \{(\sigma_2, \sigma'_2) \mid \sigma_2(\mathbf{k}) = \sigma'_2(\mathbf{k}) \land \sigma_2(\mathbf{r}) = \sigma'_2(\mathbf{r}) \land \sigma_2(\mathbf{n}) = \sigma'_2(\mathbf{n})\} \end{aligned}$$

For both transformations, we require that the common variables in the source and target have the same values. This is shown in the invariant relations α (for the transformation from C to C_1) and β (for the transformation from C_1 to C_2) next.

$$\alpha \triangleq \{(\sigma_1, \sigma) \mid \sigma_1(i) = \sigma(i) \land \sigma_1(n) = \sigma(n) \land \sigma_1(x) = \sigma(x)\}; \\ \beta \triangleq \{(\sigma_2, \sigma_1) \mid \sigma_2(k) = \sigma_1(k) \land \sigma_2(n) = \sigma_1(n) \land \sigma_2(x) = \sigma_1(x)\}$$

Thus we formalize the correctness of the two transformations as follows.

$$(C_2, \mathcal{R}_2, \mathsf{True}) \preceq_{\beta; \beta - \beta} (C_1, \mathcal{R}_1, \mathsf{True}), \ (C_1, \mathcal{R}_1, \mathsf{True}) \preceq_{\alpha; \alpha - \alpha} (C, \mathcal{R}, \mathsf{True})$$

They can be proved directly by the RGSim definition or by applying the optimization rules (the dead-code-elimination and redundancy-introduction rules). The proofs are similar to those for the previous example of invariant hoisting, and hence omitted here.

Afterwards, we can compose the proofs of these two transformations by the $\ensuremath{\mathsf{TRANS}}$ rule, and get

$$(C_2, \mathcal{R}_2, \mathsf{True}) \preceq_{\alpha \circ \beta; \alpha \circ \beta} {}_{\alpha \circ \beta} (C, \mathcal{R}, \mathsf{True}),$$

where $\alpha \circ \beta = \{(\sigma_2, \sigma) \mid \sigma_2(n) = \sigma(n) \land \sigma_2(x) = \sigma(x)\}$. That is, the optimization phases are correct when the source program is executed in an environment that does not change i nor n (as shown in \mathcal{R} and \mathcal{R}_2).

6. REFINEMENT-BASED VERIFICATION FOR CONCURRENT ALGORITHMS

The implementation of an abstract algorithm can be viewed as a transformation from an abstract operation to a concrete and executable program [Hoare 1972]. Verifying that the executable program refines the abstract operation gives us the correctness

```
A_1:
                          A_2:
   local d1;
                             local d2;
  d1 := 0;
                             d2 := 0;
   while (d1 = 0) {
                             while (d2 = 0) {
Δ
     atom{
                           0
                               atom{
       if (a = b)
                                  if (b = a)
                        d1 := 1;
                                    d2 := 1;
       if (a > b)
                                  if (b > a)
         a := a - b;
                                    b := b - a;
     }
                               }
   }
                             }
```

(a) source code

<i>C</i> ₁ :	C_2 :
local d1, t11, t12;	local d2, t21, t22;
d1 := 0;	d2 := 0;
while $(d1 = 0)$ {	while $(d2 = 0)$ {
0 t11 := a;	0 t21 := b;
1 t12 := b;	1 t22 := a;
2 if $(t11 = t12)$	$^{\parallel}$ 2 if (t21 = t22)
3 d1 := 1;	3 d2 := 1;
4 if (t11 > t12)	4 if $(t21 > t22)$
5 a := t11 - t12;	5 b := $t21 - t22;$
}	}

(b) target code

Fig. 8. Concurrent GCD.

of the implementation. In a concurrent setting, we can use RGSim to verify the finegrained implementation of an algorithm.

Similarly, RGSim also gives us a refinement-based proof method to verify the *atomicity* of concurrent object implementations. A concurrent object provides a set of methods which can be called in parallel by clients as the only way to access the object. We can define abstract atomic operations in a high-level language as specifications, and prove the concrete fine-grained implementations refine the corresponding atomic operations when executed in appropriate environments.

In this section, we discuss four examples to illustrate how we use RGSim to verify the concurrent objects and fine-grained implementation of abstract algorithms: a concurrent GCD algorithm (calculating greatest common divisors) [Feng 2009], the lockcoupling list [Herlihy and Shavit 2008], the nonblocking concurrent counter [Turon and Wand 2011], and Treiber's stack algorithm [Treiber 1986].

6.1. Concurrent GCD

We first prove the correctness of a concurrent GCD program in Figure 8(b). The program uses two threads to compute the Greatest Common Divisor (GCD) of the shared variables a and b. One thread executes C_1 which reads the values of a and b, but only updates a if a > b. The other thread executes C_2 which does the reverse. When a = b, the two threads terminate. This fine-grained GCD program is transformed from the program in Figure 8(a), where two threads atomically update a and b respectively. Here we use $atom\{\mathbb{C}\}$ to execute \mathbb{C} atomically. Its semantics follows RGSep [Vafeiadis 2008] (or see Section 6.2).

Our goal is to prove that the concrete and abstract GCD programs always obtain the same result, that is, $(C_1 \parallel C_2)$; print(a) and $(A_1 \parallel \parallel A_2)$; print(a) have the same outputs. We use print(a) at the two levels to print out the results after both threads complete their computations.

By soundness of RGSim and its compositionality, we only need to prove that the core computations for updating a (or b) are equivalent in C_1 and A_1 (or C_2 and A_2), that is, C_1^0 is equivalent to A_1^0 (and C_2^0 is equivalent to A_2^0), where C_1^0 (or C_2^0) denotes the code from line 0 to line 5 in C_1 (or C_2), and A_1^0 (or A_2^0) denotes the atomic block in A_1 (or A_2). It is natural to define the α relation as

 $\alpha \triangleq \{(\sigma, \Sigma) \mid \sigma(\mathsf{a}) = \Sigma(\mathsf{a}) \land \sigma(\mathsf{b}) = \Sigma(\mathsf{b}) \land \sigma(\mathsf{d}1) = \Sigma(\mathsf{d}1) \land \sigma(\mathsf{d}2) = \Sigma(\mathsf{d}2)\}.$

The threads' rely and guarantee conditions can be specified as follows, where the rely of one thread is just the guarantee of the other.

$$\begin{split} \mathcal{R}_1 &= \mathcal{G}_2 \ \triangleq \ \{(\sigma, \sigma') \mid \sigma'(\texttt{t11}) = \sigma(\texttt{t11}) \land \sigma'(\texttt{t12}) = \sigma(\texttt{t12}) \land \sigma'(\texttt{d1}) = \sigma(\texttt{d1}) \land \sigma'(\texttt{a}) = \sigma(\texttt{a}) \\ & \land (\sigma(\texttt{a}) \geq \sigma(\texttt{b}) \Rightarrow \sigma'(\texttt{b}) = \sigma(\texttt{b})) \} \\ \mathcal{R}_2 &= \mathcal{G}_1 \ \triangleq \ \{(\sigma, \sigma') \mid \sigma'(\texttt{t21}) = \sigma(\texttt{t21}) \land \sigma'(\texttt{t22}) = \sigma(\texttt{t22}) \land \sigma'(\texttt{d2}) = \sigma(\texttt{d2}) \land \sigma'(\texttt{b}) = \sigma(\texttt{b}) \\ & \land (\sigma(\texttt{b}) \geq \sigma(\texttt{a}) \Rightarrow \sigma'(\texttt{a}) = \sigma(\texttt{a})) \} \\ \mathbb{R}_1 &= \mathbb{G}_2 \ \triangleq \ \{(\Sigma, \Sigma') \mid \Sigma'(\texttt{d1}) = \Sigma(\texttt{d1}) \land \Sigma'(\texttt{a}) = \Sigma(\texttt{a}) \land (\Sigma(\texttt{a}) \geq \Sigma(\texttt{b}) \Rightarrow \Sigma'(\texttt{b}) = \Sigma(\texttt{b})) \} \\ \mathbb{R}_2 &= \mathbb{G}_1 \ \triangleq \ \{(\Sigma, \Sigma') \mid \Sigma'(\texttt{d2}) = \Sigma(\texttt{d2}) \land \Sigma'(\texttt{b}) = \Sigma(\texttt{b}) \land (\Sigma(\texttt{b}) \geq \Sigma(\texttt{a}) \Rightarrow \Sigma'(\texttt{a}) = \Sigma(\texttt{a})) \} \end{split}$$

Then we can operationally prove the RGSim relations between C_1^0 and A_1^0 (here α^{-1} is the inverse relation of α , as defined in Figure 6).

$$(C_1^0, \mathcal{R}_1, \mathcal{G}_1) \preceq_{\alpha; \alpha \quad \alpha} (A_1^0, \mathbb{R}_1, \mathbb{G}_1), \qquad (A_1^0, \mathbb{R}_1, \mathbb{G}_1) \preceq_{\alpha^{-1}; \alpha^{-1} \quad \alpha^{-1}} (C_1^0, \mathcal{R}_1, \mathcal{G}_1).$$

By the rules WHILE and SEQ, we get the RGS m relations between C_1 and A_2 .

$$(C_1, \mathcal{R}_1, \mathcal{G}_1) \preceq_{\alpha; \alpha \quad \alpha} (A_1, \mathbb{R}_1, \mathbb{G}_1), \qquad (A_1, \mathbb{R}_1, \mathbb{G}_1) \preceq_{\alpha^{-1}: \alpha^{-1} \quad \alpha^{-1}} (C_1, \mathcal{R}_1, \mathcal{G}_1)$$

Similarly, the relations hold between C_2 and A_2 .

$$(C_2, \mathcal{R}_2, \mathcal{G}_2) \preceq_{\alpha; \alpha \quad \alpha} (A_2, \mathbb{R}_2, \mathbb{G}_2), \qquad (A_2, \mathbb{R}_1, \mathbb{G}_1) \preceq_{\alpha^{-1}: \alpha^{-1} \quad \alpha^{-1}} (C_2, \mathcal{R}_1, \mathcal{G}_1)$$

When C_1 and C_2 (or A_1 and A_2) are parallel composed to compute the GCD together, the environment of the whole GCD program should be the identity transition set Id because the shared variables a and b cannot be modified when $C_1 \parallel C_2$ is computing their GCD. Its guarantee is just specified as True, a set of all the possible state transitions. We can prove that both (print(a), Id, True) $\leq_{\alpha;\alpha} \alpha$ (print(a), Id, True) and the reverse direction hold. Then by the rules PAR and SEQ, we can get

$$((C_1 || C_2); \texttt{print}(a), \texttt{Id}, \texttt{True}) \preceq_{\alpha; \alpha \ \alpha} ((A_1 || A_2); \texttt{print}(a), \texttt{Id}, \texttt{True}), \texttt{Id}, \texttt{True})$$

and also the reverse direction. By the soundness of RGSim (Theorem 4.3) we obtain the final result.

$$(C_1 \parallel C_2); \texttt{print}(\texttt{a}) \approx_{\mathbf{T}} (A_1 \parallel A_2); \texttt{print}(\texttt{a}),$$

This is for any **T** that respects α .

Thus we have proved that the concrete fine-grained and the abstract coarse-grained GCD programs can obtain the same results from the same inputs. It is not difficult to find that the abstract program really computes the GCD of a and b. So we can conclude

ACM Transactions on Programming Languages and Systems, Vol. 36, No. 1, Article 3, Publication date: March 2014.

(a) an abstract set

```
add(e):
                       rmv(e):
                           local x,y,z,v;
   local x,y,z,u;
                         0 <x := Head;>
 0 < x := Head;>
                         1 lock(x);
 1 lock(x);
                         2 <y := x.next;>
 2 <z := x.next;>
                         3 <v := y.data;>
 3 <u := z.data;>
                         4 while (v < e) {
 4 while (u < e) {
                         5
                             lock(y);
 5
     lock(z);
                         6
                             unlock(x);
 6
     unlock(x);
                         7
                             x := y;
 7
     x := z;
                         8
                             <y := x.next;>
 8
     <z := x.next;>
                         9
                             <v := y.data;>
 9
     <u := z.data;>
                           }
   }
                        10 if (v = e) {
10 if (u != e) {
                        11
                             lock(y);
     y := new();
11
                        12
                             <z := y.next;>
12
     y.lock := 0;
                        13
                             <x.next := z;>
13
     y.data := e;
                        14
                             unlock(x);
14
     y.next := z;
                        15
                             free(y);
15
     <x.next := y;>
                           } else {
   }
                        16
                             unlock(x);
16 unlock(x);
                           }
```

(b) the lock-coupling list-based set

Fig. 9. The set object.

that the concrete program computes their GCD as well. This example shows a way to verify a complicated program by proving that it is equivalent to a simpler program and then verifying the simpler program.

6.2. Lock-Coupling List

In this section, we prove the atomicity of the lock-coupling list-based implementation for the set object. In Figure 9(a) we define two atomic set operations, ADD(e) and RMV(e). Figure 9(b) gives a concrete implementation of the set object using a lock-coupling list. Partial correctness and atomicity of the algorithm has been verified before [Vafeiadis 2008; Vafeiadis and Parkinson 2007]. Here we show that its atomicity can also be verified using our RGSim by proving the low-level methods refine the corresponding abstract operations. We will discuss the key difference between the previous proofs and ours in Section 8.

To support dynamically allocated memory and ownership transfers, we split the states into shared and thread-local parts. We first take the generic languages in Figure 2, and instantiate the high-level program states as follows. The state Σ consists of shared memory M_s (where the object resides) and a thread pool Π , which is a mapping from thread identifiers (t \in *ThrdID*) to their memory M_l . The low-level state σ is defined similarly. We use m_s , m_l and π to represent the low-level shared memory, thread-local memory, and the thread pool respectively.

We show the high-level and low-level languages and the operational semantics in Figure 10. To allow ownership transfers between the shared memory and thread-local memory, we use **atom**{ \mathbb{C} }_A (or $\langle C \rangle_A$ at the low level) to convert the shared memory to local and then execute \mathbb{C} (or *C*) atomically. Following RGSep [Vafeiadis and Parkinson 2007], an abstract transition $\mathbb{A} \in \mathcal{P}(HMem \times HMem)$ (or $\mathcal{A} \in \mathcal{P}(LMem \times LMem)$) is used to specify the effects of the atomic operation over the shared memory, which allows us to split the resulting state back into shared and local when we exit the atomic block.² The atomic blocks are instantiations of the generic primitive operations c (or *c*) in Figure 2. We omit the annotations \mathbb{A} and \mathcal{A} in Figure 9, which are the same as the corresponding guarantees in Figure 11, as we will explain next.

In Figure 9, the abstract set is implemented by an ordered singly-linked list pointed to by a shared variable Head, with two sentinel nodes at the two ends of the list containing the values MIN_VAL and MAX_VAL respectively. Each list node is associated with a lock. Traversing the list uses "hand-over-hand" locking: the lock on one node is not released until its successor is locked. add(e) inserts a new node with value e in the appropriate position while holding the lock of its predecessor. rmv(e) redirects the predecessor's pointer while both the node to be removed and its predecessor are locked. Note that lock(x) and unlock(x) are instantiations of c. Their semantics has been explained in Section 3.1.

We define the α relation, the guarantees and the relies in Figure 11. The predicate $list(x,A)(m_s)$ represents a singly-linked list in the shared memory m_s at the location x, whose values form the sequence A. Then the mapping shared_map between the low-level and the high-level shared memory is defined by only concerning about the value sequence on the list: the concrete list should be sorted and its elements constitute the abstract set. For a thread t's local memory of the two levels, we require that the values of e are the same and enough local space is provided for add(e) and rmv(e), as defined in the mapping local_map. Then α relates the shared memory by shared_map and the local memory of each thread t by local_map.

Before defining the rely and guarantee relations, we first introduce some syntactic

sugar in Figure 11(b). We use $x \mapsto (n, v, y)$ and $x \stackrel{l}{\mapsto} (n, v, y)$ for nodes in the lowlevel shared memory m_s and the local memory m_l of the current thread respectively. *Itrue* means the thread-local memory m_l is arbitrary. The separating conjunction p * qmeans p and q hold on disjoint memory. The action $p \ltimes_t q$ represents the update of some memory (m_l, m_s) satisfying p to some memory satisfying q, and the memory of the threads other than the current thread t is unchanged. We overload the notations to the high-level machine, and use $x \mapsto v$ to mean the value of x is v in the high-level shared memory M_s .

The atomic actions of the algorithm are specified by \mathcal{G}_{lock} , \mathcal{G}_{unlock} , \mathcal{G}_{add} , \mathcal{G}_{rmv} and \mathcal{G}_{local} respectively, which are all parametrized with a thread identifier t. For example, $\mathcal{G}_{rmv}(t)$ says that when holding the locks of the node y and its predecessor x, we can transfer the node y from the shared memory to the thread's local memory. This corresponds

²It is easy to prove that if the program W (or W) does not abort, then the event trace set of W (or W) under the instrumented semantics is the same as the event trace set of W (or W) under a standard (flat) operational semantics where we erase the annotations A (or A) and merge the local and shared memory. Thus by proving the event-trace refinement under the instrumented semantics, we can get the event-trace refinement under the flat semantics.

