Owicki-Gries Reasoning for Weak Memory Models

Ori Lahav and Viktor Vafeiadis

Max Planck Institute for Software Systems (MPI-SWS)

Abstract. We show that even in the absence of auxiliary variables, the well-known Owicki-Gries method for verifying concurrent programs is unsound for weak memory models. By strengthening its non-interference check, however, we obtain OGRA, a program logic that is sound for reasoning about programs in the release-acquire fragment of the C11 memory model. We demonstrate the usefulness of this logic by applying it to several challenging examples, ranging from small litmus tests to an implementation of the RCU synchronization primitives.

1 Introduction

In 1976, Owicki and Gries [10] introduced a proof system for reasoning about concurrent programs, which formed the basis of rely/guarantee reasoning. Their system includes the usual Hoare logic rules for sequential programs, a rule for introducing auxiliary variables, and the following parallel composition rule:

$$\frac{\{P_1\}\, c_1\, \{Q_1\} \quad \{P_2\}\, c_2\, \{Q_2\} \quad \text{the two proofs are non-interfering}}{\{P_1 \wedge P_2\}\, c_1 \parallel c_2\, \{Q_1 \wedge Q_2\}}$$

This rule allows one to compose two verified programs into a verified concurrent program that assumes both preconditions and ensures both postconditions. The soundness of this rule requires that the two proofs are *non-interfering*, namely that every assertion R in the one proof is stable under any $\{P\}x := e$ (guarded) assignment in the other and vice versa; i.e., for every such pair, $R \wedge P \vdash R[e/x]$.

The Owicki-Gries system (OG) assumes a fairly simple but unrealistic concurrency model: *sequential consistency* (SC) [7]. This is essential: OG is complete for verifying concurrent programs under SC [12], and is therefore unsound under a weakly consistent memory semantics, such as TSO [9]. Auxiliary variables are instrumental in achieving completeness—without them, OG is blatantly incomplete; e.g., it cannot verify that $\{x = 0\}\, x \overset{\mathsf{at}}{:=} x + 1 \parallel x \overset{\mathsf{at}}{:=} x + 1\, \{x = 2\}$ (where "$\overset{\mathsf{at}}{:=}$" denotes atomic assignment).

Nevertheless, many useful OG proofs do not use auxiliary variables, and one might wonder whether such proofs are sound under weak memory models. This is sadly not the case. Figure 1 presents an OG proof that a certain program cannot return $a = b = 0$ whereas under all known weak memory models it can in fact do so. Intuitively speaking, the proof is invalid under weak memory because the two threads may have different views of memory before executing each command. Thus, when the second thread terminates, the first thread may perform $a := y$ reading $y = 0$ and storing 0 in a, thereby

Due to space limits, supplementary material including full proofs and further examples is available at: `http://plv.mpi-sws.org/ogra/`.

$\{x = 0 \land y = 0 \land a \neq 0\}$

Thread I	Thread II
$\{a \neq 0\}$	$\{\top\}$
$x := 1;$	$y := 1;$
$\{x \neq 0\}$	$\{y \neq 0\}$
$a := y$	$b := x$
$\{x \neq 0\}$	$\{y \neq 0 \land (a \neq 0 \lor b = x)\}$

$\{a \neq 0 \lor b \neq 0\}$

Non-interference checks are trivial. For example,

$$y \neq 0 \land (a \neq 0 \lor b = x) \land a \neq 0 \vdash y \neq 0 \land (a \neq 0 \lor b = 1)$$

and

$$y \neq 0 \land (a \neq 0 \lor b = x) \land x \neq 0 \vdash y \neq 0 \land (y \neq 0 \lor b = x)$$

show stability of the last assertion of thread II under $\{a \neq 0\}x := 1$ and $\{x \neq 0\}a := y$.

Fig. 1. OG proof that the "store buffering" program cannot return $a = b = 0$. This can also be proved in the restricted OG system with one (stable) global invariant [11]. Note that OG's "indivisible assignments" condition (3.1) is met: assignments mention at most one shared location.

invalidating the second thread's last assertion. We note that $y = 0$ was also readable by the second thread, albeit at an earlier point (before the $y := 1$ assignment). This is no accident, and this observation is essential for soundness of our proposed alternative.

In this paper we identify a stronger non-interference criterion that does not assume SC semantics. Thus, while considering the effect of an assignment $\{P\}x := y$ in thread I on the validity of an assertion R in thread II, one does not get to assume that R holds for the view of thread I while reading y. In fact, in some executions, the value read for y might even be inconsistent with R. Instead, the only allowed assumption is that some assertion that held not later than R in thread II was true while reading y. Thus our condition for checking stability of R under $\{P\}x := y$ is that $R \land P \vdash R[v/x]$ for every value v of y that is consistent with P and some non-later assertion of thread II.

We show that OG with our stronger non-interference criterion is sound under the release-acquire (RA) fragment of the C11 memory model [6], which exhibits a good balance between performance and mathematical sanity (see, e.g., [16,17]). Soundness under TSO follows, as TSO behaviors are all observable under RA (see [1]). Formalizing the aforementioned intuitions into a soundness proof for RA executions is far from trivial. Indeed, RA is defined axiomatically without an operational semantics and without the notion of a state. As a basis for the soundness proof, we introduce such a notion and study the properties of sequences of states observed by different threads.

We believe that the results of this paper may provide new insights for understanding weak memory models, as well as a *simple* and *useful* method for proving partial correctness of concurrent programs. We demonstrate the applicability of our logic (which we call OGRA) with several challenging examples, ranging from small litmus tests to an implementation of the read-copy-update (RCU) synchronization primitives [3]. We also provide support for fence instructions by implementing them as RMWs to an otherwise unused location and for a simple class of auxiliary variables, namely *ghost values*.

Related Work. Aiming to understand and verify high-performance realistic concurrent programs, program logics for weak memory models have recently received a lot of attention (see, e.g., [4,13,18,16,14]). Most of these logics concern the TSO memory model. Only two—RSL [18] and GPS [16]—can handle RA, but have a fairly complex foundation being based on separation logic. The most advanced of the two logics, GPS, has been used (with considerable ingenuity) to verify the RCU synchronization primitives [15], but simpler examples such as "read-read coherence" seem to be beyond its power (see Fig. 8). Finally, Cohen [2] studies an alternative memory model under which OG reasoning can be performed at the execution level.

