

1 Owicki-Gries Reasoning for C11 RAR

2 **Sadegh Dalvandi**

3 University of Surrey, UK
4 m.dalvandi@surrey.ac.uk

5 **Simon Doherty**

6 University of Sheffield, UK
7 s.doherty@sheffield.ac.uk

8 **Brijesh Dongol**

9 University of Surrey, UK
10 b.dongol@surrey.ac.uk

11 **Heike Wehrheim**

12 University of Paderborn, Germany
13 wehrheim@upb.de

14 — Abstract —

15 Owicki-Gries reasoning for concurrent programs uses Hoare logic together with an *interference*
16 *freedom* rule for concurrency. In this paper, we develop a new proof calculus for the C11 RAR
17 memory model (a fragment of C11 with both relaxed and release-acquire accesses) that allows all
18 Owicki-Gries proof rules for compound statements, including non-interference, to remain unchanged.
19 Our proof method features novel assertions specifying *thread-specific views* on the state of programs.
20 This is combined with a set of Hoare logic rules that describe how these assertions are affected by
21 atomic program steps. We demonstrate the utility of our proof calculus by verifying a number of
22 standard C11 litmus tests and Peterson’s algorithm adapted for C11. Our proof calculus and its
23 application to program verification have been fully mechanised in the theorem prover Isabelle.

24 **2012 ACM Subject Classification** To be done

25 **Keywords and phrases** C11, Verification, Hoare logic, Owicki-Gries, Isabelle

26 **Digital Object Identifier** 10.4230/LIPIcs.AAA.2020.0

27 **1 Introduction**

28 In 1976, Susan Owicki and David Gries proposed an extension of Hoare’s axiomatic reasoning
29 technique [15] to concurrent programs [27]. Their proof calculus allows one to reason about
30 concurrent programs with shared variables via a number of proof rules, including the rules
31 for sequential programs as introduced by Hoare plus an additional proof rule for concurrent
32 composition. This composition rule basically allows for the conjunction of pre- and post-
33 conditions of the process’ individual proofs, given that their proof outlines are *interference*
34 *free*. Interference freedom requires that an assertion in the proof of one process cannot
35 be invalidated by a statement in another process, when executed under the statement’s
36 precondition.

37 Today, concurrent programs are run on multi-core processors. Multi-core processors
38 come with *weak memory models* specifying the execution behaviour of concurrent programs.
39 Reasoning consequently needs to be adapted to the memory model under consideration.
40 Owicki-Gries reasoning is, however, fixed to the memory model of *sequential consistency*
41 (SC) [23], and is unsound for weak memory models. Recent research has thus worked towards
42 new sound proof calculi for concurrent programs. Most often, such approaches involve
43 concurrent separation logics (e.g., GPS and RSL [34, 16]). These techniques constitute a
44 radical departure from the (relatively) small and easy proof calculus of Owicki and Gries,
45 further extending already complex logics. A proposal for a (rely-guarantee variant of) the



© Anonymous;
licensed under Creative Commons License CC-BY

AAA.

Editors: AAA; Article No. 0; pp. 0:1–0:38



Leibniz International Proceedings in Informatics

LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

46 Owicki-Gries proof system has been made by Lahav and Vafeiadis [22], however, requiring a
47 strengthened non-interference check.

48 In this paper, we develop a proof method based on the Owicki-Gries proof calculus,
49 keeping all of the original proof rules including the non-interference check unchanged. Our
50 technique introduces a set of basic axioms to cope with memory accesses (reads, writes,
51 read-modify-writes) and simple assertions that describe the current configuration of the weak
52 memory state. Our proof calculus targets the weak memory model of the C11 programming
53 language [9]. Here, we deal with the release-acquire-relaxed (RAR) fragment of C11 (thereby
54 going further than prior work on Owicki-Gries reasoning for C11 [22]).

55 The key idea of our approach is the usage of novel assertions, which allow to specify
56 *thread-specific* views on shared variables. We also include a specific assertion containing
57 a modality for release-acquire (RA) synchronisation, capturing particularities of C11 RA
58 message passing. The use of non-standard assertions as a consequence necessitates the
59 introduction of new rules of assignment, formalising the effect of assignments on assertions.