(HStmts)	\mathbb{C}	::=	$skip \mid c \mid atom\{\mathbb{C}\}_A$	$ \mathbb{C}_1;;\mathbb{C}_2$
			$\textit{if} ~ \mathbb{B} \textit{ then} ~ \mathbb{C}_1 \textit{ else } \mathbb{C}_2 ~ \mid ~$	while ${\mathbb B}$ do ${\mathbb C}$
(HProg)	\mathbb{W}	::=	$t_1.\mathbb{C}_1 \ldots t_n.\mathbb{C}_n$	$(ThrdID) t \in Nat$
(HMem)	M_s, M_l	\in	$(Loc \cup PVar) \rightharpoonup HVal$	
(HThrds)	П	∈	$\mathit{ThrdID} ightarrow \mathit{HMem}$	
(HState)	Σ	∈	HThrds imes HMem	
(HAtomG)	A	∈	$\mathcal{P}(HMem \times HMem)$	

(a) the high-level language for abstract operations

(LStmts)	C	=:: 	$ \begin{aligned} \mathbf{skip} &\mid c \mid \langle C \rangle_{\mathcal{A}} \mid C_1; C_2 \\ \mathbf{if} \left(B \right) C_1 \ \mathbf{else} \ C_2 \mid \mathbf{while} \ \left(B \right) C \end{aligned} $
(LProg)	W	::=	$t_1.C_1 \parallel \ldots \parallel t_n.C_n$
(LMem)	m_s, m_l	∈	$(Loc \cup PVar) \rightarrow LVal$
(LThrds)	π	\in	$ThrdID \rightarrow LMem$
(LState)	σ	∈	LThrds imes LMem
(LAtomG)	\mathcal{A}	\in	$\mathcal{P}(LMem \times LMem)$

(b) the low-level language for concrete implementations

$$\frac{(\mathbb{C}, (\Pi \uplus \{\mathbf{t} \quad (M_{l} \uplus M_{s})\}, \emptyset)) \longrightarrow_{\mathbf{t}}^{*} (\mathbf{skip}, (\Pi \uplus \{\mathbf{t} \quad M_{l}^{\prime\prime}\}, \emptyset)) \qquad M_{l}^{\prime\prime} = M_{l}^{\prime} \uplus M_{s}^{\prime\prime} \quad (M_{s}, M_{s}^{\prime\prime}) \in \mathbb{A} }{(atom\{\mathbb{C}\}_{\mathbb{A}}, (\Pi \uplus \{\mathbf{t} \quad M_{l}\}, M_{s})) \longrightarrow_{\mathbf{t}} (\mathbf{skip}, (\Pi \uplus \{\mathbf{t} \quad M_{l}^{\prime\prime}\}, M_{s}^{\prime\prime}))} \\ \frac{(\mathbb{C}, (\Pi \uplus \{\mathbf{t} \quad (M_{l} \uplus M_{s})\}, \emptyset)) \longrightarrow_{\mathbf{t}}^{*} (\mathbf{skip}, (\Pi \uplus \{\mathbf{t} \quad M_{l}^{\prime\prime}\}, \emptyset))}{(atom\{\mathbb{C}\}_{\mathbb{A}}, (\Pi \uplus \{\mathbf{t} \quad M_{l}\}, M_{s})) \longrightarrow_{\mathbf{t}}^{*} (\mathbf{skip}, (\Pi \uplus \{\mathbf{t} \quad M_{l}^{\prime\prime}\}, \emptyset))} \\ \frac{(\mathbb{C}, (\Pi \uplus \{\mathbf{t} \quad (M_{l} \uplus M_{s})\}, \emptyset)) \longrightarrow_{\mathbf{t}}^{*} (\mathbf{skip}, (\Pi \uplus \{\mathbf{t} \quad M_{l}\}, M_{s})) \longrightarrow_{\mathbf{t}} \mathbf{abort}}{(atom\{\mathbb{C}\}_{\mathbb{A}}, (\Pi \uplus \{\mathbf{t} \quad M_{l}\}, M_{s})) \longrightarrow_{\mathbf{t}} \mathbf{abort}} \qquad (\mathbb{C}_{i}, \Sigma) \longrightarrow_{\mathbf{t}_{i}} \mathbf{abort} \\ \frac{(\mathbb{C}, (\Pi \uplus \{\mathbf{t} \quad (M_{l} \uplus M_{s})\}, \emptyset)) \longrightarrow_{\mathbf{t}}^{*} \mathbf{abort}}{(\mathbf{t}_{1}.\mathbb{C}_{1} \parallel \dots \mathbf{t}_{i}.\mathbb{C}_{i} \dots \parallel \mathbf{t}_{n}.\mathbb{C}_{n}, \Sigma) \longrightarrow \mathbf{abort}} \qquad (\mathbb{C}_{i}, \Sigma) \longrightarrow_{\mathbf{t}_{i}} \mathbb{C}_{i} \dots \mathbb{C}_{i} \otimes \mathbb{C}_{i} \otimes$$

(c) selected operational semantics rules of the high-level language

Fig. 10. The languages for concurrent objects.

to the action performed by the code of line 13 in rmv(e) in Figure 9. Every thread t is executed in the environment that any other thread t' can only perform those five actions, as defined in $\mathcal{R}(t)$. Similarly, the high-level $\mathbb{G}(t)$ and $\mathbb{R}(t)$ are defined according to the abstract ADD(e) and RMV(e). The relies and guarantees are almost the same as those in the proofs in RGSep [Vafeiadis 2008].

$$\begin{split} & \text{list}(x,A) \triangleq \lambda m_s. \ (m_s = \emptyset \land x = \textbf{null} \land A = \epsilon) \\ & \lor (\exists m'_s, v, y, A'. \ m_s = m'_s \uplus \{x \quad (_, v, y)\} \land A = v :: A' \land \text{list}(y, A')(m'_s)) \end{split} \\ & \text{sorted}(A) \triangleq \begin{cases} \textbf{true} & \text{if } A = \epsilon \lor A = a :: \epsilon \\ (a < b) \land \text{sorted}(b :: A') & \text{if } A = a :: b :: A' \end{cases} \\ & \text{elems}(A) \triangleq \begin{cases} \emptyset & \text{if } A = \epsilon \\ \{a\} \cup \text{elems}(A') & \text{if } A = a :: A' \end{cases} \\ & \text{shared_map}(m_s, M_s) \triangleq \exists m'_s, A, x. \ m_s = m'_s \uplus \{\text{Head} \ x\} \land \text{list}(x, \text{MIN_VAL} :: A :: \text{MAX_VAL})(m'_s) \\ & \land \text{sorted}(A) \land (\text{elems}(A) = M_s(S)) \end{cases} \\ & \text{local_map}(m_l, M_l) \triangleq m_l(e) = M_l(e) \land \exists m'_l. \ m_l = m'_l \uplus \{x \ ..., y \ ..., z \ ..., u \ ..., v \ ...\} \\ & \alpha \triangleq \{((\pi, m_s), (\Pi, M_s)) \mid \text{shared_map}(m_s, M_s) \land \forall t \in dom(\Pi). \text{ local_map}(\pi(t), \Pi(t))) \} \end{cases} \end{split}$$

(a) the α relation

$$\begin{split} x \mapsto (n, v, y) &\triangleq \lambda(m_l, m_s). \ (dom(m_l) = \emptyset) \land (m_s = \{x \quad (n, v, y)\}) \\ x \stackrel{l}{\mapsto} (n, v, y) &\triangleq \lambda(m_l, m_s). \ (m_l = \{x \quad (n, v, y)\}) \land (dom(m_s) = \emptyset) \\ ltrue &\triangleq \lambda(m_l, m_s). \ (dom(m_s) = \emptyset) \\ p * q &\triangleq \lambda(m_l, m_s). \ \exists m'_l, m'_s, m''_l, m''_s. \ p(m'_l, m'_s) \land q(m''_l, m''_s) \land (m_l = m'_l \uplus m''_l) \land (m_s = m'_s \uplus m''_s) \\ p \ltimes_{\mathsf{t}} q &\triangleq \{((\pi \uplus \{\mathsf{t} \ m_l\}, m_s), (\pi \uplus \{\mathsf{t} \ m'_l\}, m'_s)) \\ & \mid \exists m_{l1}, m_{s1}, m_{l2}, m_{s2}, m'_{l1}, m'_{s1} \cdot p(m_{l1}, m_{s1}) \land q(m'_{l1}, m'_{s1}) \\ & \land (m_l = m_{l1} \uplus m_{l2}) \land (m_s = m_{s1} \uplus m_{s2}) \land (m'_l = m'_{l1} \uplus m_{l2}) \land (m'_s = m'_{s1} \uplus m_{s2}) \} \end{split}$$

$$\begin{array}{ll} x \mapsto v & \triangleq & \lambda(M_l, M_s). \ (dom(M_l) = \emptyset) \land (M_s = \{x \quad v\}) \\ ltrue & \triangleq & \lambda(M_l, M_s). \ (dom(M_s) = \emptyset) \end{array}$$

(b) syntactic sugar (where $p, q \in LMem \times LMem \rightarrow Prop$)

(c) rely and guarantee relations

Fig. 11. Auxiliary definitions and specifications for the lock-coupling list.

We can prove that for any thread t, the following hold.

$$(\text{add}(e), \mathcal{R}(t), \mathcal{G}(t)) \preceq^{t}_{\alpha;\alpha} _{\alpha} (\text{ADD}(e), \mathbb{R}(t), \mathbb{G}(t)); \\ (\text{rmv}(e), \mathcal{R}(t), \mathcal{G}(t)) \preceq^{t}_{\alpha;\alpha} _{\alpha} (\text{RMV}(e), \mathbb{R}(t), \mathbb{G}(t)).$$

Here $\leq_{\alpha;\alpha}^{t} _{\alpha}$ is the RGSim relation in Definition 4.2 with the transitions \longrightarrow being replaced by \longrightarrow_{t} (defined in Figure 10(c)). The proofs are done operationally based

ACM Transactions on Programming Languages and Systems, Vol. 36, No. 1, Article 3, Publication date: March 2014.

```
atom{ x := x+1; }
(a) source code INC(x)
local d, t;
d := 0;
while (d = 0) {
    <t := x;>
    d := cas(&x,t,t+1);
}
```

(b) target code inc(x)

Fig. 12. The atomic and nonblocking counters.

on the definition of RGSim. We analyze the implementation step by step and find out the instructions which correspond to the high-level single atomic steps (i.e., the *linearization points*). For the add(e) operation, since we require the elements in the concrete list are those in the abstract set, we can pick line 15 as the linearization point of a successful call where the new node containing the value e is inserted into the list. For unsuccessful calls (e is already in the set), we choose lines 3 and 9 where the value e is read from an existing list node. Similarly, for rmv(e), we choose line 13 (for successful calls) and lines 3 and 9 (for unsuccessful calls) as linearization points. We omit the detailed proofs here.

By the compositionality and the soundness of RGSim, we know that the fine-grained operations (under the parallel environment \mathcal{R}) are simulated by the corresponding atomic operations (under the high-level environment \mathbb{R}), while \mathcal{R} and \mathbb{R} say all accesses to the set must be done through the add and remove operations. This gives us the atomicity of the concurrent implementation of the set object.

6.3. Nonblocking Counter

The next example (in Figure 12) is counters which increase the value of a shared variable x atomically. The basic requirement is that the counter should not miss any increment when several threads update x concurrently. A simple abstract counter INC(x) increases the value of x in a coarse-grained atomic block. The concrete implementation inc(x) uses the Compare-And-Swap (CAS) instruction $cas(\&x, t_1, t_2)$, which reads the value from the location of x, compares it with an expected value t_1 , writes out a new value t_2 if the two match, and returns whether the update succeeds. Next we use RGSim to prove the atomicity of inc(x).

We first define the α relation between low-level and high-level states, where only the values of x are concerned.

$$\alpha \triangleq \{((\pi, m_s), (\Pi, M_s)) \mid m_s(\mathbf{x}) = M_s(\mathbf{x})\}$$

We let the pre- and postconditions be the same as the invariant α . Both inc(x) and INC(x) guarantee that a thread t only updates its local variables and/or increases the values of x. The rely conditions of thread t allow any other thread t' to update x and thread-local variables of t'. Here we use the syntactic sugar in Figure 11 to define the rely and guarantee relations.

$\mathcal{G}(t)$	≜	$(\exists n. (\mathbf{x} \mapsto n * ltrue) \ltimes_{t} (\mathbf{x} \mapsto n + 1 * ltrue)) \lor \mathcal{G}_{local}(t)$	$\mathcal{R}(t) \triangleq \bigcup_{t' \neq t} \mathcal{G}(t')$
G(t)	\triangleq	$(\exists n. (\mathbf{x} \mapsto n) \ltimes_{t} (\mathbf{x} \mapsto n+1)) \lor \mathbb{G}_{local}(t)$	$\mathbb{R}(t) \triangleq \bigcup_{t' \neq t} \mathbb{G}(t')$

Then we can prove the RGSim relation holds.

$$(\operatorname{inc}(\mathbf{x}), \mathcal{R}(t), \mathcal{G}(t)) \preceq^{t}_{\alpha;\alpha} \alpha (\operatorname{INC}(\mathbf{x}), \mathbb{R}(t), \mathbb{G}(t))$$

It says that the fine-grained inc(x) does not have more behaviors than the atomic INC(x) in any environment, that is, inc(x) has atomicity. The proof is done operationally based on the RGSim definition. We find out the corresponding program points

in inc(x) and INC(x), and prove they are related no matter what the environments do. Also we can prove $(INC(x), \mathbb{R}(t), \mathbb{G}(t)) \preceq^t_{\alpha^{-1}:\alpha^{-1} \quad \alpha^{-1}} (inc(x), \mathcal{R}(t), \mathcal{G}(t))$, which says the implementation inc(x) has all the behaviors of INC(x). Thus inc(x) and INC(x) behave just the same.

As a simple illustration of the atomicity, we go on to show that the nonblocking inc(x) can be used by two threads concurrently without missing any increment, as if x was updated by the threads one after the other. Formally, we prove that $(inc(x); print(x)) \parallel (inc(x); print(x))$ and $(INC(x); print(x)) \parallel (INC(x); print(x))$ have the same observable event traces when the initial values of x are the same.

We can prove that $(print(x), \mathcal{R}(t), \mathcal{G}(t)) \preceq_{\alpha;\alpha \ \alpha} (print(x), \mathbb{R}(t), \mathbb{G}(t))$ and the reverse direction holds. Then by the rules SEQ, PAR and CONSEQ, we can get both

> ((inc(x); print(x)) || (inc(x); print(x)), ld, True) $\leq_{\alpha:\alpha} \alpha$ ((INC(x); print(x)) $\||$ (INC(x); print(x)), ld, True)

and the reverse direction. By the soundness of RGSim (Theorem 4.3), we know they are e-trace equivalent, and hence the transformation is correct.

 $(inc(x); print(x)) \parallel (inc(x); print(x)) \approx_{\mathbf{T}} (INC(x); print(x)) \parallel (INC(x); print(x))$

This is for any **T** that respects α . That is, no matter how the two nonblocking threads interleave, they complete their operations as if both of them were executing the abstract atomic counter.

Incrementing several shared variables. We have verified the transformation from INC(x) to inc(x) without caring about other shared resource. The FRAME rule allows us to combine several verified transformations together which work on disjoint parts of states without redoing the proofs.

For example, suppose we have another shared variable y which can be incremented as well as x. It is easy to see: $(inc(y), \mathcal{R}_1(t), \mathcal{G}_1(t)) \preceq^t_{\alpha_1;\alpha_1 \quad \alpha_1} (INC(y), \mathbb{R}_1(t), \mathbb{G}_1(t))$, where $\alpha_1 \triangleq \{((\pi, m_s), (\Pi, M_s)) \mid m_s(y) = M_s(y)\}$ and $\mathcal{R}_1(t), \mathcal{G}_1(t), \mathbb{R}_1(t)$ and $\mathbb{G}_1(t)$ are defined similarly as $\mathcal{R}(t)$, $\mathcal{G}(t)$, $\mathbb{R}(t)$ and $\mathbb{G}(t)$ except all the occurrences of x are replaced by y.

By the rules FRAME and SEQ, we can get

$$(\operatorname{inc}(x); \operatorname{inc}(y); \operatorname{print}(x), \mathcal{R}(t) \uplus \mathcal{R}_{1}(t), \mathcal{G}(t) \uplus \mathcal{G}_{1}(t)) \leq_{\beta;\beta-\beta}^{t} (\operatorname{INC}(x); \operatorname{INC}(y); \operatorname{print}(x), \mathbb{R}(t) \uplus \mathbb{R}_{1}(t), \mathbb{G}(t) \uplus \mathbb{G}_{1}(t))$$

$$(6.1)$$

where $\beta \triangleq \alpha \uplus \alpha_1 = \{((\pi, m_s), (\Pi, M_s)) \mid m_s(x) = M_s(x) \land m_s(y) = M_s(y)\}$, the rely conditions ensure that the environments cannot update any local variable used in incrementing x nor y, and the guarantees just say that the programs increment x or y or update local variables.