2 Preliminaries

In this section, we present a simplified programming language, whose semantics adheres to that of the release-acquire fragment of C11's memory model [1]. We assume a finite set of locations $\mathsf{Loc} = \{\nu_1, \dots, \nu_M\}$, a finite set Val of values with a distinguished value $0 \in \mathsf{Val}$, and any standard interpreted language for expressions containing at least all locations and values. We use x, y, z as metavariables for locations, v for values, e for expressions, and denote by $e(x_1, \dots, x_n)$ an expression in which $x_1, \dots, x_n$ are the only mentioned locations. The language's commands are given by the following grammar:

$$c ::= \texttt{skip} \mid \texttt{if}\ e(x)\ \texttt{then}\ c\ \texttt{else}\ c \mid \texttt{while}\ e(x)\ \texttt{do}\ c \mid c\,;c \mid c \parallel c \mid \\ x := v \mid x := e(y) \mid x \overset{y,z}{:=} e(y,z) \mid x \overset{\mathsf{at}}{:=} e(x)$$

To keep the presentation simple, expressions in assignments are limited to mention at most two locations, and those in conditionals and loops mention one location. Assignments of expression mentioning two locations also specify the order in which these locations should be read (if one of them is local, this has no observable effect). The command $x \overset{\mathsf{at}}{:=} e(x)$ is an atomic assignment corresponding to a primitive read-modify-write (RMW) instruction and, as such, mentions only one location.[1]

Now, as in the C11 formalization, the semantics of a program is defined to be its set of consistent executions [1]. An execution G is a triple $\langle A, L, E\rangle$ where:

- $A \subseteq \mathbb{N}$ is a finite set of *nodes*. We identify G with this set, e.g., when writing $a \in G$.
- L is a function assigning a *label* to each node, where a label is either $\langle \mathtt{S}\rangle$ ("Skip"), a triple of the form $\langle \mathtt{R}, x, v_r\rangle$ ("Read"), a triple of the form $\langle \mathtt{W}, x, v_w\rangle$ ("Write"), or a quadruple of the form $\langle \mathtt{U}, x, v_r, v_w\rangle$ ("Update"). For $\mathtt{T} \in \{\mathtt{S}, \mathtt{R}, \mathtt{W}, \mathtt{U}\}$, we denote by $G.\mathtt{T}$ the set of nodes $a \in A$ for which $\mathtt{T}$ is the first entry of $L(a)$, while $G.\mathtt{T}_x$ denotes the set of $a \in G.\mathtt{T}$ for which x is the second entry of $L(a)$. In addition, L induces the partial functions $G.loc : A \to \mathsf{Loc}$, $G.val_r : A \to \mathsf{Val}$, and $G.val_w : A \to \mathsf{Val}$ that respectively return (when applicable) the x, v_r and v_w components of a node.
- $E \subseteq (A \times A) \cup (A \times A \times \mathsf{Loc})$ is a set of *edges*, such that for every triple $\langle a, b, x\rangle \in E$ (*reads-from* edge) we have $a \in G.\mathtt{W}_x \cup G.\mathtt{U}_x$, $b \in G.\mathtt{S} \cup G.\mathtt{R}_x \cup G.\mathtt{U}_x$, and $G.val_w(a) = G.val_r(b)$ whenever $b \notin G.\mathtt{S}$.[2] The subset $E \cap (A \times A)$ is denoted by $G.po$ (*program order*), and $G.E_x$ denotes the set $\{\langle a, b\rangle \in A \times A \mid \langle a, b, x\rangle \in E\}$ (*x-reads-from*) for every $x \in \mathsf{Loc}$. Finally, $G.E_{all}$ denotes the set $G.po \cup \bigcup_{x \in \mathsf{Loc}} E_x$.

For all these notations, we often omit the "$G.$" prefix when it is clear from the context. Given an execution $G = \langle A, L, E\rangle$ and a set E' of edges we write $G \cup E'$ for the triple $\langle A, L, E \cup E'\rangle$ and $G \setminus E'$ for $\langle A, L, E \setminus E'\rangle$.

Definition 1. A node a in an execution G is *initial* (*terminal*) in G if $\langle b, a\rangle \notin E_{all}$ ($\langle a, b\rangle \notin E_{all}$) for every $b \in G$. An edge $\langle a, b\rangle \in po$ is *initial* (*terminal*) in G if a is initial (b is terminal) in G.

Definition 2. Let $G = \langle A, L, E\rangle$ and $G' = \langle A', L', E'\rangle$ be two executions with disjoint sets of nodes.

[1] Unlike usual OG [10], our assignments can mention more than one shared variable. In fact, our formal development does not differentiate between local and shared variables.

[2] Reads-from edges $\langle a, b, x\rangle$ with $b \in G.\mathtt{S}$ are used for defining visible states (see Definition 7).

$$
\begin{array}{ll}
[\![\texttt{skip}]\!] & = \mathcal{SG} \\
[\![\texttt{if } e(x) \texttt{ then } c_1 \texttt{ else } c_2]\!] & = \bigcup\{\mathcal{RG}(x,v); [\![c_i]\!] \mid v \in \mathsf{Val}, i \in \{1,2\}, [\![e]\!](v) = 0 \textit{ iff } i = 2\} \\
[\![\texttt{while } e(x) \texttt{ do } c]\!] & = \bigcup_{n\geq 0}(\bigcup\{\mathcal{RG}(x,v) \mid v \in \mathsf{Val}, [\![e]\!](v) \neq 0\}; [\![c]\!])^n; \\
 & \quad \bigcup\{\mathcal{RG}(x,v) \mid v \in \mathsf{Val}, [\![e]\!](v) = 0\} \\
[\![c_1; c_2]\!] & = [\![c_1]\!] ; [\![c_2]\!] \\
[\![c_1 \parallel c_2]\!] & = \mathcal{SG}; ([\![c_1]\!] \parallel [\![c_2]\!]); \mathcal{SG} \\
[\![x := v]\!] & = \mathcal{WG}(x,v) \\
[\![x := e(y)]\!] & = \bigcup\{\mathcal{RG}(y,v); \mathcal{WG}(x, [\![e]\!](v)) \mid v \in \mathsf{Val}\} \\
[\![x \overset{y,z}{:=} e(y,z)]\!] & = \bigcup\{\mathcal{RG}(y,v_y); \mathcal{RG}(z,v_z); \mathcal{WG}(x, [\![e]\!](v_y,v_z)) \mid v_y, v_z \in \mathsf{Val}\} \\
[\![x \overset{\mathsf{at}}{:=} e(x)]\!] & = \bigcup\{\mathcal{UG}(x,v,[\![e]\!](v)) \mid v \in \mathsf{Val}\}
\end{array}
$$

Fig. 2. Mapping of commands to sets of executions.

- The execution $G \parallel G'$ is given by $\langle A \cup A', E \cup E', L \cup L' \rangle$.
- The execution $G; G'$ is given by $(G \parallel G') \cup (O \times I)$, where O is the set of terminal nodes of G, and I is the set of initial nodes of G'.
- Given $n \geq 0$, G^n is inductively defined by $G^0 = \langle \emptyset, \emptyset, \emptyset \rangle$ and $G^{n+1} = G^n; G$.

The above operations are extended to sets of executions in the obvious way (e.g., $\mathcal{G}; \mathcal{G}' = \{G; G' \mid G \in \mathcal{G}, G' \in \mathcal{G}', G; G' \text{ is defined}\}$).

Definition 3. Given $x \in \mathsf{Loc}$ and $v \in \mathsf{Val}$, an $\langle x, v\rangle$*-read gadget* is any execution of the form $\langle \{a\}, \{a \mapsto \langle \mathtt{R}, x, v\rangle\}, \emptyset\rangle$. $\langle x, v\rangle$*-write gadgets*, $\langle x, v_r, v_w\rangle$*-update gadgets* and *skip gadgets* are defined similarly. $\mathcal{RG}(x, v)$, $\mathcal{WG}(x, v)$, $\mathcal{UG}(x, v_r, v_w)$ and $\mathcal{SG}$ denote, respectively, the sets of all $\langle x, v\rangle$-read gadgets, all $\langle x, v\rangle$-write gadgets, all $\langle x, v_r, v_w\rangle$-update gadgets, and all skip gadgets.

Using these definitions, the mapping of commands to (sets of) executions is given in Fig. 2. Note that every execution $G \in [\![c]\!]$ for some command c satisfies $G.E_{all} = G.po$, and has a unique initial node that can reach any node, and a unique terminal node that can be reached from any node. We refer to such executions as *plain*. However, many of these executions are nonsensical as they can, for instance, read values never written in the program. We restrict our attention to *consistent* executions, as defined next.

Definition 4. A relation R is called a *modification order* for a location $x \in \mathsf{Loc}$ in an execution G if the following hold: (i) R is a total strict order on $\mathtt{W}_x \cup \mathtt{U}_x$; (ii) if $\langle a, b\rangle \in E^*_{all}$ then $\langle b, a\rangle \notin R$; (iii) if $\langle a, b\rangle \in E^+_{all}$ and $\langle c, b\rangle \in E_x$ then $\langle c, a\rangle \notin R$; and (iv) if $\langle a, b\rangle, \langle b, c\rangle \in R$ and $c \in \mathtt{U}$ then $\langle a, c\rangle \notin E_x$.

Definition 5. An execution $G = \langle A, L, E\rangle$ is called:
- *complete* if for every $b \in \mathtt{R} \cup \mathtt{U}$, we have $\langle a, b\rangle \in E_{loc(b)}$ for some $a \in \mathtt{W} \cup \mathtt{U}$.
- *coherent* if E_{all} is acyclic, and there is a modification order in G for each $x \in \mathsf{Loc}$.
- *consistent* if $G \cup E'$ is complete and coherent for some $E' \subseteq A \times A \times \mathsf{Loc}$.

To illustrate these definitions, Fig. 3 depicts a consistent non-SC execution of the "store buffering" program of Fig. 1 together with the implicit variable initializations.

While our notations are slightly different, the axiomatic semantics presented above corresponds to the semantics of C11 programs (see [1]) in which all locations are atomic, reads are acquire reads, writes are release writes, and updates are acquire-release RMWs. In addition, we do not allow reads from uninitialized locations. C11's "happens-before" relation corresponds to our E^+_{all}.

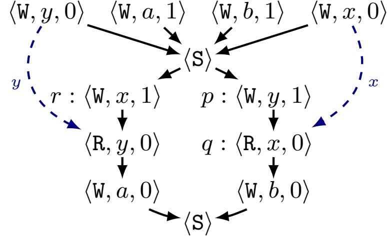

Fig. 3. Ignoring the dashed edges, this graph G is an initialized execution of the "store buffering" program (i.e., $G \in \mathcal{WG}(\top); [\![c]\!]$, Def. 8). G is consistent as it can be extended with the set E' of the two dashed reads-from edges.

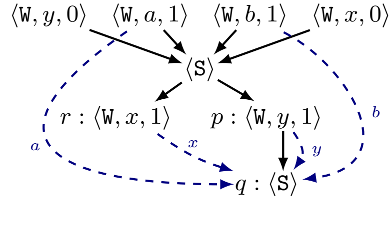

Fig. 4. Ignoring the dashed edges, we have the snapshot of $G \cup E'$ of Fig. 3 at $\langle p, q\rangle$ with respect to $\{r\}$. Adding the dashed edges results in a coherent execution; so the state $\{x \mapsto 1, y \mapsto 1, a \mapsto 1, b \mapsto 1\}$ is visible at $\langle p, q\rangle$ in $G \cup E'$.

3 An Owicki-Gries Proof System for Release-Acquire

In this section, we present OGRA—our logic for reasoning about concurrent programs under release-acquire. As usual, the basic constructs are *Hoare triples* of the form $\{P\}\, c\, \{Q\}$, where P and Q are assertions and c is a command. To define validity of such a triple (in the absence of usual operational semantics), we formalize the notion of a visible state, taken to be a function from Loc to Val.

Definition 6. A *snapshot* of an execution $G = \langle A, L, E\rangle$ at an edge $\langle a, b\rangle \in po$ with respect to a set $B \subseteq A$ of nodes, denoted by $\mathcal{S}(G, \langle a, b\rangle, B)$, is the execution $\langle A' \uplus \{b\}, L|_{A'} \cup \{b \mapsto \langle \mathtt{S}\rangle\}, E|_{A'} \cup \{\langle a, b\rangle\}\rangle$, where:

- $A' = \{a' \in A \setminus \{b\} \mid \exists c \in B \cup \{a\}.\ \langle a', c\rangle \in E^*_{all}\}$ and
- $E|_{A'} = E \cap ((A' \times A') \cup (A' \times A' \times \mathsf{Loc}))$.

Definition 7. Let G be an execution, and let $\langle a, b\rangle \in po$.

- A function $D : \mathsf{Loc} \to \mathbb{N}$ is called a $\langle G, \langle a, b\rangle\rangle$-*reader* of a state $\sigma : \mathsf{Loc} \to \mathsf{Val}$ if $D(x) \in \mathtt{W}_x \cup \mathtt{U}_x$ and $val_w(D(x)) = \sigma(x)$ for every $x \in \mathsf{Loc}$, and the execution $\mathcal{S}(G, \langle a, b\rangle, D[\mathsf{Loc}]) \cup \{\langle D(x), b, x\rangle \mid x \in \mathsf{Loc}\}$ is coherent.
- A state σ is called *visible* at $\langle a, b\rangle$ in G if there is a $\langle G, \langle a, b\rangle\rangle$-reader of σ.
- An assertion P *holds* at $\langle a, b\rangle$ in G if $\sigma \models P$ for every state σ visible at $\langle a, b\rangle$ in G.

In essence, the snapshot restricts the execution to the edge $\langle a, b\rangle$, all nodes in B, and all prior nodes and edges, and replaces the label of b by a skip. For a state to be visible at $\langle a, b\rangle$, additional reads-from edges should be added. For an example, see Fig. 4.

Definition 8. For a state σ, let $\mathcal{WG}(\sigma)$ be $\mathcal{WG}(\nu_1, \sigma(\nu_1)) \parallel \ldots \parallel \mathcal{WG}(\nu_M, \sigma(\nu_M))$, the set of all σ-initializations. Given an assertion P, $\mathcal{WG}(P) = \bigcup\{\mathcal{WG}(\sigma) \mid \sigma \models P\}$.