Similarly, we can get the reverse direction of Eq. (6.1). Then by the soundness of RGSim, we can conclude the combined transformation is correct

$$inc(x); inc(y); print(x) \approx_{\mathbf{T}} INC(x); INC(y); print(x),$$

for any **T** that respects β .

```
PUSH(v):
                     POP():
                         local r;
                        atom {
                      0
                            if (A = \epsilon) {
                              r := EMPTY;
   atom {
0
                            }else {
      A := v::A;
                              r := head(A);
    3
                              A := tail(A);
                            }
                          }
                          return r;
```

(a) an abstract stack

```
push(v) :
                            pop():
                                local r, d, x, t;
                               d := 0;
                             1
                               while (d = 0) {
   local d, x, t;
                            2
                                  <t := S;>
0 x := new Cell();
                            3
                                  if (t = null) {
1
  x.data := v;
                             4
                                    r := EMPTY;
2
  d := 0;
                            5
                                    d := 1;
3
   while (d = 0) {
                                  }else {
4
      <t := S;>
                            6
                                    r := t.data;
5
      x.next := t;
                            7
                                    x := t.next;
6
      d := cas(\&S,t,x);
                            8
                                    d := cas(\&S,t,x);
   }
                                  }
                                }
                                return r;
```

(b) Treiber's nonblocking implementation

Fig. 13. The stack object.

6.4. Treiber's Nonblocking Stack

The last example is to verify the atomicity of Treiber's nonblocking stack. The stack object provides two operations in its interface. The abstract PUSH(v) and POP(), defined in Figure 13(a), atomically operate on a value sequence. We implement the abstract stack by a singly-linked list pointed to by a shared variable S, and PUSH(v) and POP() by the nonblocking code push(v) and pop() respectively. As shown in Figure 13(b), the nonblocking implementation uses CAS instructions to obtain fine-grained atomicity.

We use RGSim to prove the atomicity of the nonblocking stack, that is, push(v) refines PUSH(v) and pop() refines POP() when they are executed in appropriate environments.

We define the α relation, the guarantees and the relies in Figure 14. The mapping shared_map between the low-level and the high-level shared memory is defined by only considering the value sequence on the stack. It requires that the concrete shared memory m_s contains a submemory \hat{m}_s of a linked list as the stack, and the concrete stack has the same value sequence as the abstract one. As in the lock-coupling list example

 $\begin{aligned} & \mathsf{shared_map}(m_s, M_s) \triangleq \exists \widehat{m}_s. \ \mathsf{list}(m_s(\mathtt{S}), M_s(\mathtt{A}))(\widehat{m}_s) \land \widehat{m}_s \subseteq m_s \backslash \{\mathtt{S}\} \\ & \mathsf{local_map}(m_l, M_l) \triangleq m_l(\mathtt{v}) = M_l(\mathtt{v}) \land \exists m_l'. \ m_l = m_l' \uplus \{\mathtt{d} \ _, \mathtt{x} \ _, \mathtt{t} \ _, \mathtt{r} \ _\} \\ & \alpha \triangleq \{((\pi, m_s), (\Pi, M_s)) \mid \mathsf{shared_map}(m_s, M_s) \land \forall \mathtt{t} \in dom(\Pi). \ \mathsf{local_map}(\pi(\mathtt{t}), \Pi(\mathtt{t}))\} \end{aligned}$

 $\begin{array}{lll} \mathcal{G}_{\mathrm{push}}(\mathsf{t}) &\triangleq \exists v, x, y. \ (\mathsf{S} \mapsto y \ast x \stackrel{l}{\mapsto} (v, y) \ast \mathit{ltrue}) \ltimes_{\mathsf{t}} (\mathsf{S} \mapsto x \ast x \mapsto (v, y) \ast \mathit{ltrue}) \\ \mathcal{G}_{\mathrm{pop}}(\mathsf{t}) &\triangleq \exists x, v, y. \ (\mathsf{S} \mapsto x \ast x \mapsto (v, y) \ast \mathit{ltrue}) \ltimes_{\mathsf{t}} (\mathsf{S} \mapsto y \ast x \mapsto (v, y) \ast \mathit{ltrue}) \\ \mathcal{G}(\mathsf{t}) &\triangleq \mathcal{G}_{\mathrm{push}}(\mathsf{t}) \cup \mathcal{G}_{\mathrm{pop}}(\mathsf{t}) \cup \mathcal{G}_{\mathrm{local}}(\mathsf{t}) \\ \end{array}$

$$\begin{array}{ll} \mathbb{G}_{push}(t) \triangleq \exists A, v. \ ((A \mapsto A) * \textit{ltrue}) \ltimes_{t} ((A \mapsto v :: A) * \textit{ltrue}) \\ \mathbb{G}_{pop}(t) & \triangleq \exists A, v. \ ((A \mapsto v :: A) * \textit{ltrue}) \ltimes_{t} ((A \mapsto A) * \textit{ltrue}) \\ \mathbb{G}(t) & \triangleq \mathbb{G}_{push}(t) \cup \mathbb{G}_{pop}(t) \cup \mathbb{G}_{local}(t) \\ \end{array}$$

Fig. 14. Auxiliary definitions and specifications for the nonblocking stack.

in Section 6.2, we use the predicate $list(x,A)(\widehat{m}_s)$ to represent a singly-linked list in the shared memory \widehat{m}_s whose head node's address is x and values form a sequence A. Since S is a shared variable containing the address of the top node, it itself is not in the domain of \widehat{m}_s . For the local memory, local_map defines the mapping of each thread. The value of v in the low-level local memory should be the same as in the high-level local memory, and the low-level local memory should provide enough additional space needed by the object operations (i.e., the local variables d, x, t and r). Then α relates the shared memory by shared_map and the local memory of each thread t by local_map.

Each thread guarantees that it performs push, pop, and local operations only, and its environment includes the operations made by all the other threads. The guarantees reflect the ownership transfers in push and pop operations. For example, $\mathcal{G}_{push}(t)$ says that the node x is transferred from the thread-local memory to the shared memory. The definitions use the syntactic sugar in Figure 11.

We could prove the nonblocking stack operations are simulated by the corresponding atomic operations.

$$(\operatorname{push}(v), \mathcal{R}(t), \mathcal{G}(t)) \preceq^{t}_{\alpha;\alpha \quad \alpha} (\operatorname{PUSH}(v), \mathbb{R}(t), \mathbb{G}(t)); (r := \operatorname{pop}(), \mathcal{R}(t), \mathcal{G}(t)) \preceq^{t}_{\alpha;\alpha \quad \alpha} (r := \operatorname{POP}(), \mathbb{R}(t), \mathbb{G}(t)).$$

This gives us the atomicity of the nonblocking implementation of the stack object.

7. VERIFYING CONCURRENT GARBAGE COLLECTORS

In this section, we explain in detail how to reduce the problem of verifying concurrent garbage collectors to transformation verification, and use RGSim to develop a general GC verification framework. We apply the framework to prove the correctness of the Boehm et al. concurrent GC algorithm [Boehm et al. 1991].

7.1. Correctness of Concurrent GCs

A concurrent GC is executed by a dedicated thread and performs the collection work in parallel with user threads (mutators), which access the shared heap via read, write, and allocation operations. To ensure that the GC and the mutators share a coherent view of the heap, the heap operations from mutators may be instrumented with extra operations, which provide an interaction mechanism to allow arbitrary mutators to cooperate with the GC. These instrumented heap operations are called barriers (e.g., read barriers, write barriers, and allocation barriers).

The GC thread and the barriers constitute a concurrent garbage collecting system, which provides a higher-level user-friendly programming model for garbage-collected languages (e.g., Java). In this high-level model, programmers feel they access the heap

using regular memory operations, and are freed from manually disposing objects that are no longer in use. They do not need to consider the implementation details of the GC and the existence of barriers.

We could verify the GC system by using a Hoare-style logic to prove that the GC thread and the barriers satisfy their specifications. However, we say this is an indirect approach because it is unclear if the specified correct behaviors would indeed preserve the mutators' intended behaviors and generate the abstract view for high-level programmers. Usually this part is examined by experts and then trusted.

Here we propose a more direct approach. We view a concurrent garbage collecting system as a transformation **T** from a high-level garbage-collected language to a low-level language. A standard atomic memory operation at the source level is transformed into the corresponding barrier code at the target level. In the source level, we assume there is an *abstract GC thread* that magically turns unreachable objects into reusable memory. The abstract collector *AbsGC* is transformed into the concrete GC code $C_{\rm gc}$ running concurrently with the target mutators. That is,

$$\mathbf{T}(\mathsf{t}_{\mathsf{gc}}AbsGC ||| \mathsf{t}_1.\mathbb{C}_1 ||| \dots ||| \mathsf{t}_n.\mathbb{C}_n) \triangleq \mathsf{t}_{\mathsf{gc}}.C_{\mathsf{gc}} || \mathsf{t}_1.\mathbf{T}(\mathbb{C}_1) || \dots || \mathsf{t}_n.\mathbf{T}(\mathbb{C}_n),$$

where $\mathbf{T}(\mathbb{C})$ simply translates some memory access instructions in \mathbb{C} into the corresponding barriers, and leaves the rest unchanged. Note that here we introduce an abstract GC and assume a finite memory at the source level. This is because at the target level we assume a finite memory to model the real machine; and if the source-level memory is infinite, the bijective mapping between the memory at the two levels would become much complicated.

Then we reduce the correctness of the concurrent garbage collecting system to Correct(T), saying that any mutator program will not have unexpected behaviors when executed using this system.

7.2. A General Verification Framework

The compositionality of RGSim allows us to develop a general framework to prove Correct(T), which is much more difficult using monolithic proof methods. By the parallel compositionality of RGSim (the PAR rule in Figure 7), we can decompose the refinement proofs into proofs for the GC thread and each mutator thread. For a mutator thread, we can further decompose the refinement proof into proof for each primitive instruction, using the compositionality of RGSim (the rules SEQ, IF and WHILE in Figure 7).

Verifying the GC. The semantics of the abstract GC thread can be defined by a binary state predicate AbsGCStep.

$$\frac{(\Sigma, \Sigma') \in \mathsf{AbsGCStep}}{(\mathsf{t}_{\mathsf{gc}}AbsGC, \Sigma) \longrightarrow (\mathsf{t}_{\mathsf{gc}}AbsGC, \Sigma')}$$

That is, the abstract GC thread always makes AbsGCStep to change the high-level state. We can choose different AbsGCStep for different GCs, but usually AbsGCStep guarantees not modifying reachable objects in the heap.

Thus for the GC thread, we need to show that $C_{\rm gc}$ is simulated by AbsGC when executed in their environments. This can be reduced to unary rely-guarantee reasoning about $C_{\rm gc}$ by proving $\mathcal{R}_{\rm gc}$; $\mathcal{G}_{\rm gc} \vdash \{p_{\rm gc}\}C_{\rm gc}\{q_{\rm gc}\}$ in a standard rely-guarantee logic with proper $\mathcal{R}_{\rm gc}$, $\mathcal{G}_{\rm gc}$, $p_{\rm gc}$ and $q_{\rm gc}$, as long as $\mathcal{G}_{\rm gc}$ is a concrete representation of AbsGCStep. The judgment says given an initial state satisfying the precondition $p_{\rm gc}$, if the environment's behaviors satisfy $\mathcal{R}_{\rm gc}$, then each step of $C_{\rm gc}$ satisfies $\mathcal{G}_{\rm gc}$, and the postcondition $q_{\rm gc}$ holds at the end if $C_{\rm gc}$ terminates. In general, the collector never terminates, thus

we can let $q_{\rm gc}$ be **false**. $\mathcal{G}_{\rm gc}$ and $p_{\rm gc}$ should be provided by the verifier, where $p_{\rm gc}$ needs to be general enough so that it can be satisfied by any possible low-level initial state. $\mathcal{R}_{\rm gc}$ encodes the possible behaviors of mutators, which can be derived, as we will show shortly.

Verifying mutators. For the mutator thread, since **T** is syntax-directed on \mathbb{C} , we can reduce the refinement problem for arbitrary mutators to the refinement on each primitive instruction only, following the compositionality of RGSim. The proof needs proper rely/guarantee conditions. Let \mathbb{G}_{c}^{t} and $\mathcal{G}_{T(c)}^{t}$ denote the guarantees of the source instruction c and the target code **T**(c) for the mutator thread t respectively. Then we can define the general guarantees for the thread.

$$\mathcal{G}(\mathbf{t}) \triangleq \bigcup_{\mathbf{c}} \mathcal{G}_{\mathbf{T}(\mathbf{c})}^{\mathbf{t}}; \qquad \mathbb{G}(\mathbf{t}) \triangleq \bigcup_{\mathbf{c}} \mathbb{G}_{\mathbf{c}}^{\mathbf{t}}.$$
 (7.1)

Its rely conditions should include all the possible guarantees made by other threads, and the GC's abstract and concrete behaviors respectively.

$$\mathcal{R}(\mathsf{t}) \triangleq \mathcal{G}_{gc} \cup \bigcup_{\mathsf{t}' \neq \mathsf{t}} \mathcal{G}(\mathsf{t}'); \qquad \mathbb{R}(\mathsf{t}) \triangleq \mathsf{AbsGCStep} \cup \bigcup_{\mathsf{t}' \neq \mathsf{t}} \mathbb{G}(\mathsf{t}'). \tag{7.2}$$

The \mathcal{R}_{gc} used to verify the GC code can now be defined.

$$\mathcal{R}_{gc} \triangleq \bigcup_{t} \mathcal{G}(t) \tag{7.3}$$

The refinement proof also needs definitions of binary relations α , ζ and γ . The invariant α relates the low-level and the high-level states and needs to be preserved by each low-level step. In general, a high-level state Σ can be mapped to a low-level state σ by giving a concrete local store for the GC thread, adding additional structures in the heap (to record information for collection), renaming heap cells (for copying GCs), etc. The relations ζ and γ are parametrized over the thread id t. For each mutator thread t, $\zeta(t)$ and $\gamma(t)$ need to hold at the beginning and the end of each basic transformation unit (every high-level primitive instruction in this case) respectively. We let $\gamma(t)$ be the same as $\zeta(t)$ to support sequential compositions. We require lnitRel_T($\zeta(t)$) (see Figure 6), that is, $\zeta(t)$ holds over the initial states. In addition, the target and the source boolean expressions should be evaluated to the same value under ζ -related states, as required in the IF and WHILE rules in Figure 7.

$$\mathsf{Good}_{\mathbf{T}}(\zeta(\mathfrak{t})) \triangleq \mathsf{Init}\mathsf{Rel}_{\mathbf{T}}(\zeta(\mathfrak{t})) \land \forall \mathbb{B}. \ \zeta(\mathfrak{t}) \subseteq (\mathbf{T}(\mathbb{B}) \Leftrightarrow \mathbb{B})$$
(7.4)

THEOREM 7.1 (VERIFYING CONCURRENT GARBAGE COLLECTING SYSTEMS). If there exist \mathbb{G}_{c}^{t} , $\mathcal{G}_{\mathbf{T}(c)}^{t}$, $\zeta(t)$, α , \mathcal{G}_{gc} and p_{gc} (for any c and t) such that the following hold (where $\mathcal{G}(t)$, $\mathbb{G}(t)$, $\mathcal{R}(t)$, $\mathbb{R}(t)$ and \mathcal{R}_{gc} are defined in (7.1), (7.2) and (7.3), and $\mathsf{Good}_{\mathbf{T}}(\zeta(t))$ defined in (7.4) holds):

- (1) (Correctness of **T** on mutator instructions) $\forall t, c. (\mathbf{T}(c), \mathcal{R}(t), \mathcal{G}(t)) \preceq^{t}_{\alpha; \zeta(t) - \zeta(t)} (c, \mathbb{R}(t), \mathbb{G}(t));$
- (2) (Verification of the GC code) $\mathcal{R}_{gc}; \mathcal{G}_{gc} \vdash \{p_{gc}\}C_{gc}\{\mathbf{false}\};$
- (3) (Side conditions) $\mathcal{G}_{gc} \circ \alpha^{-1} \subseteq \alpha^{-1} \circ (AbsGCStep)^*; and \forall \sigma, \Sigma, \sigma = \mathbf{T}(\Sigma) \implies p_{gc} \sigma;$

then Correct(T).

That is, to verify a concurrent garbage collecting system, we need to do the following.

— Define the α and $\zeta(t)$ relations, and prove the correctness of **T** on high-level primitive instructions. Since **T** preserves the syntax on most instructions, it's often immediate to prove the target instructions are simulated by their sources. But for

ACM Transactions on Programming Languages and Systems, Vol. 36, No. 1, Article 3, Publication date: March 2014.

instructions that are transformed to barriers, we need to verify that the barriers implement both the source instructions (by RGSim) and the interaction mechanism (shown in their guarantees).

— Find some proper \mathcal{G}_{gc} and p_{gc} , and verify the GC code by R-G reasoning. We require the GC's guarantee \mathcal{G}_{gc} should not contain more behaviors than AbsGCStep (the first side condition), and C_{gc} can start its execution from any state σ transformed from a high-level one (the second side condition).

To prove Theorem 7.1, we first prove the following from (2) and (3).

$$(C_{\text{gc}}, \mathcal{R}_{\text{gc}}, \mathcal{G}_{\text{gc}}) \preceq_{\alpha; \zeta_{\sigma c}} (AbsGC, \text{True}, \text{AbsGCStep})$$

Here $\zeta_{gc} \triangleq \{(\sigma, \Sigma) \mid \sigma = \mathbf{T}(\Sigma)\}$. The proof directly follows the RGSim definition. Then with (1) and the compositionality of RGSim, we can get the following by induction over the program structure.

$$\begin{aligned} \forall \mathbb{C}_1, \dots, \mathbb{C}_n. \ (\mathsf{t}_{\mathrm{gc}}.C_{\mathrm{gc}} \, \| \, \mathsf{t}_1.\mathbf{T}(\mathbb{C}_1) \, \| \dots \| \, \mathsf{t}_n.\mathbf{T}(\mathbb{C}_n), \mathsf{Id}, \mathsf{True}) \\ \leq_{\alpha; \zeta} \ \zeta \ (\mathsf{t}_{\mathrm{gc}}.AbsGC \| \| \mathsf{t}_1.\mathbb{C}_1 \| \| \dots \| \| \mathsf{t}_n.\mathbb{C}_n, \mathsf{Id}, \mathsf{True}) \end{aligned}$$

Here $\zeta \triangleq \zeta_{gc} \cap \bigcap_t \zeta(t)$. Finally, from the soundness of RGSim (Corollary 4.4), we can conclude Correct(**T**).

7.3. Application: Boehm et al. Concurrent GC Algorithm

We illustrate the applications of the framework (Theorem 7.1) by proving the correctness of a mostly concurrent mark-sweep garbage collector proposed by Boehm et al. [1991]. Variants of the algorithm have been used in practice (e.g., by IBM [Barabash et al. 2005]).

7.3.1. Overview of the GC Algorithm. The GC runs both the mark and sweep phases concurrently with the mutators. In the mark phase, it does a depth-first tracing and marks the objects which are reachable from the *roots* (i.e., the mutators' local pointer variables that may contain references to the heap objects). Later in the sweep phase, it scans the heap and reclaims unmarked objects. During the tracing, the connectivity between objects might be changed by the mutators, thus a write barrier is required to notify the collector of those modified objects. Boehm et al.'s algorithm gives each object a dirty bit (called a *card*) and its write barrier dirties the card of the object being updated. Then, between the mark and sweep phases, the GC runs a short stop-the-world phase, where it suspends all the mutators and retraces from the dirty objects which have been marked (called *card-cleaning*). Thus all reachable objects have been marked before the sweep phase, ensuring the correctness of the GC.

We show the code of the GC thread in Figure 15. We assume each object contains m pointer fields pt_1, \ldots, pt_m , a data field, and two auxiliary color and dirty fields. The color field has three possible values and is used for two purposes: for marking, we use BLACK for a marked object and WHITE for an unmarked one; and for allocation, we use BLUE for an unallocated object which will neither be traced nor be reclaimed, but can be allocated later. New objects are created BLACK, and when reclaiming an object, we just set its color to BLUE. The dirty field is the card bit whose value can be 0 (not dirty) or 1 (dirty). We also assume the total number of threads is N and the heap domain is [1..M].

To make the GC code more readable, we divide it into several methods in Figure 15, which should be viewed as macros. The GC thread executes Collection() and repeats the collection cycle (the loop body in the method) forever. In each collection cycle, it first clears the dirty cards and resets the colors of all the objects (the method

ACM Transactions on Programming Languages and Systems, Vol. 36, No. 1, Article 3, Publication date: March 2014.

```
1 constant int WHITE, BLACK, BLUE; // colors
                                               50 MarkAndPush(i) {
2 constant int N; // total number of threads
                                                51 local c;
3 constant int M; // size of heap
                                               52 if (i != 0) {
 4
                                               53
                                                      c := i.color;
5 Collection() {
                                               54
                                                      if (c = WHITE) {
6
    local mstk;
                                               55
                                                        i.color := BLACK;
7
    while (true) {
                                               56
                                                         push(i, mstk);
8
       Initialize();
                                                57
                                                       }
9
       Trace();
                                                58
                                                     }
                                                59 }
10
       CleanCard();
       atomic{ ScanRoot(); CleanCard(); }
                                                60
11
                                                61 CleanCard() {
12
       Sweep();
13
     }
                                                62
                                                   local i, c, d;
14 }
                                                63
                                                   i := 1;
                                                     while (i <= M) {
15
                                                64
16 Initialize() {
                                                65
                                                       c := i.color;
17
     local i, c;
                                                66
                                                       d := i.dirty;
                                                      if (d = 1) {
18
   i := 1;
                                                67
                                                         i.dirty := 0;
   while (i <= M) {
                                                68
19
                                                         if (c = BLACK) {
20
      i.dirty := 0;
                                                69
                                                70
21
      c := i.color;
                                                           push(i, mstk);
22
      if (c = BLACK) { i.color := WHITE; }
                                                71
                                                         7
                                                72
23
      i := i + 1;
                                                       }
24
     }
                                               73
                                                       i := i + 1;
25 }
                                                74
                                                     }
26
                                                75
                                                     TraceStack();
                                                76 }
27 Trace() {
    local t, rt, i;
                                                77
28
29
     t := 1;
                                                78 ScanRoot() {
30
   while (t <= N) {
                                                79
                                                   local t, rt, i;
31
     rt := get_root(t);
                                                80
                                                     t := 1;
32
      foreach i in rt do {
                                               81
                                                     while (t <= N) {
33
        MarkAndPush(i);
                                                82
                                                       rt := get_root(t);
34
       }
                                                83
                                                       foreach i in rt do {
35
       t := t + 1;
                                                84
                                                         MarkAndPush(i);
36
       TraceStack();
                                               85
                                                      }
37
     }
                                                86
                                                       t := t + 1;
38 }
                                                87
                                                     }
39
                                                88 }
40 TraceStack() {
                                               89
                                               90 Sweep() {
41
     local i, j;
42
     while (!is_empty(mstk)) {
                                               91
                                                   local i, c;
43
      i := pop(mstk);
                                               92
                                                   i := 1;
44
       j := i.pt<sub>1</sub>; MarkAndPush(j);
                                               93
                                                     while (i <= M) {
45
                                               94
                                                       c := i.color;
46
       j := i.pt<sub>m</sub>; MarkAndPush(j);
                                                95
                                                       if (c = WHITE) { free(i); }
47
     }
                                                96
                                                      i := i + 1;
48 }
                                                97
                                                     }
                                                98 }
49
```



```
update(x, fd, E) { // fd ∈ {pt1, ..., ptm}
atomic{ x.fd := E; aux := x; }
atomic{ x.dirty := 1; aux := 0; }
}
```

Fig. 16. The write barrier for Boehm et al. GC.

call of Initialize()). After the initialization, the GC enters the mark phase by calling Trace(). The command $rt := get_root(t)$ (line 31) allows the GC to read the values of all the pointer variables in the thread t's store at once to a set rt, and foreach *i* in rt do C allows to execute C for every value *i* in rt. Our atomic get_root tries to reflect the real-world GC implementation [Barabash et al. 2005], where the GC stops a mutator thread to scan its roots. A mark stack mstk is used to do the depth-first tracing in the method TraceStack(). For simplicity, we assume there are primitive commands push(x, mstk) and x := pop(mstk) to manipulate mstk. The stop-the-world phase (line 11) is implemented by $atomic{C}$. Here the roots are re-scanned in ScanRoot(), because the write barrier is not applied to the roots and we should assume conservatively that they have been modified. In the sweep phase (the call of Sweep() at line 12), the GC can use free(x) to reclaim the object x. Usually in practice, there is also a concurrent card-cleaning phase (the call of CleanCard() at line 10) before the stop-the-world card-cleaning (at line 11) to reduce the pause time of the latter.

The write barrier is shown in Figure 16, where the dirty field is set after modifying the object's pointer field. Here we use a write-only auxiliary variable aux for each mutator thread to record the current object that the mutator is updating. We add aux for the purpose of verification only, which can be safely deleted after the proof is completed. We use aux to help specify some fine-grained and temporal property of the write barrier in the guarantees. For instance, a mutator should ensure that after it sets a pointer field of an object x to another object y, it must first set x's dirty field before updating other pointers (in particular, those pointing to y). Otherwise, the GC may not know that y is newly reachable from x and may finally reclaim y. In Figure 16 we set aux to the object x when its pointer field is updated, and specify in the mutator's guarantee ($\mathcal{G}_{set.dirty}^t$ in Figure 25(b)) that when aux = x, it must set x's dirty field. The GC does not use read barriers nor allocation barriers. Allocation can be implemented using a standard concurrent list algorithm. To be more focused on verifying the GC algorithm itself, we model allocation as an abstract instruction $x := \mathbf{new}()$ which can magically find an unallocated (BLUE) object in the heap.

7.3.2. The Transformation. We first present the detailed high-level and low-level languages and state models in Figures 17 and 18 respectively, which are instantiations of the generic languages in Figure 2.

- An object has *m* pointer fields and a data field from the high-level view, whereas a concrete object also has two auxiliary fields color and dirty for the collection.
- The behaviors of the high-level abstract GC thread are defined in Figure 19(a), saying that the mutator stores and the reachable objects in the heap remain unmodified. Here Reachable(l)(Π , H) means the object at the location l is reachable in H from the roots in Π .
- The low-level concrete GC thread could use privileged commands, such as $x := get_root(y)$ and free(x), to control the mutator threads and manage the heap.
- High-level mutators can use x := y.fd to read a field of an object, x.fd $:= \mathbb{E}$ to write the value of \mathbb{E} to a field of an object and x := new() to allocate a new object. If

(a) the language

 $\begin{array}{rcl} (Loc) & l &\in \{L_1, \dots, L_M, \mathbf{nil}\}\\ (HVal) & \mathbf{V} &\in Int \cup Loc\\ (HStore) & \mathbf{S} &\in PVar \rightarrow HVal\\ (HObj) & \mathbf{O} &\in HField \rightarrow HVal\\ (HHeap) & \mathbf{H} &\in Loc \rightarrow HObj\\ (HThrds) & \Pi &\in MutID \rightarrow HStore\\ (HState) & \boldsymbol{\Sigma} &\in HThrds \times HHeap \end{array}$

(b) program states

Fig. 17. The high-level language and state model.

the instruction $x.fd := \mathbb{E}$ updates a pointer field (i.e., $fd \in \{pt_1, \ldots, pt_m\}$), then it will be transformed to the write barrier in Figure 16. Note here \mathbb{E} is restricted to be either **nil** (null pointers) or pointer variables.

- The high-level language is typed in the sense that heap locations and integers are regarded as distinct kinds (or types) of values. We present the high-level operational semantics in Figure 19(b). Here we use SameType(V, V') to mean that the two values V and V' are of the same type.
- On the low-level machine, we allow the GC to perform pointer arithmetic, so we do not distinguish locations and integers. A low-level value v can be an integer, a set, or a sequence of integers. We use $\mathcal{P}(_)$ for the power set and $Seq(_)$ for the set of sequences. Every low-level variable is given an extra bit to preserve its high-level type information (0 for nonpointers and 1 for pointers), so that the GC can easily get the roots. The low-level mutators are still prohibited from pointer arithmetic. An expression E is evaluated (shown in Figure 20) under the store with an extra tag *tag* to indicate whether it is used as an object location in the heap (tag = 1 if E is used as a heap location; and tag = 0 otherwise). When tag = 2, we do not care about the usage of the expression, and such an expression will be used in the GC code since the GC has the privilege to use an integer as an address and vice versa. We present part of the low-level operational semantics rules in Figure 21. To formulate the semantics of **foreach** x **in** y **do** C, we assume x and y are temporary variables and not updated by C. At the beginning of each iteration, we set x to an

 $(LExpr) E ::= x | n | E+E | E-E | \dots$ $(LBExp) B ::= true | false | E=E | !B | is_empty(x) | \dots$ (LInstr) c ::= print(E) | x := E | x := y.fd | x.fd := E | x := new() $| x := get_root(y) | free(x) | push(x,y) | x := pop(y)$ $(LStmts) C ::= skip | c | C_1; C_2 | if (B) C_1 else C_2 | while (B) C$ $| atomic{C} | foreach x in y do C$

 $(LProg) W ::= \mathsf{t}_{\mathsf{gc}}.C_{\mathsf{gc}} || \mathsf{t}_1.C_1 || \dots || \mathsf{t}_n.C_n$

 $(LField) fd \in \{pt_1, \dots, pt_m, data, color, dirty\}$

(a) the language

 $\begin{array}{rcl} (LVal) & v &\in Int \cup \mathcal{P}(Int) \cup Seq(Int) \\ (LStore) & s &\in PVar \rightarrow LVal \times \{0,1\} \\ (LObj) & o &\in LField \rightarrow LVal \\ (LHeap) & h &\in [1..M] \rightarrow LObj \\ (LThrds) & \pi &\in (MutID \cup \{t_{gc}\}) \rightarrow LStore \\ (LState) & \sigma &\in LThrds \times LHeap \end{array}$

(b) program states

Fig. 18. The low-level language and state model.

arbitrary item in the set y, and after executing C we remove that item from y. The **foreach** loop terminates when y becomes empty.

— We do not provide infinite heaps; instead there are only M valid high-level locations and the low-level heap domain is [1..M]. High-level mutators can use **nil** for null pointers and it will be translated to 0 on the low-level machine. We assume there is a bijective function from high-level locations to low-level integers

 $Loc2Int: Loc \leftrightarrow [0..M]$

which satisfies Loc2Int(nil) = 0.

The transformation **T** is defined as follows. For *code*, the high-level abstract GC thread is transformed to the GC thread shown in Figure 15. Each instruction $x.fd := \mathbb{E}$ in mutators is transformed to the write barrier update($x, fd, \mathbf{T}(\mathbb{E})$), where fd is a pointer field of x. **T** over expressions \mathbb{E} returns 0 if \mathbb{E} is **nil**, and keeps the syntax otherwise. Other instructions and the program structures of mutators are unchanged.

We also need to transform the initial high-level state to the low level. The transformation $\mathbf{T}(\Sigma)$ is defined in Figure 22.

- First we require the high-level initial state to be *well-formed* (WfState(Σ)), that is, reachable locations cannot be dangling pointers.
- -High-level locations are transformed to integers by the bijective function Loc2Int.

$$\begin{array}{lll} \operatorname{Root}(t,S) & \triangleq \lambda \Sigma. \ \Sigma = (\Pi \uplus \{t & \operatorname{S}_t\}, \operatorname{H}) \land S = \{l \mid \exists x. \operatorname{S}_t(x) = l\} \\ \operatorname{Edge}(l_1,l_2) & \triangleq \lambda \Sigma. \ \Sigma = (\Pi,\operatorname{H}) \land \exists fd \in \{\operatorname{pt}_1,\ldots,\operatorname{pt}_m\}. \ \operatorname{H}(l_1)(fd) = l_2 \\ \operatorname{Path}_k(l_1,l_2) & \triangleq \begin{cases} l_1 = l_2 & \text{if } k = 0 \\ \exists l_3. \ \operatorname{Edge}(l_1,l_3) \land \operatorname{Path}_{k-1}(l_3,l_2) & \text{if } k > 0 \end{cases} \\ \operatorname{Path}(l_1,l_2) & \triangleq \exists k. \ \operatorname{Path}_k(l_1,l_2) \\ \operatorname{Reachable}(t,l) & \triangleq \exists S, l'. \ \operatorname{Root}(t,S) \land l' \in S \land \operatorname{Path}(l',l) \land l \neq \mathbf{nil} \\ \operatorname{Reachable}(l) & \triangleq \exists t \in [1..N]. \ \operatorname{Reachable}(t,l) \\ \operatorname{AbsGCStep} & \triangleq \{((\Pi,\operatorname{H}),(\Pi,\operatorname{H}')) \mid \forall l. \ \operatorname{Reachable}(l)(\Pi,\operatorname{H}) \Longrightarrow \operatorname{H}(l) = \operatorname{H}'(l)\} \end{array}$$