An execution G is called *initialized* if $G = (G_1; G_2) \cup E$ for some $G_1 \in \mathcal{WG}(\top)$, plain execution G_2, and set $E \subseteq A_1 \times A_2 \times \mathsf{Loc}$ of edges. It can be shown that if G is coherent and initialized, then at least one state is visible at every program order edge.

Definition 9. A Hoare triple $\{P\}\, c\, \{Q\}$ is *valid* if Q holds at the terminal edge of $G \cup E'$ in $G \cup E'$ for every execution $G = \langle A, L, E\rangle$ in $\mathcal{WG}(P); [\![c]\!]; \mathcal{SG}$ and set $E' \subseteq A \times A \times \mathsf{Loc}$, such that $G \cup E'$ is a complete and coherent execution.

OG-style reasoning is often judged as non-compositional because it refers to non-interference of proof outlines that cannot be checked based solely on the two input Hoare triples. A straightforward remedy is to use a *rely/guarantee-style presentation* of OG, that permits compositional reasoning. In this case, the rely component, denoted by $\mathcal{R}$, consists of a set of assertions that are assumed to be stable under assignments performed by other threads. In turn, the guarantee component, denoted by $\mathcal{G}$, is a set of *guarded assignments*, that is assignments together with their immediate preconditions. Roughly speaking, a validity of an OG judgment $\mathcal{R}; \mathcal{G} \Vdash \{P\}\, c\, \{Q\}$ amounts to: "every terminating run of c starting from a state in P ends in a state in Q, and performs only assignments in $\mathcal{G}$, where each of which is performed while satisfying its guard; and moreover, the above holds in parallel to any run of a program c', provided that the assertions in $\mathcal{R}$ are stable under each of the assignments performed by c'."

Now, as demonstrated in the introduction, reasoning under RA requires a richer rely condition, as stability of an assertion in thread I under a guarded assignment of the form $\{P\}x := e(y)$ in thread II should be checked for all values readable for y in some non-later point of thread I. Similarly, stability under $\{P\}x \overset{y,z}{:=} e(y, z)$ should cover all values readable for y and z in two non-later points. Hence, we take $\mathcal{R}$ to consist of pairs of assertions, where the first component of each pair describes the current state and the second summarizes all non-later states. This leads us to the following definitions.

Definition 10. An *OG judgment* $\mathcal{R}; \mathcal{G} \Vdash \{P\}\, c\, \{Q\}$ extends a Hoare triple with two extra components:

- A finite set $\mathcal{R}$ of pairs of the form $R\upharpoonleft C$, where R and C are assertions. We write $\mathcal{R}^R$ for $\bigvee\{R \mid R\upharpoonleft _ \in \mathcal{R}\}$ and $\mathcal{R}^C$ for $\bigwedge\{C \mid _\upharpoonleft C \in \mathcal{R}\}$. We also write $\mathcal{R} \leq \mathcal{R}'$ for such sets if for every $R\upharpoonleft C \in \mathcal{R}$ there exists C' such that $R\upharpoonleft C' \in \mathcal{R}'$ and $C \vdash C'$.
- A finite set $\mathcal{G}$ of *guarded assignments*, i.e., pairs of the form $\{R\}c$, where R is an assertion and c is an assignment command. We write $\mathcal{G} \leq \mathcal{G}'$ for such sets if for every $\{R\}c \in \mathcal{G}$ there exists R' such that $\{R'\}c \in \mathcal{G}'$ and $R \vdash R'$.

Definition 11. A pair $R\upharpoonleft C$ is *stable* under $\{P\}c$ if one of the following holds:

- c has the form $x := v$ and $R \wedge P \vdash R[v/x]$;
- c has the form $x := e(y)$ and $R \wedge P \vdash R[[\![e]\!](v_y)/x]$ for every $v_y \in \mathsf{Val}$ such that $C \wedge P \nvdash y \neq v_y$ (i.e., for every $v_y \in \mathsf{Val}$ such that $C \wedge P \wedge y = v_y$ is satisfiable);
- c has the form $x \overset{y,z}{:=} e(y, z)$ and $R \wedge P \vdash R[[\![e]\!](v_y, v_z)/x]$ for every $v_y, v_z \in \mathsf{Val}$, such that $C \wedge P \nvdash y \neq v_y$ and $C \wedge P \nvdash z \neq v_z$; or
- c has the form $x \overset{\mathsf{at}}{:=} e(x)$ and $R \wedge P \vdash R[e/x]$.

The proof system for deriving OGRA's judgments is given in Fig. 5. The rules are essentially those of Owicki and Gries [10] with minor adjustments due to our rely/guarantee style presentation and the more complex form of the $\mathcal{R}$ component. (To assist the reader, the supplementary material includes a similar presentation of usual OG.) Typically, we require the preconditions and postconditions to be included in $\mathcal{R}$, and make sure their second components keep track of (at least) all non-later assertions: for example, all the assignment rules require $\{P\upharpoonleft P, Q\upharpoonleft(P \vee Q)\} \leq \mathcal{R}$.

The rule for parallel composition (PAR) allows composing non-interfering judgments. Its precondition is the conjunction of the preconditions of the threads, while its

$$(\text{CONSEQ})\ \frac{\begin{array}{c}\mathcal{R};\mathcal{G} \Vdash \{P\}\, c\, \{Q\} \\ P' \vdash P \quad Q \vdash Q' \quad \mathcal{R} \leq \mathcal{R}' \quad \mathcal{G} \leq \mathcal{G}'\end{array}}{\mathcal{R}';\mathcal{G}' \Vdash \{P'\}\, c\, \{Q'\}}
\qquad
(\text{SEQ})\ \frac{\begin{array}{c}\mathcal{R}_1;\mathcal{G}_1 \Vdash \{P\}\, c_1\, \{R\} \\ \mathcal{R}_2;\mathcal{G}_2 \Vdash \{R\}\, c_2\, \{Q\} \quad \mathcal{R}_1^R \vdash \mathcal{R}_2^C\end{array}}{\mathcal{R}_1 \cup \mathcal{R}_2;\mathcal{G}_1 \cup \mathcal{G}_2 \Vdash \{P\}\, c_1; c_2\, \{Q\}}$$

$$(\text{SKIP})\ \frac{\{P \nearrow P\} \leq \mathcal{R}}{\mathcal{R};\emptyset \Vdash \{P\}\, \texttt{skip}\, \{P\}}
\qquad
(\text{PAR})\ \frac{\begin{array}{c}\mathcal{R}_1;\mathcal{G}_1 \Vdash \{P_1\}\, c_1\, \{Q_1\} \quad \mathcal{R}_2;\mathcal{G}_2 \Vdash \{P_2\}\, c_2\, \{Q_2\} \\ Q_1 \wedge Q_2 \vdash Q \quad \mathcal{R}_1;\mathcal{G}_1 \text{ and } \mathcal{R}_2;\mathcal{G}_2 \text{ are non-interfering}\end{array}}{\mathcal{R}_1 \cup \mathcal{R}_2 \cup \{Q \nearrow (\mathcal{R}_1^R \vee \mathcal{R}_2^R \vee Q)\};\mathcal{G}_1 \cup \mathcal{G}_2 \Vdash \{P_1 \wedge P_2\}\, c_1 \parallel c_2\, \{Q\}}$$