 $(a) \ definition \ of \ {\tt AbsGCStep}$

$$\begin{split} & \frac{\mathbf{S}(x) = l \quad \mathbf{H}(l) = \mathbf{O} \quad \llbracket \mathbb{E} \rrbracket_{\mathbf{S}} = \mathbf{V} \quad \mathbf{O}(\mathbf{fd}) = \mathbf{V}' \quad \mathbf{SameType}(\mathbf{V}, \mathbf{V}') \\ \hline & \overline{(x.\mathbf{fd} := \mathbb{E}, (\Pi \uplus \{\mathbf{t} \ \mathbf{S}\}, \mathbf{H})) \longrightarrow_{\mathbf{t}} (\mathbf{skip}, (\Pi \uplus \{\mathbf{t} \ \mathbf{S}\}, \mathbf{H}|l \quad \mathbf{O}(\mathbf{fd} \ \mathbf{V}\})))} \\ & \frac{x \notin dom(\mathbf{S}) \quad \text{or} \quad \mathbf{S}(x) \notin dom(\mathbf{H}) \quad \text{or} \quad \llbracket \mathbb{E} \rrbracket_{\mathbf{S}} = \bot \quad \text{or} \quad \neg \mathbf{SameType}(\mathbf{H}(\mathbf{S}(x))(\mathbf{fd}), \llbracket \mathbb{E} \rrbracket_{\mathbf{S}}) \\ \hline & \overline{(x.\mathbf{fd} := \mathbb{E}, (\Pi \uplus \{\mathbf{t} \ \mathbf{S}\}, \mathbf{H})) \longrightarrow_{\mathbf{t}} \mathbf{abort}} \\ \hline & \frac{l \notin dom(\mathbf{H}) \quad l \neq \mathbf{nil} \quad \mathbf{S}(x) = l' \quad \mathbf{S}' = \mathbf{S}\{x \quad l\} \quad \mathbf{H}' = \mathbf{H} \uplus \{l \quad \{\mathbf{pt}_1 \quad \mathbf{nil}, \dots, \mathbf{pt}_m \quad \mathbf{nil}, \mathbf{data} \quad 0\}\} \\ \hline & \overline{(x := \mathbf{new}(), (\Pi \uplus \{\mathbf{t} \ \mathbf{S}\}, \mathbf{H})) \longrightarrow_{\mathbf{t}} (\mathbf{skip}, (\Pi \uplus \{\mathbf{t} \ \mathbf{S}'\}, \mathbf{H}'))} \\ & \frac{\neg (\exists l.l \notin dom(\mathbf{H}) \land l \neq \mathbf{nil}) \quad \mathbf{S}(x) = l' \quad \mathbf{S}' = \mathbf{S}\{x \quad \mathbf{nil}\}}{(x := \mathbf{new}(), (\Pi \uplus \{\mathbf{t} \ \mathbf{S}\}, \mathbf{H})) \longrightarrow_{\mathbf{t}} (\mathbf{skip}, (\Pi \uplus \{\mathbf{t} \ \mathbf{S}'\}, \mathbf{H}))} \\ \hline & \frac{x \notin dom(\mathbf{S}) \quad \text{or} \quad \neg \exists l.\mathbf{S}(x) = l}{(x := \mathbf{new}(), (\Pi \uplus \{\mathbf{t} \ \mathbf{S}\}, \mathbf{H})) \longrightarrow_{\mathbf{t}} (\mathbf{skip}, (\Pi \uplus \{\mathbf{t} \ \mathbf{S}'\}, \mathbf{H}))} \\ \hline & \frac{x \notin dom(\mathbf{S}) \quad \text{or} \quad \neg \exists l.\mathbf{S}(x) = l}{(\mathbf{t} \ \mathbf{S}\}, \mathbf{H}) \longrightarrow_{\mathbf{t}} \mathbf{t} \mathbf{sbort}} \quad \underbrace{(\mathbb{C}_i, \Sigma) \longrightarrow_{\mathbf{t}_i} \mathbf{abort}}_{(\mathbf{t}_{\mathbf{t}} \ldots \mathbb{T}_i, \mathbb{C}_i \dots |||\mathbf{t}_n, \mathbb{C}_n, \Sigma) \longrightarrow \mathbf{abort}} \end{split}$$

$$\frac{(\mathbb{C}_{i},\Sigma)\longrightarrow_{\mathfrak{t}_{i}}(\mathbb{C}'_{i},\Sigma') \quad \text{or} \quad (\Sigma,\Sigma') \in \mathsf{AbsGCStep} \land \mathbb{C}'_{i} = \mathbb{C}_{i}}{(\mathfrak{t}_{gc}.AbsGC |||\mathfrak{t}_{1}.\mathbb{C}_{1}||| \dots \mathfrak{t}_{i}.\mathbb{C}_{i} \dots |||\mathfrak{t}_{n}.\mathbb{C}_{n},\Sigma) \longrightarrow (\mathfrak{t}_{gc}.AbsGC |||\mathfrak{t}_{1}.\mathbb{C}_{1}||| \dots \mathfrak{t}_{i}.\mathbb{C}'_{i} \dots |||\mathfrak{t}_{n}.\mathbb{C}_{n},\Sigma')}$$

(b) selected operational semantics rules

Fig. 19. A high-level garbage-collected machine.

$$\begin{split} \llbracket n \rrbracket_{(s,tag)} &= \begin{cases} n & \text{if } tag = 0 \text{ or } tag = 2 \\ 0 & \text{if } tag = 1 \text{ and } n = 0 \\ \bot & \text{otherwise} \end{cases} \\ \llbracket x \rrbracket_{(s,tag)} &= \begin{cases} n & \text{if } s(x) = (n,b) \text{ and } (tag = b \lor tag = 2) \\ \bot & \text{otherwise} \end{cases} \\ \llbracket E_1 + E_2 \rrbracket_{(s,tag)} &= \begin{cases} n_1 + n_2 & \text{if } \llbracket E_1 \rrbracket_{(s,tag)} = n_1 \text{ and } \llbracket E_2 \rrbracket_{(s,tag)} = n_2 \text{ and } (tag = 0 \lor tag = 2) \\ \bot & \text{otherwise} \end{cases} \\ \llbracket \text{is_empty}(x) \rrbracket_{(s,tag)} &= \begin{cases} \text{true} & \text{if } tag = 0 \text{ and } s(x) = (\epsilon, 0) \\ \texttt{false} & \text{if } tag = 0 \text{ and } s(x) = (n::A, 0) \\ \bot & \text{otherwise} \end{cases} \end{cases}$$

Fig. 20. Expression evaluation on the low-level machine.

— Variables are transformed to the low level using an extra bit to preserve the highlevel type information (0 for nonpointers and 1 for pointers). Usually we use v^{np} and v^{p} short for (v, 0) and (v, 1) respectively.

$$\begin{split} \frac{\mathfrak{t} \in [1.N] \quad s(x) = (.,b) \quad [\mathbb{E}]_{(s,b)} = n \quad s' = s(x \quad (n,b))}{(x := E, (\pi \uplus [\mathfrak{t} = s], h)) \longrightarrow_{\mathfrak{t}} (\mathfrak{skip}, (\pi \uplus [\mathfrak{t} = s'], h))} \\ \frac{s(x) = (.,b) \quad [\mathbb{E}]_{(s,2)} = n \quad s' = s(x \quad (n,b))}{(x := E, (\pi \uplus [\mathfrak{tg} = s], h)) \longrightarrow_{\mathfrak{tg}} (\mathfrak{skip}, (\pi \uplus [\mathfrak{tg} = s'], h))} \\ \frac{s(y) = (n_y, 1) \quad h(n_y)(fd) = n \quad s(x) = (.,b)}{fd \in [\mathfrak{pt}_1, \dots, \mathfrak{pt}_m] \Longrightarrow b = 1 \quad fd \in (\mathfrak{dat}) \Longrightarrow b = 0 \quad s' = s(x \quad (n,b))}{(x := y, fd, (\pi \uplus [\mathfrak{t} = s], h)) \longrightarrow_{\mathfrak{t}} (\mathfrak{skip}, (\pi \uplus [\mathfrak{tg} = s'], h))} \\ \frac{s(x) = (n, 1) \quad h(n) = o \quad fd \in (\mathfrak{pt}_1, \dots, \mathfrak{pt}_m] \Longrightarrow [\mathbb{E}]_{(s,2)} = n'}{(xfd := E, (\pi \uplus [\mathfrak{t} = s], h)) \longrightarrow_{\mathfrak{t}} (\mathfrak{skip}, (\pi \uplus [\mathfrak{t} = s], h(n \quad o(fd \quad n'])))} \\ \frac{s(y) = (t, 0) \quad s(x) = (., 0) \quad \pi(t) = \mathfrak{st} \quad S = (n \mid \exists x.\mathfrak{st}(x) = (n, 1) \quad s' = \mathfrak{st}(x \quad (S, 0))}{(\mathfrak{t} = \mathfrak{gt.troot}(y), (\pi \uplus [\mathfrak{tgg} = s], h)) \longrightarrow_{\mathfrak{tgg}} (\mathfrak{skip}, (\pi \uplus [\mathfrak{tgg} = s'], h))} \\ \frac{s(y) = (t, 0) \quad s(x) = (., 0) \quad \pi(t) = \mathfrak{st} \quad S = (n \mid \exists x.\mathfrak{st}(x) = (n, 1) \quad s' = \mathfrak{st}(x \quad (S, 0))}{(\mathfrak{foreach } x \text{ in } y \text{ do } C, (\pi \uplus [\mathfrak{tgg} = s], h)) \longrightarrow_{\mathfrak{tgg}} (\mathfrak{skip}, (\pi \uplus [\mathfrak{tgg} = s], h))} \\ \frac{s(x) = (., b) \quad s(y) = ((n_1, \dots, n_k), 0) \quad s' = \mathfrak{st}(x \quad (n_1, b))}{(\mathfrak{atomic}(C), (\pi \uplus [\mathfrak{t} = s], h)) \longrightarrow_{\mathfrak{tgg}} (\mathfrak{skip}, (\pi \uplus [\mathfrak{tgg} = s], h))} \\ \frac{\mathfrak{t} \in [1.N] \quad s(x) = (., 1) \quad h(n)(\operatorname{color}) = \operatorname{BLUE} \quad s' = \mathfrak{st}(x \quad (n, 1))}{(\mathfrak{atomic}(C), (\pi \uplus [\mathfrak{t} = s], h)) \longrightarrow_{\mathfrak{tgg}} (\mathfrak{skip}, (\pi \uplus [\mathfrak{t} = s], h)) \longrightarrow_{\mathfrak{t}} \mathfrak{abort}} \\ \frac{\mathfrak{t} \in [1.N] \quad \mathfrak{s}(x) = (., 1) \quad h(n)(\operatorname{color}) = \operatorname{BLUE} \quad s' = \mathfrak{st}(x \quad (n, 1))}{(x := \operatorname{new}(), (\pi \uplus [\mathfrak{t} = s], h)) \longrightarrow_{\mathfrak{t}} \mathfrak{skip}, (\pi \uplus [\mathfrak{t} = s'], h')} \\ \frac{\mathfrak{t} \in [1.N] \quad \mathfrak{s}(x) = (., 1) \quad h(n)(\operatorname{color}) = \operatorname{BLUE} \quad s' = \mathfrak{st}(x \quad (0, 1))}{(x := \operatorname{new}(), (\pi \uplus [\mathfrak{t} = s], h)) \longrightarrow_{\mathfrak{t}} \mathfrak{skip}, (\pi \uplus [\mathfrak{t} = s'], h')} \\ \frac{\mathfrak{t} \in [1.N] \quad \mathfrak{s}(x) = (., 1) \quad h(n)(\operatorname{color}) = \operatorname{BLUE} \quad s' = \mathfrak{st}(x \quad (0, 1))}{(x := \operatorname{new}(), (\pi \uplus [\mathfrak{t} = s], h)) \longrightarrow_{\mathfrak{t}} \mathfrak{skip}, (\pi \uplus [\mathfrak{t} = s'], h')} \\ \frac{\mathfrak{t} \in [1.N] \quad \mathfrak{s}(x) = (., 1) \quad h(n)(\operatorname{color}) = \operatorname{BLUE} \quad s' = \mathfrak{st}(x \quad (0, 1))}{(x := \operatorname{new}(), (\pi \uplus [\mathfrak{t} = s], h)) \longrightarrow_{\mathfrak{t$$

Fig. 21. Selected operational semantics rules on the low-level machine.

— High-level objects are transformed to the low level by adding the color and dirty fields with initial values WHITE and 0 respectively. Other addresses in the low-level heap domain [1..*M*] are filled out using unallocated objects whose colors are BLUE and all the other fields are initialized by 0.

$$\mathbf{T}(\Sigma) \triangleq \begin{cases} (\{\mathbf{t} \quad \mathbf{T}(S) \mid (\mathbf{t} \quad S) \in \Pi\} \uplus \{\mathsf{t}_{gc} \quad s_{gc_init}\}, \mathbf{T}(H)) & \text{if } \Sigma = (\Pi, H) \land \mathsf{WfState}(\Sigma) \\ \bot & \text{otherwise} \end{cases}$$

where

WfState(Π , H) $\triangleq \forall l$. Reachable(l)(Π , H) $\Longrightarrow l \in dom(H)$

$$s_{ ext{gc_init}} \triangleq \{ ext{mstk} \quad \epsilon^{ ext{np}}, ext{rt} \quad \emptyset^{ ext{np}}, ext{i} \quad 0^p, ext{j} \quad 0^p, ext{c} \quad 0^{ ext{np}}, ext{d} \quad 0^{ ext{np}} \}$$

$$\mathbf{T}(\mathbf{S})(x) \triangleq \begin{cases} n^{\mathrm{np}} & \text{if } \mathbf{S}(x) = n \\ n^{\mathrm{p}} & \text{if } \mathbf{S}(x) = l \wedge \mathsf{Loc2Int}(l) = n \\ 0^{\mathrm{p}} & \text{if } x = \mathsf{aux} \\ \bot & \text{if } x \notin dom(\mathbf{S}) \wedge x \neq \mathsf{aux} \end{cases}$$

$$\mathbf{T}(\mathbf{H})(i) \triangleq \begin{cases} \{ \mathsf{pt}_1 \quad n_1, \dots, \mathsf{pt}_m \quad n_m, \mathsf{data} \quad n, \mathsf{color} \quad \mathsf{WHITE}, \mathsf{dirty} \quad 0 \} \\ & \text{if } \exists l. \ l \in dom(\mathbf{H}) \wedge \mathsf{Loc2Int}(l) = i \wedge 1 \leq i \leq M \\ & \wedge \mathbf{H}(l) = \{ \mathsf{pt}_1 \quad l_1, \dots, \mathsf{pt}_m \quad l_m, \mathsf{data} \quad n \} \\ & \wedge \mathsf{Loc2Int}(l_1) = n_1 \wedge \dots \wedge \mathsf{Loc2Int}(l_m) = n_m \end{cases} \\ \{ \mathsf{pt}_1 \quad 0, \dots, \mathsf{pt}_m \quad 0, \mathsf{data} \quad 0, \mathsf{color} \quad \mathsf{BLUE}, \mathsf{dirty} \quad 0 \} \\ & \text{if } \exists l. \ l \notin dom(\mathbf{H}) \wedge \mathsf{Loc2Int}(l) = i \wedge 1 \leq i \leq M \end{cases}$$

Fig. 22. The transformation **T** on initial states for Boehm et al. GC.

$$\begin{split} \mathsf{store_map}(s,\mathsf{S}) &\triangleq \forall x \neq \mathsf{aux.} (\forall n. s(x) = n^{\mathrm{np}} \Longleftrightarrow \mathsf{S}(x) = n) \\ & \land (\forall n. s(x) = n^{\mathrm{p}} \Longleftrightarrow \exists l. \ \mathsf{Loc2Int}(l) = n \land \mathsf{S}(x) = l) \\ \mathsf{obj_map}(o,\mathsf{O}) &\triangleq \exists n_1, \dots, n_m, n, c, l_1, \dots, l_m. \ \mathsf{Loc2Int}(l_1) = n_1 \land \dots \land \mathsf{Loc2Int}(l_m) = n_m \\ & \land o = \{\mathsf{pt}_1 \quad n_1, \dots, \mathsf{pt}_m \quad n_m, \mathsf{data} \quad n, \mathsf{color} \quad c, \mathsf{dirty} \quad _\} \land c \neq \mathsf{BLUE} \\ & \land \mathsf{O} = \{\mathsf{pt}_1 \quad l_1, \dots, \mathsf{pt}_m \quad l_m, \mathsf{data} \quad n\}) \\ \mathsf{unalloc}(o,\mathsf{H},l) &\triangleq (o = \{\mathsf{pt}_1 \quad _, \dots, \mathsf{pt}_m \quad _, \mathsf{data} \quad _, \mathsf{color} \quad \mathsf{BLUE}, \mathsf{dirty} \quad _\}) \land l \notin dom(\mathsf{H}) \\ \mathsf{heap_map}(h,\mathsf{H}) &\triangleq \forall i, l. \ 1 \leq i \leq M \land \mathsf{Loc2Int}(l) = i \implies \mathsf{obj_map}(h(i), \mathsf{H}(l)) \lor \mathsf{unalloc}(h(i), \mathsf{H}, l) \\ \alpha &\triangleq \{((\pi \uplus \{\mathsf{tgc} \quad _\}, h), (\Pi, \mathsf{H})) \mid \forall \mathsf{t.} \ \mathsf{store_map}(\pi(\mathsf{t}), \Pi(\mathsf{t})) \land \mathsf{heap_map}(h, \mathsf{H}) \land \mathsf{WfState}(\Pi, \mathsf{H})\} \end{split}$$

Fig. 23. The α relation for Boehm et al. GC.

— The concrete GC thread is given an initial store s_{gc_init} where its local variables are initialized by 0 (for integer and pointer variables), ϵ (for the mark stack mstk) or Ø (for the root set rt).

To prove Correct(T) in our framework, we apply Theorem 7.1, prove the refinement between low-level and high-level mutators, and verify the GC code using a unary Rely-Guarantee-based logic.

7.3.3. Refinement Proofs for Mutator Instructions. We first define the α and ζ (t) relations. In α (see Figure 23), the relations between low-level and high-level stores and heaps are enforced by store_map and heap_map respectively. Their definitions reflect the state transformations we describe before, ignoring the values of those high-level-invisible structures (e.g., the GC's local variables, the color and dirty fields for nonblue objects and all the fields of blue objects). α also requires the well-formedness of high-level states. Here we still use Loc2Int to relate integers and locations.

n

ACM Transactions on Programming Languages and Systems, Vol. 36, No. 1, Article 3, Publication date: March 2014.

For each mutator thread t, the $\zeta(t)$ relation enforced at the beginning and the end of each transformation unit (each high-level instruction) is stronger than α . It requires that the value of the auxiliary variable aux (see Figure 16) be a null pointer (0^p).

$$\zeta(\mathfrak{t}) \triangleq \alpha \cap \{((\pi, h), (\Pi, \mathbf{H})) \mid \pi(\mathfrak{t})(\mathtt{aux}) = 0^p\}.$$

To define the guarantees of the mutator instructions, we first introduce some separation logic assertions in Figure 24 to describe states. Following Parkinson et al. [2006], we treat program variables as resource and use $\mathsf{own}_p(x)$ and $\mathsf{own}_{np}(x)$ for the current thread's ownerships of pointers and nonpointers respectively. They are interpreted under (π, s, h) , where s is the store of the current thread, π consists of the stores of all the other threads and h is the shared heap. We use $E_1.fd \mapsto E_2$ to specify a singleobject single-field heap with E_2 stored in the field fd of the object E_1 . The separating conjunction p * q means p and q hold on disjoint states. We define the disjoint union of states in Figure 24(c). We use $f_1 \uplus f_2$ as usual to denote the union of two partial functions when their domains are disjoint. Since heaps are curried functions that first map locations to objects, which then map field names to values, they can be transformed to an uncurried form by the uncurry operator. We then use $h_1 \oplus h_2$ to denote the union when their domains of $uncurry(h_1)$ and $uncurry(h_2)$ are disjoint. The disjoint union of states is defined based on the disjoint unions of the shared heaps and the stores for each thread. We use $E_1.fd \hookrightarrow E_2$ for $E_1.fd \mapsto E_2 * \mathbf{true}$ and $\circledast_{x \in S}.p(x)$ for iterated separating conjunction over the set S. We overload the notations to the high-level machine and use \mathbb{E}_1 .fd $\Rightarrow \mathbb{E}_2$ for a single-object single-field heap at the high level.