$$(\text{ASSN}_0)\ \frac{P \vdash Q[v/x] \quad \{P \nearrow P, Q \nearrow (P \vee Q)\} \leq \mathcal{R}}{\mathcal{R};\{\{P\}x := v\} \Vdash \{P\}\, x := v\, \{Q\}}
\qquad
(\text{ASSN}_1)\ \frac{P \vdash Q[e(y)/x] \quad \{P \nearrow P, Q \nearrow (P \vee Q)\} \leq \mathcal{R}}{\mathcal{R};\{\{P\}x := e(y)\} \Vdash \{P\}\, x := e(y)\, \{Q\}}$$

$$(\text{ASSN}_2)\ \frac{\begin{array}{c}P \vdash Q[e(y,z)/x] \quad \{P \nearrow P, Q \nearrow (P \vee Q)\} \leq \mathcal{R} \\ \{(P \wedge (y = v)) \nearrow P \mid v \in \mathsf{Val}\} \leq \mathcal{R}\end{array}}{\mathcal{R};\{\{P\}x \overset{y,z}{:=} e(y,z)\} \Vdash \{P\}\, x \overset{y,z}{:=} e(y,z)\, \{Q\}}
\qquad
(\text{ASSN}_{at})\ \frac{P \vdash Q[e(x)/x] \quad \{P \nearrow P, Q \nearrow (P \vee Q)\} \leq \mathcal{R}}{\mathcal{R};\{\{P\}x \overset{\mathsf{at}}{:=} e(x)\} \Vdash \{P\}\, x \overset{\mathsf{at}}{:=} e(x)\, \{Q\}}$$

$$(\text{ITE})\ \frac{\begin{array}{c}\{P \nearrow P\} \leq \mathcal{R} \quad P \vdash \mathcal{R}^C \\ \mathcal{R};\mathcal{G} \Vdash \{P \wedge (e(x) \neq 0)\}\, c_1\, \{Q\} \\ \mathcal{R};\mathcal{G} \Vdash \{P \wedge (e(x) = 0)\}\, c_2\, \{Q\}\end{array}}{\mathcal{R};\mathcal{G} \Vdash \{P\}\, \texttt{if}\ e(x)\ \texttt{then}\ c_1\ \texttt{else}\ c_2\, \{Q\}}
\qquad
(\text{WHILE})\ \frac{\begin{array}{c}P \nearrow _ \in \mathcal{R} \quad \mathcal{R}^R \vdash \mathcal{R}^C \quad P \wedge (e(x) = 0) \vdash Q \\ \mathcal{R};\mathcal{G} \Vdash \{P \wedge (e(x) \neq 0)\}\, c\, \{P\}\end{array}}{\mathcal{R} \cup \{Q \nearrow (\mathcal{R}^R \vee Q)\};\mathcal{G} \Vdash \{P\}\, \texttt{while}\ e(x)\ \texttt{do}\ c\, \{Q\}}$$

Fig. 5. Owicki-Gries proof system for release-acquire.

postcondition, Q, is any stable assertion implied by the conjunction of the thread post-conditions. (The asymmetry is because of the second components of the $\mathcal{R}$ entries: the states prior to the end of the parallel compositions are the union of those of each thread, and hence the stability of Q does not necessarily follow from that of Q_1 and Q_2.) Non-interference is checked for every rely condition of one thread and guarded assignment in the guarantee component of the other:

Definition 12. $\mathcal{R}_1;\mathcal{G}_1$ and $\mathcal{R}_2;\mathcal{G}_2$ are *non-interfering* if every $R \nearrow C \in \mathcal{R}_i$ is stable under every $\{P\}c \in \mathcal{G}_j$ for $i \neq j$.

The consequence rule (CONSEQ) allows strengthening the precondition ($P' \vdash P$), weakening the postcondition ($Q \vdash Q'$), increasing the set of assertions required to be stable ($\mathcal{R} \leq \mathcal{R}'$), and increasing the set of allowed guarded assignments ($\mathcal{G} \leq \mathcal{G}'$).

The sequential composition rule (SEQ) collects the assertions and allowed assignments of both commands, and checks that $\mathcal{R}_1^R \vdash \mathcal{R}_2^C$. This ensures that stability of c_2's assertions would take into account all the states of c_1, that now become previous states.

The next interesting rule is ASSN_2 concerning assignments with expressions reading two variables. The rule requires that the value of the first variable being read (y) is stable assuming P also holds. This check is needed because of the way we interpret assertions as snapshot reads differs from the way that programs read the variables (one at a time): the stability check ensures that the difference is not observable. Note that the stability of y is trivial in case that there are no assignments to it in other threads.

Finally, the rules for conditionals and while-loops are standard: as with the SEQ rule, we require that the second component of $\mathcal{R}$ has taken into account all earlier states, and include the initial precondition in the set of stable assertions.

We can now state our main theorem, namely the soundness of OGRA.

$\{x = 0\}$

Thread I	Thread II
$\{\top\}$	$\{x \neq 0 \to m = 42\}$
$m := 42;$	`while` $x = 0$ `do skip;`
$\{m = 42\}$	$\{m = 42\}$
$x := 1$	$a := m$
$\{\top\}$	$\{a = 42\}$

$\{a = 42\}$

Fig. 6. Proof outline for a simple message passing idiom.

$\{f = 0\}$

Thread I	Thread II
$\{f \in \{0,2\}\}$	$\{f \in \{0,1\}\}$
$x := 1;$	$y := 1;$
$\{f \in \{0,2\} \wedge x = 1\}$	$\{f \in \{0,1\} \wedge y = 1\}$
$f \overset{\text{at}}{:=} 10f + 1;$	$f \overset{\text{at}}{:=} 10f + 2;$
$\{f \in \{1,12,21\} \wedge x = 1\}$	$\{f \in \{2,12,21\} \wedge y = 1\}$
$a := y$	$b := x$
$\left\{\begin{array}{l} f \in \{1,12,21\} \wedge x = 1 \wedge \\ (f = 21 \to a = y) \end{array}\right\}$	$\left\{\begin{array}{l} f \in \{2,12,21\} \wedge y = 1 \wedge \\ (f = 12 \to b = x) \end{array}\right\}$

$\{a = 1 \vee b = 1\}$

Fig. 7. Proof outline for "store buffering" with fences.

$\{x = a = c = 0\}$

Thread I	Thread II	Thread III	Thread IV
$\left\{\begin{array}{l} (x \neq 1 \wedge a \neq 1) \\ \upharpoonright x \neq 1 \end{array}\right\}$	$\left\{\begin{array}{l} (x \neq 2 \wedge c \neq 2) \\ \upharpoonright x \neq 2 \end{array}\right\}$	$\{\top\}$	$\{\top\}$
$x := 1$	$x := 2$	$a := x;$	$c := x;$
$\{\top\}$	$\{\top\}$	$\{\top\}$	$\{\top\}$
		$b := x$	$d := x$
		$\{a = 1 \wedge b = 2 \to x = 2\}$	$\{c = 2 \wedge d = 1 \to x = 1\}$

$\{a = 1 \wedge b = 2 \wedge c = 2 \to d \neq 1\}$

Fig. 8. Proof outline for read-read coherence test (example CoRR2 in [8]).

Theorem 1. *If $\mathcal{R}; \mathcal{G} \Vdash \{P\}\, c\, \{Q\}$ is derivable, then $\{P\}\, c\, \{Q\}$ is valid.*

Before proving this theorem, we provide a few example derivations. The derivations are presented in a proof outline fashion. For each thread, the set $\mathcal{R}$ consists of all the assertions in its proof outline, with the second component being $\top$ (all values are possible) unless mentioned otherwise. The set $\mathcal{G}$ consists of all the assignments in the proof outline guarded by their immediate preconditions.

Our first example, shown in Fig. 6, is a simple message passing idiom. Thread I initializes a message m to 42 and then raises a flag x; thread II waits for x to have a non-zero value and then reads m, which should have value 42. To prove this, thread II assumes the invariant $x \neq 0 \to m = 42$ that holds initially and is stable.

Our next example, shown in Fig. 7, is a variant of the "store buffering" program (see Fig. 1) that uses fences to restore sequential consistency. Fence instructions are implemented as RMWs to a distinguished location f. The RA semantics enforces the corresponding update nodes to be linearly ordered by E^*_{all}, so this implementation imposes a synchronization between every pair of fences. These fences are stronger than C11's SC fences, as they restore full SC when placed between every pair of consecutive instructions. While any atomic assignment to f will have this effect, we choose commands that record the exact order in which the fences are linearized. By referring to this order in the proof, we can easily show that the outcome $a = b = 0$ is not possible.

Our third example, shown in Fig. 8, is a coherence test, demonstrating that threads cannot observe writes to the same location happen in different orders. The program consists of two independent writes to x and two readers: the goal is to prove that the first reader cannot read the one write and then the other, while the second reads them in the reverse order. The key to showing this are the assertions at the end of the reader threads saying that the value of x cannot change after both assignments have been observed. For these assertions to be stable, the writers correspondingly assert that the assignments to x happen before the corresponding reader observes x to have that value. Formally, the precondition of the $x := 1$ assignment is $(x \neq 1 \wedge a \neq 1) \upharpoonright x \neq 1$. This is stable under the $a := x$ assignment because 1 is not a readable value for x (we have: $x \neq 1 \not\models x \neq v$ iff $v \neq 1$).

3.1 Soundness Proof

We present the main steps in the proof of Theorem 1. Annotations play a crucial role. An *annotation* is a function that assigns an assertion to every pair in $\mathbb{N} \times \mathbb{N}$. An annotation Θ is *valid* for an execution G if $\Theta(\langle a, b\rangle)$ holds at $\langle a, b\rangle$ in G for every $\langle a, b\rangle \in po$.

The proof consists of two parts. First, we show that derivability of a judgment $\mathcal{R}; \mathcal{G} \Vdash \{P\}\, c\, \{Q\}$ allows us to construct annotations of executions of c, that are *locally valid* and *stable*, as defined below. Then, we prove that such annotations, for complete and coherent executions, must also be valid. Theorem 1 is obtained as a corollary.

Definition 13. An annotation Θ is *locally valid* for an execution G if the following hold for every $\langle a, b\rangle \in po$, where $P = \bigwedge_{\langle a', a\rangle \in po} \Theta(\langle a', a\rangle)$ and $Q = \Theta(\langle a, b\rangle)$:

- If $L(a) = \langle \mathtt{S}\rangle$ and a is not initial then $P \vdash Q$.
- If $L(a) = \langle \mathtt{R}, x, v\rangle$ then $P \wedge (x = v) \vdash Q$.
- If $L(a) = \langle \mathtt{W}, x, v\rangle$ then either $P \vdash Q[v/x]$, or there is a unique node a' such that $\langle a', a\rangle \in po$, and we have $a' \in \mathtt{R}$ and $P \wedge (loc(a') = val_r(a')) \vdash Q[v/x]$.
- If $L(a) = \langle \mathtt{U}, x, v_r, v_w\rangle$ then $P \wedge (x = v_r) \vdash Q[v_w/x]$.

Definition 14. Let G be an execution. An edge $\langle b_1, b_2\rangle \in po$ is called *G-before* an edge $\langle a_1, a_2\rangle \in po$ if either $\langle b_1, b_2\rangle = \langle a_1, a_2\rangle$ or $\langle b_2, a_1\rangle \in po^*$.

Definition 15. Let G be an execution. A node $c \in G$ *interferes* with $\langle a, b\rangle \in po$ in G for an annotation Θ if the following hold:

- $\langle c, a\rangle \notin po^*$ and $\langle b, c\rangle \notin po^*$ (c is *parallel* to $\langle a, b\rangle$ in G).
- For all $c' \in \mathtt{R}$ with $\langle c', c\rangle \in po$ and $\langle c', a\rangle \notin po^*$, we have $\Theta(\langle a', b'\rangle) \wedge \Theta(\langle c', c\rangle) \nvdash loc(c') \neq val_r(c')$ for some $\langle a', b'\rangle \in po$ such that $\langle a', b'\rangle$ is G-before $\langle a, b\rangle$ and $\langle b', c'\rangle \notin po^*$.

Definition 16. An annotation Θ is *stable* for an execution G if the following hold for every $\langle a, b\rangle \in po$ and node $c \in \mathtt{W} \cup \mathtt{U}$ that interferes with $\langle a, b\rangle$ in G for Θ, where $R = \Theta(\langle a, b\rangle)$ and $P = \bigwedge_{\langle c', c\rangle \in po} \Theta(\langle c', c\rangle)$:

- If $L(c) = \langle \mathtt{W}, x, v\rangle$ then $P \wedge R \vdash R[v/x]$.
- If $L(c) = \langle \mathtt{U}, x, v_r, v_w\rangle$ then $P \wedge (x = v_r) \wedge R \vdash R[v_w/x]$.

Definition 17. A Hoare triple $\{P\}\, c\, \{Q\}$ is *safe* if for every $G \in \mathcal{SG}; [\![c]\!]; \mathcal{SG}$, there is an annotation Θ that is locally valid and stable for G, and assigns some assertion P', such that $P \vdash P'$, to the initial edge of G, and some assertion Q', such that $Q' \vdash Q$, to its terminal edge.