In Figure 24(d), we define two forms of actions. $p \ltimes_t q$ represents the update over the current thread t's store and the shared heap, which is defined similarly as in Figure 11(b). $p \ltimes_t q$ provided p' ensures that the context p' is not changed by the action.

In Figure 25, we give the guarantees of the high-level mutator instructions and the transformed code, which are defined following their operational semantics. We use $(x^{p} = n)$ short for $(x = n) \land \text{own}_{p}(x)$ and $(x^{np} = n)$ for $(x = n) \land \text{own}_{np}(x)$. When the context is clear, we omit the superscript. The predicates blueobj and newobj denote a blue object and a newly allocated object, which are defined in Figure 27. Each action just accesses the local store of the mutator and will not touch the GC store.

The refinement between the write barrier at the low level and the pointer update instruction at the high level is formulated as

$$(\texttt{update}(x, \texttt{fd}, E), \mathcal{R}(\texttt{t}), \mathcal{G}_{\texttt{write_barrier}}^{\texttt{t}}) \preceq^{\texttt{t}}_{\alpha; \zeta(\texttt{t}) \quad \zeta(\texttt{t})} (x.\texttt{fd} := \mathbb{E}, \mathbb{R}(\texttt{t}), \mathbb{G}_{\texttt{write_pt}}^{\texttt{t}}),$$

where $\mathcal{G}_{write_barrier}^{t} \triangleq \mathcal{G}_{write_pt}^{t} \cup \mathcal{G}_{set_dirty}^{t}$, that is, the guarantee of the low-level two-step write barrier. $\mathbb{G}_{write_pt}^{t}$ is the guarantee of the high-level atomic write operation. Recall $\mathcal{R}(t)$ and $\mathbb{R}(t)$ are defined in Eq. (7.2) in Section 7.2. Since the transformation of other high-level instructions is identity, the corresponding refinement proofs are simple. For example, we can prove

$$(x := \mathbf{new}(), \mathcal{R}(\mathsf{t}), \mathcal{G}_{\mathrm{new}}^{\mathsf{t}} \cup \mathcal{G}_{\mathrm{assgn-pt}}^{\mathsf{t}}) \preceq_{\alpha;\zeta(\mathsf{t})}^{\mathsf{t}} \zeta(\mathsf{t}) (x := \mathbf{new}(), \mathbb{R}(\mathsf{t}), \mathbb{G}_{\mathrm{new}}^{\mathsf{t}} \cup \mathbb{G}_{\mathrm{assgn-pt}}^{\mathsf{t}})$$

7.3.4. Rely-Guarantee Reasoning about the GC Code. We use a unary logic to verify the GC thread. The proof details here are orthogonal to our simulation-based proof (but it is RGSim that allows us to derive Theorem 7.1, which then links proofs in the unary logic with relational proofs). Thus shortly we only give a sketch of the assertion language, the unary logic, the precondition and the guarantee of the GC thread, the key invariants, and the proof structure.

(a) state assertions

(b) shorthand notations for some state assertions (\uplus and \oplus defined below)

$$\begin{array}{ll} f_{1} \bot f_{2} & \triangleq \ dom(f_{1}) \cap dom(f_{2}) = \emptyset \\ f_{1} \uplus f_{2} & \triangleq \ \begin{cases} f_{1} \cup f_{2} \quad \text{if} \ f_{1} \bot f_{2} \\ \bot \quad \text{otherwise} \\ \\ h_{1} \oplus h_{2} & \triangleq \ \begin{cases} curry(uncurry(h_{1}) \cup uncurry(h_{2})) \quad \text{if} \ uncurry(h_{1}) \bot uncurry(h_{2}) \\ \bot \quad & \text{otherwise} \\ \end{cases} \\ \pi_{1} \oplus \pi_{2} & \triangleq \ \begin{cases} t \quad (\pi_{1}(t) \uplus \pi_{2}(t)) \mid t \in dom(\pi_{1}) \\ \downarrow \quad & \text{otherwise} \\ \\ f \quad (\pi_{1}(t) \boxminus \pi_{2}(t)) \mid t \in dom(\pi_{2}) \land \forall t \in dom(\pi_{1}). \ \pi_{1}(t) \bot \pi_{2}(t) \\ \bot \quad & \text{otherwise} \\ \end{cases} \\ \sigma_{1} \oplus \sigma_{2} & \triangleq \ \begin{cases} (\pi, h) \quad \text{if} \ \sigma_{1} = (\pi_{1}, h_{1}) \land \sigma_{2} = (\pi_{2}, h_{2}) \land \pi_{1} \oplus \pi_{2} = \pi \land h_{1} \oplus h_{2} = h \\ \bot \quad & \text{otherwise} \end{cases}$$

(c) disjoint unions

$$\begin{array}{l} p \ltimes_{\mathsf{t}} q \ \triangleq \ \{((\pi \uplus \{\mathsf{t} \quad s\}, h), (\pi \uplus \{\mathsf{t} \quad s'\}, h')) \ \mid \ \exists s_1, h_1, s_2, h_2, s'_1, h'_1. \ p(\pi, s_1, h_1) \land q(\pi, s'_1, h'_1) \land (s = s_1 \uplus s_2) \land (h = h_1 \uplus h_2) \land (s' = s'_1 \uplus s_2) \land (h' = h'_1 \uplus h_2)\} \\ p \ltimes_{\mathsf{t}} q \text{ provided } p' \ \triangleq \ (p \ltimes_{\mathsf{t}} q) \cap ((p \ast p') \ltimes_{\mathsf{t}} (q \ast p')) \end{array}$$

(d) actions

Fig. 24. Semantics of basic assertions.

The unary program logic we use to verify the GC thread is a standard rely-guarantee logic adapted to the target language. The assertions are defined in Figure 24 and discussed before. We show the inference rules in Figure 26. Rules on the top half are for

$$\begin{split} & \mathbb{G}_{assgn.int}^{t} & \triangleq \exists x, n, n'. (x = n \land emp_{h}) \ltimes_{t} (x = n' \land emp_{h}) \\ & \mathbb{G}_{assgn.pt}^{t} & \triangleq \exists x, l, l'. (x = l \land emp_{h}) \ltimes_{t} (x = l' \land emp_{h}) \\ & \text{provided} (l' = \mathbf{nil} \lor \exists y, y = l' \lor \exists y, \text{fd} \Rightarrow l') \\ & \mathbb{G}_{write.data}^{t} & \triangleq \exists x, n, n'. (x. data \Rightarrow n) \ltimes_{t} (x. data \Rightarrow n') \\ & \mathbb{G}_{write.pt}^{t} & \triangleq \exists x, fd, l, l'. (x. fd \Rightarrow l) \ltimes_{t} (x. fd \Rightarrow l') \text{ provided} (l' = \mathbf{nil} \lor \exists y, y = l') \\ & \mathbb{G}_{new}^{t} & \triangleq \exists x. (x = _ \land emp_{h}) \ltimes_{t} (x = l \land l. pt_{1} \Rightarrow \mathbf{nil} * \dots * l. pt_{m} \Rightarrow \mathbf{nil} * l. data \Rightarrow 0) \\ & \mathbb{G}(t) & \triangleq \mathbb{G}_{assgn.int}^{t} \cup \mathbb{G}_{assgn.pt}^{t} \cup \mathbb{G}_{write.data}^{t} \cup \mathbb{G}_{write.pt}^{t} \cup \mathbb{G}_{new}^{t} \end{aligned}$$

(a) high-level guarantees

$$\begin{array}{lll} \mathcal{G}^{t}_{assgn_{int}} & \triangleq \exists x, n, n'. \ (x^{np} = n \land emp_{h}) \ltimes_{t} (x^{np} = n' \land emp_{h}) \ \text{provided} \ (aux^{p} = 0) \\ \mathcal{G}^{t}_{assgn_{ipt}} & \triangleq \exists x, n, n'. \ (x^{p} = n \land emp_{h}) \ltimes_{t} (x^{p} = n' \land emp_{h}) \ \text{provided} \\ & (aux^{p} = 0 \ast (n' = 0 \lor \exists y, y^{p} = n' \lor \exists y, fd. \ fd \in \{pt_{1}, \ldots, pt_{m}\} \land y.fd \mapsto n' \lor n = n')) \\ \mathcal{G}^{t}_{write_data} & \triangleq \exists x, n, n'. \ (x.data \mapsto n) \ltimes_{t} (x.data \mapsto n') \ \text{provided} \ (aux^{p} = 0) \\ \mathcal{G}^{t}_{write_pt} & \triangleq \exists x, fd, n, n'. \ (aux^{p} = 0 \ast x.fd \mapsto n) \ltimes_{t} (aux^{p} = x \ast x.fd \mapsto n') \\ & \quad \text{provided} \ ((n' = 0 \lor \exists y, y^{p} = n') \land fd \in \{pt_{1}, \ldots, pt_{m}\}) \\ \mathcal{G}^{t}_{set_dirty} & \triangleq \exists n. \ (aux^{p} = n \ast n.dirty \mapsto _) \ltimes_{t} (aux^{p} = 0 \ast n.dirty \mapsto 1) \\ \mathcal{G}^{t}_{new} & \triangleq \exists x, n, n'. \ (x^{p} = n \ast blueobj(n')) \ltimes_{t} (x^{p} = n' \ast newobj(n')) \ \text{provided} \ (aux^{p} = 0) \\ \mathcal{G}(t) & \triangleq \ \mathcal{G}^{t}_{assgn_int} \cup \ \mathcal{G}^{t}_{assgn_pt} \cup \ \mathcal{G}^{t}_{write_data} \cup \ \mathcal{G}^{t}_{write_pt} \cup \ \mathcal{G}^{t}_{new} \end{array}$$

(b) low-level guarantees



sequential reasoning. Most are exactly the same as separation logic [Reynolds 2002] and omitted here. The figure only shows some rules we added for the GC-specific commands (e.g., $x := get_root(y)$) and some particular heap manipulation rules adapted to our low-level machine model (e.g., **free**(x) just sets the object's color to BLUE). The concurrency rules in the bottom half follow standard rely-guarantee reasoning. The soundness of the logic with respect to the operational semantics is straightforward and we omit the proofs here.

To verify the GC code, we first give the precondition and the guarantee of the GC. The GC starts its executions from a low-level *well-formed* state, that is, $p_{gc} \triangleq$ wfstate. Just corresponding to the high-level WfState definition (see Figure 22), the low-level wfstate predicate says that none of the reachable objects is BLUE, as

 $\mathsf{wfstate} \ \triangleq \ \circledast_{x \in [1..M]}.\mathsf{obj}(x) \land (\forall x. \ \mathsf{reachable}(x) \Rightarrow \mathsf{not_blue}(x)) \,,$

where obj(x) means x is a low-level heap location with the pt_1, \ldots, pt_m , data, color and dirty fields, reachable(x) is defined similarly to the high-level definition in Figure 19, and not_blue(x) is defined in Figure 27. We define \mathcal{G}_{gc} as follows.

$$\begin{array}{l} \mathcal{G}_{\mathrm{gc}} \triangleq \{ ((\pi \uplus \{ \mathrm{tgc} \quad s\}, h), (\pi \uplus \{ \mathrm{tgc} \quad s'\}, h')) \\ & \mid \forall n. \ \mathrm{reachable}(n)(\pi, h) \\ \implies \lfloor h(n) \rfloor = \lfloor h'(n) \rfloor \land h(n). \mathrm{color} \neq \mathrm{BLUE} \land h'(n). \mathrm{color} \neq \mathrm{BLUE} \} \end{array}$$

$$\frac{\{(x^{np} = X') * (1 \le y^{np} \le N)\}x := \mathbf{get}_{\mathbf{root}}(y)\{(x^{np} = X) * (1 \le y^{np} \le N \land \operatorname{root}(y, X))\}}{\{x.\operatorname{color} \mapsto _\}\mathbf{free}(x)\{x.\operatorname{color} \mapsto \mathsf{BLUE}\}} (FREE)$$

$$\frac{\{y, O_{np}; O_p \Vdash x = X \land y = Y\}\mathbf{push}(x, y)\{y, O_{np}; O_p \Vdash x = X \land y = X :: Y\}}}{\{y, O_{np}; O_p \Vdash x = X \land y = X' :: Y\}x := \mathbf{pop}(y)\{y, O_{np}; O_p \Vdash x = X' \land y = Y\}} (POP)$$

$$\frac{\{p\}C\{q\} \quad (p \ltimes q) \Rightarrow \mathcal{G}}{\mathsf{ld}; \mathcal{G} \vdash \{p\}\mathbf{atomic}\{C\}\{q\}} (ATOM) \qquad \frac{\mathsf{ld}; \mathcal{G} \vdash \{p\}\mathbf{atomic}\{C\}\{q\}}{\mathcal{R}; \mathcal{G} \vdash \{p\}\mathbf{atomic}\{C\}\{q\}} (ATOM-R)$$

$$\frac{p \Rightarrow \mathsf{own}_{np}(y) * \mathbf{true}}{\mathcal{R}; \mathcal{G} \vdash \{p * \mathsf{own}(x) \land x \in y\} C; y := y \setminus \{x\} \{p * \mathsf{own}(x)\}}{\mathcal{R}; \mathcal{G} \vdash \{p * \mathsf{own}(x)\} \mathbf{foreach} x \mathbf{in} y \mathbf{do} C\{p * \mathsf{own}(x) \land y = \emptyset\}}$$
(P-FOREACH)

Fig. 26. Selected inference rules for GC verification.

The GC guarantees not modifying the mutator stores. For any mutator-reachable object, the GC does not update its fields coming from the high-level mutator, nor does it reclaim the object. Here $\lfloor _ \rfloor$ lifts a low-level object to a new one that contains mutator data only.

 $[o] \triangleq \{ pt_1 \quad o(pt_1), \dots, pt_m \quad o(pt_m), data \quad o(data) \}$

We could prove that \mathcal{G}_{gc} does not contain more behaviors than AbsGCStep.

$$\mathcal{G}_{gc} \circ \alpha^{-1} \subseteq \alpha^{-1} \circ \mathsf{AbsGCStep}$$

We present the proof of the top-level collection cycle in Figure 28. One of the key invariants used in the proofs is reach_inv (defined in Figure 27). It says, if a reachable BLACK object x points to a WHITE object y, then either x is dirty or a mutator is going to dirty x (the predicate todirty(x, y) holds). The latter occurs when the mutator thread t has done the first step of its write barrier update(x, fd, y). We have t.aux = x and from the mutator's guarantees (Figure 25(b)), we know t must be going to dirty x.

Since each instruction in the GC code is executed atomically, we need to stabilize the pre- and postconditions when verifying it (e.g., see the ATOM-R rule in Figure 26). For example, when reading a pointer field of an object to a local variable, the postcondition should be stabilized since mutators might update the field.

$$\begin{array}{l} \left\{ \exists X, Y. \ (j = Y) * (i.pt_1 \hookrightarrow X) \right\} \\ \mathcal{R}_{gc}; \mathcal{G}_{gc} \vdash \begin{array}{l} j := i.pt1; \\ \left\{ \exists X. \ (j = X) * ptfd_sta(i.pt_1, X) \right\} \end{array}$$

Here $ptfd_sta(i.pt_1, X)$ says either the pt_1 field of i is X, or i is (or is going to be) marked as dirty. Similarly, when reading the color of an object, the postcondition should take

ACM Transactions on Programming Languages and Systems, Vol. 36, No. 1, Article 3, Publication date: March 2014.

(**Dm**)

obj(x)	≜	$x.pt_1 \mapsto * x.pt_m \mapsto * x.data \mapsto * x.color \mapsto * x.dirty \mapsto * true$
blueobj(x)	\triangleq	$x.pt_1 \mapsto * x.pt_m \mapsto * x.data \mapsto * x.color \mapsto BLUE * x.dirty \mapsto$
newobj(x)	\triangleq	$x.\mathrm{pt}_1\mapsto 0*\ldots*x.\mathrm{pt}_m\mapsto 0*x.\mathrm{data}\mapsto 0*x.\mathrm{color}\mapsto \mathtt{BLACK}*x.\mathtt{dirty}\mapsto 0$
black(x)	\triangleq	$x.\texttt{color} \hookrightarrow \texttt{BLACK}$
white(x)	\triangleq	$x.\texttt{color} \hookrightarrow \texttt{WHITE}$
dirty(x)	\triangleq	$x.\texttt{dirty} \hookrightarrow 1$
not_blue(x)	\triangleq	$\exists c. \ (x. \texttt{color} \hookrightarrow c \land c \neq \texttt{BLUE})$
not_white(x)	≜	$\exists c. \ (x. \texttt{color} \hookrightarrow c \land c \neq \texttt{WHITE})$
not_dirty(x)	≜	$x.\texttt{dirty} \hookrightarrow 0$
root(t, S)	\triangleq	$\lambda(\pi, s, h)$. $\exists s_t. s_t = \pi(t) \land S = \{n \mid \exists x. s_t(x) = (n, 1) \land x \neq aux\}$
edge(x, y)	≜	$\exists fd \in \{ \mathrm{pt}_1, \dots, \mathrm{pt}_m \}. \ (x.fd \hookrightarrow y)$
nath, (r, y)	Δ	$\int x = y \qquad \qquad \text{if } k = 0$
part _k (x,y)	_	$\exists z. \operatorname{edge}(x,z) \wedge \operatorname{path}_{k-1}(z,y) \text{if } k > 0$
path(x,y)	≜	$\exists k. \operatorname{path}_k(x,y)$
reachable(t, x)	≜	$\exists S, y. \operatorname{root}(t, S) \land y \in S \land \operatorname{path}(y, x) \land x \neq 0$
reachable(x)	≜	$\exists t \in [1N]$. reachable (t, x)
wfstate	≜	$\circledast_{x \in [1M]}$.obj $(x) \land (\forall x. reachable(x) \Longrightarrow not_blue(x))$
white_edge(x, fd, y)	≙	$(x.fd \hookrightarrow y) \land white(y) \land fd \in \{pt_1, \dots, pt_m\}$
white_edge (x, y)	≜	$\exists fd.$ white_edge(x, fd, y)
todirty(x, n)	≜	$\exists t, S. (t.aux = x \land root(t, S) \land n \in S)$
instk(n,A)	≜	$\exists n', A'. A = n' :: A' \land (n = n' \lor instk(n, A'))$
$stk_black(A)$	≜	$\forall x. \operatorname{instk}(x, A) \Longrightarrow \operatorname{black}(x)$
reach_inv	≜	$\forall x, y. reachable(x) \land black(x) \land white_edge(x, y)$
		\implies dirty(x) \lor todirty(x, y)
$reach_stk(A)$	≜	$\forall x, y. reachable(x) \land black(x) \land white_edge(x, y)$
		\implies dirty(x) \lor todirty(x, y) \lor instk(x, A)
reach_tomk(A, x_p, S_f, x_n)	≜	$\forall x, fd, y. reachable(x) \land black(x) \land white_edge(x, fd, y)$
- ,		\implies dirty(x) \lor todirty(x, y) \lor instk(x, A) \lor (x = x _p \land fd \in S _f) \lor (y = x _n)
reach_black	\triangleq	$\forall x. \operatorname{reachable}(x) \Longrightarrow \operatorname{black}(x)$
$ptfd_sta(x.fd, y)$	\triangleq	$\exists n. (x.fd \hookrightarrow n) \land (y = n \lor dirty(x) \lor n = 0 \lor todirty(x, n))$
newobj_sta(x)	\triangleq	$obj(x) \land black(x) \land \forall fd \in \{pt_1, \dots, pt_m\}. ptfd_sta(x.fd, 0)$
rt_black(t)	\triangleq	$\exists S. \operatorname{root}(t, S) \land \forall n \in S. \operatorname{black}(n)$
rt_black	≜	$\forall t \in [1N]$. rt_black(t)
mark_rt_till(n)	\triangleq	$\forall t \in [1n]$. rt_black(t)
$clear_color_till(n)$	\triangleq	$\forall x \in [1n]$. (x.color \hookrightarrow BLACK \Longrightarrow newobj_sta(x))
$clear_dirty_till(n)$	≜	$\forall x \in [1n]$. not_dirty(x)
$reclaim_till(n)$	≜	$\forall x \in [1n]$. not_white(x)

Fig. 27. Useful assertions for verifying Boehm et al. GC.

into account the mutators' possible update of the color field in allocation and the updates of pointer fields after allocation.

$$\mathcal{R}_{\text{gc}}; \mathcal{G}_{\text{gc}} \vdash \begin{cases} \exists X, Y. \ (c = X) * (i.color \hookrightarrow Y) \end{cases}$$

$$c := i.color;$$

$$\exists X, Y. \ (c = X) * (i.color \hookrightarrow Y)$$

$$\land (X = Y \lor X = \text{BLUE} \land \text{newobj_sta}(i))$$

Here newobj_sta(i) says i points to a new object whose color field is BLACK, and each pointer field is either 0 or the object is dirty. Both the predicates $ptfd_sta$ and $newobj_sta$ are defined in Figure 27.

```
{wfstate}
Collection() {
   local mstk: Seq(Int);
    Loop Invariant: {wfstate *(own_{np}(mstk) \land mstk = \epsilon)}
   while (true) {
       Initialize();
      {(wfstate \land reach_inv) * (own<sub>np</sub>(mstk) \land mstk = \epsilon)}
      Trace();
      \{(wfstate \land reach_inv) * (own_{np}(mstk) \land mstk = \epsilon)\}
      CleanCard();
      \{(wfstate \land reach_inv) * (own_{np}(mstk) \land mstk = \epsilon)\}
       atomic{
          ScanRoot();
         \{\exists X.(\texttt{wfstate} \land \texttt{reach\_stk}(X) \land \texttt{stk\_black}(X) \land \texttt{rt\_black}) * (\texttt{own}_{np}(\texttt{mstk}) \land \texttt{mstk} = X)\}
          CleanCard();
       }
      \{(\texttt{wfstate} \land \texttt{reach\_black}) * (\texttt{own}_{np}(\texttt{mstk}) \land \texttt{mstk} = \epsilon)\}
      Sweep();
   }
}
{false}
```

Fig. 28. Proof outline of Collection().

The module MarkAndPush(i) will be called several times in the GC code, so we first give its general specification here. When the object i is white, MarkAndPush(i) colors it black and pushes it onto the mark stack.

Here as defined in Figure 27, reach_tomk(A, x_p, S_f, x_n) means, if a reachable BLACK object x points to a WHITE object y via the field fd, then one of the following cases holds.

- (1) $dirty(x) \lor todirty(x)$: x is (or is going to be) marked as dirty, as required in reach_inv;
- (2) instk(x,A): *x* is on the stack *A*;
- (3) $x = x_p \wedge fd \in S_f$: *x* is x_p , and *fd* is a field in S_f ;

(4)
$$y = x_n : y \text{ is } x_n$$
.

The case (2) will be useful during tracing when some objects have been colored black and pushed onto the stack. We define reach_stk to express that only cases (1) and (2)are satisfied. We will discuss the uses of the last two cases later.

Every collection cycle in Figure 28 begins from a well-formed state with an empty mark stack in the GC's local store. Then the GC does the following in order.

(1) Concurrent Initializing (Initialize(), shown in Figure 29). We use $clear_color_till(n)$ to mean that the GC has done color-clearing from locations 1 to n in the heap, but there might still be black objects since the mutators could allocate an black object after the GC's clearing. We could prove reach_inv holds when the GC has cleared the colors of all the objects in the heap, as shown in the following lemma.

```
{wfstate}
Initialize() {
    local i: [1..M], c: {BLACK, WHITE, BLUE};
    i := 1;
   Loop Invariant: {(wfstate \land clear_color_till(i - 1) \land 1 \leq i \leq M + 1) * own<sub>np</sub>(c)}
    while (i <= M) { ... } // See Figure 15 for the full code
}
{wfstate < reach_inv} // using Lemma 7.2
                                              Fig. 29. Proof outline of Initialize().
\{(wfstate \land reach_inv) * (own_{np}(mstk) \land mstk = \epsilon)\}
Trace() {
    local t: [1..N], rt: Set(Int), i: [0..M];
    t := 1;
                                   \begin{array}{l} (\texttt{wfstate} \land \texttt{reach\_inv}) * (\texttt{own}_{np}(\texttt{mstk}) \land \texttt{mstk} = \epsilon) \\ * (\texttt{own}_{np}(\texttt{t}) \land 1 \leq \texttt{t} \leq N+1) * \texttt{own}_{np}(\texttt{rt}) * \texttt{own}_{p}(\texttt{i}) \end{array} 
   Loop Invariant:
    while (t <= N) {
       rt := get_root(t);
      Foreach Invariant: {FInv}
        foreach i in rt do {
         \{FInv \land i \in rt\} || using Lemma 7.3
           MarkAndPush(i);
         \{FInv \land i \in rt\} || using Lemma 7.4
        }
       t := t + 1;
         \exists X. (wfstate \land reach\_stk(X) \land stk\_black(X)) * (own_{np}(mstk) \land mstk = X)
          *\;(\mathsf{own}_{np}(\mathtt{t}) \land 1 \leq \mathtt{t} \leq N+1) * \mathsf{own}_{np}(\mathtt{rt}) * \mathsf{own}_{p}(\mathtt{i})
        TraceStack();
        (\texttt{wfstate} \land \texttt{reach\_inv}) * (\texttt{own}_{np}(\texttt{mstk}) \land \texttt{mstk} = \epsilon) * (\texttt{own}_{np}(\texttt{t}) \land 1 \leq \texttt{t} \leq N+1)
          * \, \texttt{OWN}_{np}(\texttt{rt}) * \texttt{OWN}_p(\texttt{i})
}
\{(\texttt{wfstate} \land \texttt{reach\_inv}) * (\texttt{own}_{np}(\texttt{mstk}) \land \texttt{mstk} = \epsilon)\}
```

```
 \begin{array}{l} \text{where } \textit{FInv} \triangleq \exists \textit{X}. \ (\texttt{wfstate} \land \texttt{reach\_stk}(\textit{X}) \land \texttt{stk\_black}(\textit{X})) \ast (\texttt{own}_{np}(\texttt{mstk}) \land \texttt{mstk} = \textit{X}) \\ \ast (\texttt{own}_{np}(\texttt{t}) \land 1 \leq \texttt{t} \leq \textit{N}) \ast (\texttt{own}_{np}(\texttt{rt}) \land \forall n \in \texttt{rt}. \ 0 \leq n \leq \textit{M}) \ast \texttt{own}_{p}(\texttt{i}) \\ \end{array}
```

Fig. 30. Proof outline of Trace().

LEMMA 7.2. wfstate \land clear_color_till(M) \implies reach_inv.

That is, after initialization, if a BLACK reachable object x points to a WHITE object y, then x must be a newly allocated object whose pointer field is updated and dirty bit is (or is going to be) set to 1.

(2) Concurrent mark phase (Trace(), shown in Figure 30).

(a) The GC first calls MarkAndPush(i) to mark and push every root object. We need the following two lemmas to relate the unified pre- and postconditions of MarkAndPush(i) in (7.5) and the actual pre- and postconditions when calling the module.

LEMMA 7.3. reach_stk(X) \implies reach_tomk(X, 0, \emptyset, i).

LEMMA 7.4. reach_tomk($X, 0, \emptyset, 0$) \implies reach_stk(X).

Then by the CONSEQ rule, we can reuse the proof of MarkAndPush(i).

```
\{\exists X. (wfstate \land \mathsf{reach\_stk}(X) \land \mathsf{stk\_black}(X)) * (\mathsf{own}_{np}(\texttt{mstk}) \land \texttt{mstk} = X)\}
TraceStack() {
     local i: [1..M], j: [0..M];
                                         \begin{cases} \exists X. \; (\texttt{wfstate} \land \texttt{reach\_stk}(X) \land \texttt{stk\_black}(X)) \\ * \; (\texttt{own}_{np}(\texttt{mstk}) \land \texttt{mstk} = X) * \texttt{own}_{p}(\texttt{i}) * \texttt{own}_{p}(\texttt{j}) \end{cases} 
   Loop Invariant:
     while (!is_empty(mstk)) {
         i := pop(mstk);
         \exists X'. (wfstate \land reach\_stk(i :: X') \land stk\_black(X') \land obj(i))
          * \; (\texttt{own}_{np}(\texttt{mstk}) \land \texttt{mstk} = X') * \texttt{own}_p(j)
          j := i.pt1;
          \exists X'. (wfstate \land reach_stk(i :: X') \land stk_black(X') \land obj(i)
                           \wedge \, \texttt{ptfd\_sta}(\texttt{i.pt}_1,\texttt{j}) \wedge (\texttt{j} = 0 \lor \texttt{obj}(\texttt{j}))) \ast (\texttt{own}_{np}(\texttt{mstk}) \wedge \texttt{mstk} = X')
          \exists X'. (\texttt{wfstate} \land \texttt{reach\_tomk}(X',\texttt{i},\{\texttt{pt}_2,\ldots,\texttt{pt}_m\},\texttt{j}) \land \texttt{stk\_black}(X') \land (\texttt{j} = 0 \lor \texttt{obj}(\texttt{j}))
                            \wedge \ 1 \leq \mathtt{i} \leq M) \ast (\mathtt{OWN}_{\mathtt{np}}(\mathtt{mstk}) \wedge \mathtt{mstk} = X')
                                                                                                                                      // using Lemma 7.5
         MarkAndPush(j);
          \exists X'. (wfstate \land reach\_tomk(X', i, \{pt_2, \dots, pt_m\}, 0) \land stk\_black(X') \land (j = 0 \lor not\_white(j))
                           \wedge \ 1 \leq \mathtt{i} \leq M) \ast \ (\mathtt{OWN}_{\mathtt{np}}(\mathtt{mstk}) \wedge \mathtt{mstk} = X')
         j := i.ptm; MarkAndPush(j);
         \exists X'. (wfstate \land reach\_tomk(X', i, \emptyset, 0) \land stk\_black(X') \land (j = 0 \lor not\_white(j)))
           * (\mathsf{OWN}_{np}(\mathsf{mstk}) \land \mathsf{mstk} = X')
     }
}
\{(\texttt{wfstate} \land \texttt{reach\_inv}) * (\texttt{own}_{np}(\texttt{mstk}) \land \texttt{mstk} = \epsilon)\}
```

Fig. 31. Proof outline of TraceStack().

- (b) Then the GC calls the module TraceStack() (Figure 31) to perform the depth-first traversal. The loop invariant reach_stk holds at each time before the GC pops an object from the mark stack. Suppose the top object i on the mark stack points to a white object x. The GC does the following in order.
 - (i) Pop i. Then the black object i that points to *x* is not on the stack now.
 - (ii) Read the pt₁ field of i to a local variable j. As we explained before, i.pt₁ might not equal j since mutators could update this field. We only know that ptfd_sta(i.pt₁, j) holds. Then, x might be either j, or pointed to by i via fields pt₂,..., pt_m. Thus we get reach_tomk(mstk, i, {pt₂,..., pt_m}, j) holds. Formally, the following lemma holds. LEMMA 7.5.
 - (A) reach_stk(i :: X) \iff reach_tomk($X, i, \{pt_1, \dots, pt_m\}, 0$);
 - (B) reach_tomk($X, i, S_f, 0$) \implies reach_tomk(X, i, S_f, j);
 - $\begin{array}{ll} (C) \ \mbox{reach_tomk}(X,i,S_f,j) \land \mbox{ptfd_sta}(i.fd,j) \land fd \in S_f \\ \implies \mbox{reach_tomk}(X,i,S_f \backslash \{fd\},j). \end{array}$
 - (iii) MarkAndPush(j). We can reuse the proof of this module again.
 - (iv) Mark and push other children. The proof is similar to the preceding two steps, so we omit the discussions. Finally, reach_stk holds because no reachable white object needs to rely on the reachability from i (it could be reachable from a child of i which is on the stack now).

```
{(wfstate \land reach_inv) * (own<sub>np</sub>(mstk) \land mstk = \epsilon)}
CleanCard() {
     local i: [1..M], c: {BLACK, WHITE, BLUE}, d: {1, 0};
     i := 1;
    \begin{array}{l} \textbf{Loop Invariant:} \\ \left\{ \begin{array}{l} \exists X. \; (\texttt{wfstate} \land \texttt{reach\_stk}(X) \land \texttt{stk\_black}(X)) * (\texttt{own}_{np}(\texttt{mstk}) \land \texttt{mstk} = X) \\ * \; (\texttt{own}_p(\texttt{i}) \land \texttt{1} \leq \texttt{i} \leq M + \texttt{1}) * \texttt{own}_{np}(\texttt{c}) * \texttt{own}_{np}(\texttt{d}) \\ \texttt{while} \; (\texttt{i} <= \texttt{M}) \; \{ \; \dots \; \} \; /\!\!/ \; \texttt{See Figure 15 for the full code} \end{array} \right. 
       \exists X. (wfstate \land \mathsf{reach\_stk}(X) \land \mathsf{stk\_black}(X)) * (\mathsf{own}_{np}(\texttt{mstk}) \land \texttt{mstk} = X)
        * \textit{ Own}_p(\texttt{i}) * \textit{ Own}_{np}(\texttt{c}) * \textit{ Own}_{np}(\texttt{d})
     TraceStack();
  \{(\texttt{wfstate} \land \texttt{reach\_inv}) * (\texttt{own}_{np}(\texttt{mstk}) \land \texttt{mstk} = \epsilon) * \texttt{own}_{p}(\texttt{i}) * \texttt{own}_{np}(\texttt{c}) * \texttt{own}_{np}(\texttt{d})\}
}
{(wfstate \land reach_inv) * (own<sub>np</sub>(mstk) \land mstk = \epsilon)}
                                                            Fig. 32. Proof outline of CleanCard().
\{(wfstate \land reach_inv) * (own_{np}(mstk) \land mstk = \epsilon)\}
ScanRoot() {
     local t: [1..N], rt: Set(Int), i: [0..M];
     t := 1;
    Loop Invariant:
       \exists X. \; (\mathsf{wfstate} \land \mathsf{reach\_stk}(X) \land \mathsf{stk\_black}(X) \land \mathsf{mark\_rt\_till}(\mathtt{t}-1) \land 1 \leq \mathtt{t} \leq N+1)
       * (\mathsf{OWN}_{np}(\texttt{mstk}) \land \texttt{mstk} = X) * \mathsf{OWN}_{p}(\texttt{i}) * \mathsf{OWN}_{np}(\texttt{rt})
     while (t <= N) {
         rt := get_root(t);
        Foreach Invariant:
          \exists X, Y. (wfstate \land \mathsf{reach\_stk}(X) \land \mathsf{stk\_black}(X) \land \mathsf{mark\_rt\_till}(\texttt{t}-1) \land 1 \leq \texttt{t} \leq N \land \mathsf{root}(\texttt{t}, Y)
                    \land \forall n \in (Y \setminus \texttt{rt}). \ \texttt{black}(n) \land \texttt{rt} \subseteq Y) * (\texttt{own}_{np}(\texttt{mstk}) \land \texttt{mstk} = X) * \texttt{own}_{p}(\texttt{i})
          foreach i in rt do { MarkAndPush(i); }
          t := t + 1;
     }
}
\{\exists X. (wfstate \land reach\_stk(X) \land stk\_black(X) \land rt\_black) * (own_{np}(mstk) \land mstk = X)\}
```



After tracing, we can ensure reach_inv still holds. That is, if a black object x points to a white object, then x must be (or is going to be) dirty and its pointer field is updated by the mutators.