Theorem 2. *If $\mathcal{R}; \mathcal{G} \Vdash \{P\}\, c\, \{Q\}$ is derivable for some $\mathcal{R}, \mathcal{G}$, then $\{P\}\, c\, \{Q\}$ is safe.*

Proof (Outline). Call a judgment $\mathcal{R}; \mathcal{G} \Vdash \{P\}\, c\, \{Q\}$ *good* if for every execution $G \in \mathcal{SG}; [\![c]\!]; \mathcal{SG}$, there exists an annotation Θ that satisfies the conditions given in Definition 17, as well as the following ones:

- $\mathcal{R}$ covers Θ for G, i.e., for every $\langle a_1, a_2\rangle \in po$, there exist $P_1 \nearrow C_1, \dots, P_n \nearrow C_n \in \mathcal{R}$ such that $\bigwedge P_i \dashv\vdash \Theta(\langle a_1, a_2\rangle)$ and $\Theta(\langle b_1, b_2\rangle) \vdash \bigwedge C_i$ for every $\langle b_1, b_2\rangle \in po$ that is G-before $\langle a_1, a_2\rangle$ (in particular, for $\langle b_1, b_2\rangle = \langle a_1, a_2\rangle$).

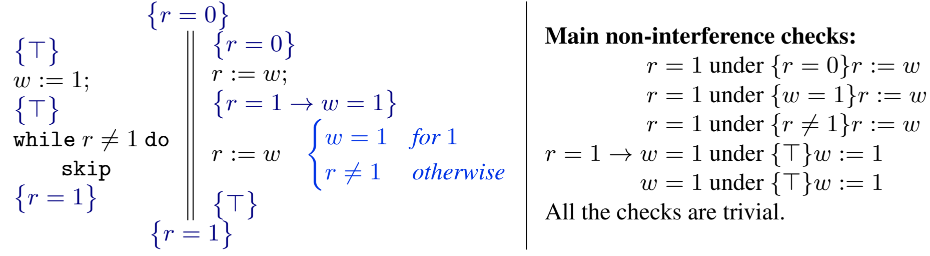

Fig. 9. Simplified RCU example illustrating the use of the stronger assignment rule.

- $\mathcal{G}$ covers Θ for G, i.e., for every $a_2 \in \mathtt{W} \cup \mathtt{U}$, there exist an edge $\langle a_1, a_2 \rangle \in po$ and an assertion P', such that $\Theta(\langle a_1, a_2 \rangle) \vdash P'$, and one of the following holds:
  - $L(a_2) = \langle \mathtt{W}, x, v \rangle$ and $\{P'\} x := v \in \mathcal{G}$.
  - $L(a_2) = \langle \mathtt{W}, x, v \rangle$, $L(a_1) = \langle \mathtt{R}, y, v_y \rangle$, and $\{P'\} x := e(y) \in \mathcal{G}$ for some expression $e(y)$ such that $[\![e]\!](v_y) = v$.
  - $L(a_2) = \langle \mathtt{W}, x, v \rangle$, $L(a_1) = \langle \mathtt{R}, z, v_z \rangle$, $\Theta(\langle a_1, a_2 \rangle) \vdash y = v_y$ for some $v_y \in \mathsf{Val}$, and $\{P'\} x \overset{y,z}{:=} e(y, z) \in \mathcal{G}$ for some expression $e(y, z)$ such that $[\![e]\!](v_y, v_z) = v$.
  - $L(a_2) = \langle \mathtt{U}, x, v_r, v_w \rangle$ and $\{P'\} x \overset{\mathsf{at}}{:=} e(x) \in \mathcal{G}$ for some expression $e(x)$ such that $[\![e]\!](v_r) = v_w$.

Next, by induction on the derivation, one shows that every derivable judgment $\mathcal{R}; \mathcal{G} \Vdash \{P\}\, c\, \{Q\}$ is good, and so $\{P\}\, c\, \{Q\}$ is safe. The non-interference condition is needed for showing that two annotations of executions G_1 and G_2 can be joined to a stable annotation of the parallel composition of G_1 and G_2. □

It remains to establish the link from safety of a Hoare triple to its validity.

Theorem 3. *Let G be a complete coherent initialized execution. If an annotation Θ is locally valid and stable for G, then it is valid for G.*

The proof (given in the full version of this paper) requires analyzing the relations between states that are visible on consecutive edges and parallel edges in the RA memory model. An alternative equivalent formulation of coherence, based on a new "write-before" relation, is particularly useful for this task.

3.2 A Stronger Assignment Rule

Consider the program shown in Fig. 9, which contains an idiom found in the RCU implementation (verified in the supplementary material). Thread II reads w and writes its value to r twice, while thread I sets w to 1 and then waits for r to become 1. The challenge is to show that after thread I reads $r = 1$, the value of r does not change; i.e. that $r = 1$ is stable under the $r := w$ assignments. For the first $r := w$ assignment, this is easy because its precondition is inconsistent with $r = 1$. For the second assignment, however, there is not much we can do. Stability requires us to consider *any value* for w readable at some point by thread I. Our idea is to do a case split on the value that w reads. If w reads the value 1, then it writes $r := 1$, and so $r = 1$ is unaffected. If w reads a different value, then from the assignment's precondition, we can derive $r \neq 1$, which contradicts the $r = 1$ assertion.

$$\{\top\}$$
$$\begin{array}{l} x := 2; \\ y := 1 \end{array} \;\Big\|\; \begin{array}{l} y := 2; \\ x := 1 \end{array}$$
$$\{x \neq 2 \lor y \neq 2\}$$

Fig. 10. Auxiliary variables are necessary under SC.

$$\{x = \langle 0,0 \rangle\}$$
$$\begin{array}{l} \{x \in \{\langle 0,0\rangle, \langle 1,2\rangle\}\} \\ x \overset{\mathsf{at}}{:=} \langle x_{\mathrm{fst}} + 1, x_{\mathrm{snd}} + 1\rangle \\ \{x \in \{\langle 1,1\rangle, \langle 2,3\rangle\}\} \end{array} \;\Big\|\; \begin{array}{l} \{x \in \{\langle 0,0\rangle, \langle 1,1\rangle\}\} \\ x \overset{\mathsf{at}}{:=} \langle x_{\mathrm{fst}} + 1, x_{\mathrm{snd}} + 2\rangle \\ \{x \in \{\langle 1,2\rangle, \langle 2,3\rangle\}\} \end{array}$$
$$\{x = \langle 2,3 \rangle\}$$

Fig. 11. Verification of the parallel increment example.

To support such case splits, we provide the following stronger assignment rule. For simplicity, we consider only assignments of the form $x := e(y)$.

$$(\text{ASSN}'_1)\; \frac{\begin{array}{c} P \vdash Q[e(y)/x] \qquad \{P \upharpoonleft P, Q \upharpoonleft (P \lor Q)\} \leq \mathcal{R} \\ \text{For every } v \in \mathsf{Val}: \qquad P \land (y = v) \vdash P_v \qquad \{P_v \upharpoonleft P\} \leq \mathcal{R} \end{array}}{\mathcal{R}; \{\{P_v\} x := e(y) \mid v \in \mathsf{Val}\} \Vdash \{P\}\, x := e(y)\, \{Q\}}$$

The previous assignment rule is an instance of this rule by taking $P_v = P$ for all v.