- (3) Concurrent card-cleaning (CleanCard(), as shown in Figure 32). We reuse the proof of TraceStack() via the frame rule. We can conclude reach_inv is maintained at the end of this phase.
- (4) Stop-the-world card-cleaning.
 - (a) The GC first rescans the roots (ScanRoot(), shown in Figure 33) as if they were dirty. Then reach_stk and rt_black hold. rt_black says all the root objects are black. Moreover, all the objects on the stack are black (stk_black). The atomic MarkAndPush(i) is proved similarly to the concurrent one (7.5) with the same pre- and postconditions.
 - (b) Then the GC cleans the cards (the atomic CleanCard(), shown in Figure 34) without the interference from the mutators. At the end, the mark stack is empty and all the reachable objects are black (denoted by reach_black). The

```
 \begin{split} & \{\exists X. (wfstate \land reach\_stk(X) \land stk\_black(X) \land rt\_black) * (own_{np}(mstk) \land mstk = X)\} \\ & \text{CleanCard}() \ \\ & \text{local i: [1..M], c: } \{BLACK, WHITE, BLUE\}, d: \{1, 0\}; \\ & \text{i := 1;} \\ & \textit{Loop Invariant:} \\ & \exists X. (wfstate \land reach\_stk(X) \land stk\_black(X) \land rt\_black \land clear\_dirty\_till(i-1) \\ & \land 1 \leq i \leq M+1) * (own_{np}(mstk) \land mstk = X) * own_{np}(c) * own_{np}(d) \\ & \text{while (i <= M) } \{ \dots \} /\!\!/ \text{See Figure 15 for the full code} \\ & \exists X. (wfstate \land reach\_stk(X) \land stk\_black(X) \land rt\_black \land clear\_dirty\_till(M)) \\ & * (own_{np}(mstk) \land mstk = X) * own_{np}(c) * own_{np}(d) \\ & \text{TraceStack();} \\ & (wfstate \land reach\_inv \land rt\_black \land clear\_dirty\_till(M)) \\ & * (own_{np}(mstk) \land mstk = \epsilon) * own_{np}(c) * own_{np}(d) \\ & \} \\ & \{(wfstate \land reach\_black) * (own_{np}(mstk) \land mstk = \epsilon)\} \\ \end{split}
```



```
 \label{eq:state_reach_black} $$ weep() { $$ local i: [1..M], c: {BLACK, WHITE, BLUE}; $$ i := 1; $$ Loop Invariant: {(wfstate \land reach_black \land reclaim_till(i - 1) \land 1 \le i \le M + 1) * own_{np}(c)}$$ while (i <= M) { ... } // See Figure 15 for the full code $$ wfstate \land reach_black \land reclaim_till(M)} $$
```

Fig. 35. Proof outline of Sweep().

proof for the atomic TraceStack() is similar to the proof of the concurrent one and omitted here.

(5) Concurrent sweep phase (Sweep(), shown in Figure 35). No matter how the mutators interleave with the GC, all the white objects remain unreachable. Thus the reclamation is safe that guarantees \mathcal{G}_{gc} . After sweep, the state is still well-formed.

8. RELATED WORK AND CONCLUSION

There is a large body of work on refinements and verification of program transformations. Here we only focus on the work most closely related to the typical applications discussed in this article.

Verifying compilation and optimizations of concurrent programs. Compiler verification for concurrent programming languages can date back to work in Wand [1995] and Gladstein and Wand [1996], which is about functional languages using messagepassing mechanisms. Recently, Lochbihler [2010] presented a verified compiler for Java threads and proved semantics preservation by a weak bisimulation. He views every heap update as an observable move, thus does not allow the target and the source to have different granularities of atomic updates. To achieve parallel compositionality, he requires the relation to be preserved by any transitions of shared states, that is, the environments are assumed arbitrary. As we explained in Section 2.2, this is a too strong requirement in general for many transformations, including the examples in this article.

Burckhardt et al. [2010] present a proof method for verifying concurrent program transformations on relaxed memory models. The method relies on a compositional

trace-based denotational semantics, where the values of shared variables are always considered arbitrary at any program point. In other words, they also assume arbitrary environments.

Following Leroy's CompCert project [Leroy 2009], Sevčík et al. [2011] verify compilation from a C-like concurrent language to x86 by simulations. They focus on correctness of a particular compiler, and there are two phases in their compiler whose proofs are not compositional. Here we provide a general, compiler-independent compositional proof technique to verify concurrent transformations.

We apply RGSim to justify concurrent optimizations, following Benton [2004] who presents a declarative set of rules for sequential optimizations. Also the proof rules of RGSim for sequential compositions, conditional statements, and loops coincide with those in relational Hoare logic [Benton 2004] and relational separation logic [Yang 2007].

Proving linearizability or atomicity of concurrent objects. Filipović et al. [2010] show linearizability can be characterized in terms of an observational refinement, where the latter is defined similarly to our Correct(T). There is no proof method given to verify the linearizability of fine-grained object implementations.

Turon and Wand [2011] propose a refinement-based proof method to verify concurrent objects. They first propose a simple refinement based on Brookes' fully abstract trace semantics [Brookes 1996], which is compositional but cannot handle complex algorithms (as discussed in Section 2.2). Their fenced refinement then uses rely conditions to filter out illegal environment transitions. The basic idea is similar to ours, and the refinement can also be used to verify Treiber's stack algorithm. However, it is "not a congruence for parallel composition". In their settings, both the concrete (fine-grained) and the abstract (atomic) versions of object operations need to be expressed in the same language. They also require that the fine-grained implementation should have only one update action over the shared state to correspond to the high-level atomic operation. These requirements and the lack of parallel compositionality limit the applicability of their method. It is unclear if the method can be used for general verification of transformations, such as concurrent GCs.

Elmas et al. [2010] prove linearizability by incrementally rewriting the fine-grained implementation to the atomic abstract specification. Their behavioral simulation used to characterize linearizability is an event-trace subset relation with requirements on the orders of method invocations and returns. Their rules heavily rely on movers (i.e., operations that can commute over any operation of other threads) and always rewrite programs to instructions, thus are designed specifically for atomicity verification. Compositionality is not considered in their work.

In his thesis [2008], Vafeiadis proves linearizability of concurrent objects in RGSep logic by introducing abstract objects and abstract atomic operations as auxiliary variables and code. The refinement between the concrete implementation and the abstract operation is implicitly embodied in the unary verification process, but is not spelled out formally in the metatheory (e.g., the soundness).

Verifying concurrent GCs. Vechev et al. [2006] define transformations to generate concurrent GCs from an abstract collector. Afterwards, Pavlovic et al. [2010] present refinements to derive concrete concurrent GCs from specifications. These methods focus on describing the behaviors of variants (or instantiations) of a correct abstract collector (or a specification) in a single framework, assuming all the mutator operations are atomic. By comparison, we provide a general correctness notion and a proof method for verifying concurrent GCs and the interactions with mutators (where the barriers could be fine grained). Furthermore, the correctness of their transformations

or refinements is expressed in a GC-oriented way (e.g., the target GC should mark no less objects than the source), which cannot be used to justify other transformations.

Kapoor et al. [2011] verify Dijkstra's GC using concurrent separation logic. To validate the GC specifications, they also verify a representative mutator in the same system. In contrast, we reduce the problem of verifying a concurrent GC to verifying a transformation, ensuring semantics preservation for *all* mutators. Our GC verification framework is inspired by McCreight et al. [2007], who propose a framework for separate verification of stop-the-world and incremental GCs and their mutators, but their framework does not handle concurrency.

Conclusion and future work. We propose RGSim to verify concurrent program transformations. By describing explicitly the interference with environments, RGSim is compositional, and can support many widely used transformations. We have applied RGSim to reason about optimizations, prove atomicity of fine-grained concurrent algorithms, and verify concurrent garbage collectors.

The compositionality of RGSim allows us to decompose the refinement for a large program to refinements for basic transformation units (which are usually instructions). However, for those transformation units, we have to refer to the semantics of RGSim (Definition 4.2) rather than syntactic rules to verify them, since Figure 7 provides only compositionality rules, with no rules for primitive instructions. This makes the proofs a bit tedious and complicated. Also, RGSim cannot verify the atomicity of concurrent algorithms with helping mechanism or speculations, such as the RDCSS algorithm [Vafeiadis 2008]. Finally, as we mentioned in Section 4.3, RGSim cannot ensure preservation of termination when establishing refinements. In the future, we would like to extend RGSim with a more complete set of proof rules and with the support of liveness verification. We also hope to further test its applicability with more applications, such as verifying STM implementations and compilers. It is also interesting to explore the possibility of building tools to automate the verification process.

ACKNOWLEDGMENTS

We would like to thank Matthew Parkinson and anonymous referees for their suggestions and comments on earlier versions of this article; Pierre Castéran and Sandrine Blazy for their generous help on the Coq implementation and understanding co-induction; Aaron Turon for the insightful discussions on comparing their and our work.

REFERENCES

- Martín Abadi and Leslie Lamport. 1991. The existence of refinement mappings. *Theor. Comput. Sci.* 82, 2, 253–284.
- Martín Abadi and Gordon Plotkin. 2009. A model of cooperative threads. In Proceedings of the 36th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL'09). ACM Press, New York, 29–40.
- Katherine Barabash, Ori Ben-Yitzhak, Irit Goft, Elliot K. Kolodner, Victor Leikehman, Yoav Ossia, Avi Owshanko, and Erez Petrank. 2005. A parallel, incremental, mostly concurrent garbage collector for servers. ACM Trans. Program. Lang. Syst. 27, 6, 1097–1146.
- Nick Benton. 2004. Simple relational correctness proofs for static analyses and program transformations. In Proceedings of the 31st ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL'04). ACM Press, New York, 14–25.
- Nick Benton and Chung-Kil Hur. 2009. Biorthogonality, step-indexing and compiler correctness. In Proceedings of the 14th ACM SIGPLAN International Conference on Functional Programming (ICFP'09). ACM Press, New York, 97–108.
- Hans-Juergen Boehm. 2005. Threads cannot be implemented as a library. In Proceedings of the ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI'05). ACM Press, New York, 261–268.

- Hans-Juergen Boehm and Sarita V. Adve. 2008. Foundations of the C++ concurrency memory model. In Proceedings of the ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI'08). ACM Press, New York, 68–78.
- Hans-Juergen Boehm, Alan J. Demers, and Scott Shenker. 1991. Mostly parallel garbage collection. In Proceedings of the ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI'91). ACM Press, New York, 157–164.
- Stephen D. Brookes. 1996. Full abstraction for a shared-variable parallel language. Inf. Comput. 127, 2, 145–163.
- Sebastian Burckhardt, Madanlal Musuvathi, and Vasu Singh. 2010. Verifying local transformations on relaxed memory models. In Proceedings of the 19th Joint European Conference on Theory and Practice of Software and the International Conference on Compiler Construction (CC'10/ETAPS'10). Springer, 104–123.
- Coq Development Team. 2010. The Coq proof assistant reference manual. The Coq release v8.3. http://coq.inria.fr/V8.3/refman/
- David Dice, Ori Shalev, and Nir Shavit. 2006. Transactional locking ii. In Proceedings of the 20th International Conference on Distributed Computing (DISC'06). Springer, 194–208.
- Tayfun Elmas, Shaz Qadeer, Ali Sezgin, Omer Subasi, and Serdar Tasiran. 2010. Simplifying linearizability proofs with reduction and abstraction. In Proceedings of the 16th International Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS'10). Springer, 296–311.
- Xinyu Feng. 2009. Local rely-guarantee reasoning. In Proceedings of the 36th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL'09). ACM Press, New York, 315–327.
- Ivana Filipović, Peter O'Hearn, Noam Rinetzky, and Hongseok Yang. 2010. Abstraction for concurrent objects. *Theor. Comput. Sci.* 411, 51–52, 4379–4398.
- David S. Gladstein and Mitchell Wand. 1996. Compiler correctness for concurrent languages. In Proceedings of the 1st International Conference on Coordination Languages and Models (COORDINATION'96). Lecture Notes in Computer Science, vol. 1061, Springer, 231–248.
- Maurice Herlihy and Nir Shavit. 2008. The Art of Multiprocessor Programming. Morgan Kaufmann, San Fransisco.
- Charles A. R. Hoare. 1972. Proof of correctness of data representations. Acta Inf. 1, 4, 271-281.
- Chung-Kil Hur and Derek Dreyer. 2011. A Kripke logical relation between ML and assembly. In Proceedings of the 38th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL'11). ACM Press, New York, 133–146.
- Cliff B. Jones. 1983. Tentative steps toward a development method for interfering programs. ACM Trans. Program. Lang. Syst. 5, 4, 596-619.
- Kalpesh Kapoor, Kamal Lodaya, and Uday Reddy. 2011. Fine-grained concurrency with separation logic. J. Philos. Logic 40, 5, 583-632.
- Xavier Leroy. 2009. A formally verified compiler back-end. J. Autom. Reason. 43, 4, 363-446.
- Hongjin Liang, Xinyu Feng, and Ming Fu. 2012. A rely-guarantee-based simulation for verifying concurrent program transformations. In Proceedings of the 39th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL'12). ACM Press, New York, 455–468.
- Andreas Lochbihler. 2010. Verifying a compiler for java threads. In Proceedings of the 19th European Conference on Programming Languages and Systems (ESOP'10). Springer, 427–447.
- Andrew McCreight, Zhong Shao, Chunxiao Lin, and Long Li. 2007. A general framework for certifying garbage collectors and their mutators. In Proceedings of the ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI'07). ACM Press, New York, 468–479.
- Matthew Parkinson, Richard Bornat, and Cristiano Calcagno. 2006. Variables as resource in hoare logics. In Proceedings of the 21st Annual IEEE Symposium on Logic in Computer Science (LICS'06). IEEE Computer Society, 137–146.
- Dusko Pavlovic, Peter Pepper, and Douglas R. Smith. 2010. Formal derivation of concurrent garbage collectors. In Proceedings of the 10th International Conference on Mathematics of Program Construction (MPC'10). 353–376.
- John C. Reynolds. 2002. Separation logic: A logic for shared mutable data structures. In Proceedings of the 17th Annual IEEE Symposium on Logic in Computer Science (LICS'02). IEEE Computer Society, 55–74.
- Jaroslav Ševčík, Viktor Vafeiadis, Francesco Zappa Nardelli, Suresh Jagannathan, and Peter Sewell. 2011. Relaxed-memory concurrency and verified compilation. In Proceedings of the 38th ACM SIGPLANSIGACT Symposium on Principles of Programming Languages (POPL'11). ACM Press, New York, 43–54.

- R. Kent Treiber. 1986. System programming: Coping with parallelism. Tech. rep. RJ 5118, IBM Almaden Research Center.
- Aaron Turon and Mitchell Wand. 2011. A separation logic for refining concurrent objects. In Proceedings of the 38th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL'11). ACM Press, New York, 247–258.
- Viktor Vafeiadis. 2008. Modular fine-grained concurrency verification. Tech. rep. UCAM-CL-TR-726, University of Cambridge, Computer Laboratory.
- Viktor Vafeiadis and Matthew J. Parkinson. 2007. A marriage of rely/guarantee and separation logic. In Proceedings of the 18th International Conference on Concurrency Theory (CONCUR'07). Springer, 256–271.
- Martin T. Vechev, Eran Yahav, and David F. Bacon. 2006. Correctness-preserving derivation of concurrent garbage collection algorithms. In Proceedings of the ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI'06). ACM Press, New York, 341–353.
- Mitchell Wand. 1995. Compiler correctness for parallel languages. In Proceedings of the 7th International Conference on Functional Programming Languages and Computer Architecture (FPCA'95). ACM Press, New York, 120–134.

Hongseok Yang. 2007. Relational separation logic. Theor. Comput. Sci. 375, 1-3, 308-334.

Received January 2013; revised September 2013; accepted November 2013