4 Discussion and Further Research

While OGRA's non-interference condition appears to be restrictive, we note that it is unsound for weaker memory models, such as C11's relaxed accesses because it can prove, e.g., message passing, see Fig. 6. We also observe that OGRA's non-interference check coincides with the standard OG one for assignments of values ($x := v$) and atomic assignments ($x \overset{\mathsf{at}}{:=} e(x)$). Moreover, the non-interference check is irrelevant for assignments to variables that do not occur in the proof outlines of other threads. Therefore, standard OG (without auxiliary variables) is sound under RA provided that all $x := e(y)$ and $x \overset{y,z}{:=} e(y,z)$ assignments write to variables that do not appear in the proof outlines of other threads. Figures 6 and 7 provide two such cases in point. In addition, this entails, for instance, that the program in Fig. 10 cannot be verified in standard OG without auxiliary variables, as $x = 2 \land y = 2$ is a possible outcome for this program under RA.

OG's auxiliary variables, in general, are unsound under weak memory because they can be used to record the exact thread interleavings and establish completeness under SC [12]. A simple form of auxiliary state, which we call *ghost values*, however, *is* sound. The idea is as follows: given a program c, one may choose a domain G of "ghost" values, together with a function $\alpha : \mathsf{G} \to \mathsf{Val}$, and obtain a program c' by substituting each expression $e(x_1, \ldots, x_n)$ in c by an expression $e'(x_1, \ldots, x_n)$ such that $\alpha([\![e']\!](g_1, \ldots, g_n)) = [\![e]\!](\alpha(g_1), \ldots, \alpha(g_n))$ for all $g_1, \ldots, g_n \in \mathsf{G}$. The validity of $\{P'\}\, c'\, \{Q'\}$ entails the validity of $\{P\}\, c\, \{Q\}$, provided that the following hold:

- If a state satisfies P then some corresponding ghost state satisfies P';
- If a state does not satisfy Q then any corresponding ghost state does not satisfy Q';

where a ghost state $\sigma' : \mathsf{Loc} \to \mathsf{G}$ *corresponds* to a state $\sigma : \mathsf{Loc} \to \mathsf{Val}$ iff $\alpha(\sigma'(x)) = \sigma(x)$ for every $x \in \mathsf{Loc}$. This solution suffices, for instance, to reason about the parallel increment example, as shown in Fig. 11. There we took $\mathsf{G} = \mathsf{Val} \times \mathbb{N}$, with α being the first projection mapping. The second component tracks which of the assignments has already happened (0: none, 1: the first thread, 2: the second thread, otherwise: both). As a result, we obtain the validity of $\{x = 0\}\, x \overset{\mathsf{at}}{:=} x + 1 \parallel x \overset{\mathsf{at}}{:=} x + 1\, \{x = 2\}$.

Analyzing soundness of other restricted forms of auxiliary variables is left for future work. Such extensions seem to be a prerequisite for obtaining a program logic that is both sound and complete under RA. Automation of proof search is another future goal. Our initial experiments show that, at least for the examples in this paper, HSF [5] is successful in automatically finding proofs in OGRA.

Acknowledgments. We would like to thank the ICALP'15 reviewers for their feedback. This work was supported by EC FET project ADVENT (308830).

References

1. Batty, M., Owens, S., Sarkar, S., Sewell, P., Weber, T.: Mathematizing C++ concurrency. In: POPL 2011. pp. 55–66. ACM (2011)
2. Cohen, E.: Coherent causal memory. CoRR abs/1404.2187 (2014)
3. Desnoyers, M., McKenney, P.E., Stern, A.S., Dagenais, M.R., Walpole, J.: User-level implementations of read-copy update. IEEE Trans. Parallel Distrib. Syst. 23(2), 375–382 (2012)
4. Ferreira, R., Feng, X., Shao, Z.: Parameterized memory models and concurrent separation logic. In: Gordon, A.D. (ed.) ESOP 2010. LNCS, vol. 6012, pp. 267–286. Springer, Heidelberg (2010)
5. Grebenshchikov, S., Lopes, N.P., Popeea, C., Rybalchenko, A.: Synthesizing software verifiers from proof rules. In: PLDI 2012. pp. 405–416. ACM (2012)
6. ISO/IEC 14882:2011: Programming language C++ (2011)
7. Lamport, L.: How to make a multiprocessor computer that correctly executes multiprocess programs. IEEE Trans. Computers 28(9), 690–691 (1979)
8. Maranget, L., Sarkar, S., Sewell, P.: A tutorial introduction to the ARM and POWER relaxed memory models. http://www.cl.cam.ac.uk/~pes20/ppc-supplemental/test7.pdf (2012)
9. Owens, S., Sarkar, S., Sewell, P.: A better x86 memory model: x86-TSO. In: Berghofer, S., Nipkow, T., Urban, C., Wenzel, M. (eds.) TPHOLs 2009. LNCS, vol. 5674, pp. 391–407. Springer, Heidelberg (2009)
10. Owicki, S., Gries, D.: An axiomatic proof technique for parallel programs I. Acta Informatica 6(4), 319–340 (1976)
11. Owicki, S., Gries, D.: Verifying properties of parallel programs: An axiomatic approach. Commun. ACM 19(5), 279–285 (May 1976)
12. Owicki, S.S.: Axiomatic Proof Techniques for Parallel Programs. Ph.D. thesis, Cornell University, Ithaca, NY, USA (1975)
13. Ridge, T.: A rely-guarantee proof system for x86-TSO. In: Leavens, G.T., O'Hearn, P.W., Rajamani, S.K. (eds.) VSTTE 2010. LNCS, vol. 6217, pp. 55–70. Springer, Heidelberg (2010)
14. Sieczkowski, F., Svendsen, K., Birkedal, L., Pichon-Pharabod, J.: A separation logic for fictional sequential consistency. In: Vitek, J. (ed.) ESOP 2015. LNCS, vol. 9032, pp. 736–761. Springer, Heidelberg (2015)
15. Tassarotti, J., Dreyer, D., Vafeiadis, V.: Verifying read-copy-update in a logic for weak memory. In: PLDI 2015. ACM (2015)
16. Turon, A., Vafeiadis, V., Dreyer, D.: GPS: Navigating weak memory with ghosts, protocols, and separation. In: OOPSLA 2014. pp. 691–707. ACM (2014)
17. Vafeiadis, V., Balabonski, T., Chakraborty, S., Morisset, R., Zappa Nardelli, F.: Common compiler optimisations are invalid in the C11 memory model and what we can do about it. In: POPL 2015. pp. 209–220. ACM (2015)
18. Vafeiadis, V., Narayan, C.: Relaxed separation logic: A program logic for C11 concurrency. In: OOPSLA 2013. pp. 867–884. ACM (2013)