60 We build our proof calculus on top of an *operational* semantics for C11 RAR. The
61 semantics is a mixture of the operational semantics proposed by Doherty et al. [12] (for RAR)
62 and Kaiser et al.'s semantics [16] for RA plus non-atomics. Correctness of this novel proposal
63 is shown by proving it to coincide with the semantics in [12] which in turn has been proven
64 to coincide with the standard axiomatic semantics of Batty et al. [9]. We have formalised
65 our semantics within the theorem prover Isabelle [28] and mechanically proved soundness of
66 all of our new rules for C11 assertions. Moreover, we provide mechanical proofs¹ of several
67 litmus tests from the literature (message passing, load buffering, read-read coherence) as
68 well as a version of Peterson's algorithm adapted for C11 memory [12, 36].

69 *Overview.* The paper is organised as follows. In the next section we start with an example
70 explaining the behaviour of concurrent programs on C11, motivating our novel assertions.
71 Section 3 defines the syntax of C11 RAR programs and Section 4 its semantics. We present
72 the proof calculus and its novel assertions in Section 5 via proofs of correctness for some
73 standard litmus tests, and a case study of Peterson's algorithm in Section 6. Section 7
74 describes our Isabelle mechanisation, Section 8 discusses related work and the last section
75 concludes.

76 **2 Deductive Reasoning for Weak Memory**

77 In this section, we illustrate the basic principles of C11 synchronisation and our verification
78 method by considering the message-passing example (Figures 1 and 2). The two programs
79 are almost identical and consist of two threads executing in parallel, accessing shared variables.
80 The assertions in curly brackets at the end specify the programs' postconditions.

81 The programs comprise two shared variables: d (that stores some data) and f (that stores
82 a flag). In both programs, both d and f are initially 0. thread 1 updates d to 5, then updates
83 f to 1. Thread 2 waits for f to be set to 1, then reads from d . Under sequential consistency,
84 one would expect that the final value of $r2$ is 5, since the loop in thread 2 only terminates
85 after f has been updated to 1 in thread 1, which in turn happens after d has been set to 5.
86 However, the C11 semantics allows the behaviour in Figure 2, where thread 2 may read a
87 stale value of d , and hence only the weaker postcondition $r2 = 0 \vee r2 = 5$ holds. To regain
88 the expected behaviour, one must introduce additional synchronisation in the program. In

¹ The Isabelle files may be downloaded from: <https://www.dropbox.com/sh/4yr2w7792qwyw09/AACsWUXtZbK3PvyfJkqjyDYa> within the file ECOOP-2020-Isabelle.zip.

<pre> Init: $d := 0; f := 0;$ Thread 1 \parallel Thread 2 $d := 5;$ \parallel do $r1 \leftarrow^A f$ $f :=^R 1;$ \parallel until $r1 = 1;$ \parallel $r2 \leftarrow d;$ $\{r2 = 5\}$ </pre>	<pre> Init: $d := 0; f := 0;$ Thread 1 \parallel Thread 2 $d := 5;$ \parallel do $r1 \leftarrow f$ $f := 1;$ \parallel until $r1 = 1;$ \parallel $r2 \leftarrow d;$ $\{r2 = 0 \vee r2 = 5\}$ </pre>
---	---

■ **Figure 1** Message-passing litmus test

■ **Figure 2** Unsynchronised message passing

89 particular, the write to f by thread 1 must be a *releasing write* (i.e., $f :=^R 1$) and the read
90 of f in thread 2 must be an *acquiring read* (i.e., $r1 \leftarrow^A f$) as in Figure 1.

91 In sequential consistency all threads have a single common view of the shared state,
92 namely all threads see the latest write that occurs for each variable. When a new write is
93 executed, the views of all threads are updated so that they see this write. In contrast, each
94 thread in C11 programs has its own view of each variable, which is affected by synchronisation
95 annotations. Thus, for the program in Figure 2, after initialisation, all threads see the initial
96 writes (i.e., $d = 0, f = 0$). The assignments in thread 1 only change thread 1's view, and
97 leave thread 2's view unchanged. Thus, after execution of $f := 1$, thread 2 has access to two
98 values for d (i.e., $d \in \{0, 5\}$) and f (i.e., $f \in \{0, 1\}$). Even if thread 2 reads $f = 1$, its view of
99 d remains unchanged and it continues to have access to both values of d .

100 The program in Figure 1 has a similar semantics for initialisation and execution of thread 1,
101 i.e., its execution does not affect the view of thread 2. However, due to the release-acquire
102 synchronisation on f (notation R and A), after thread 2 reads $f = 1$, its view for d will be
103 updated so that the stale value $d = 0$ is no longer available for it to read. One way to explain
104 this behaviour is by thinking of thread 1 as *passing its knowledge of the write to d* to thread
105 2 via the variable f , which is synchronised using the release-acquire annotations.

106 This intuition is captured formally using a semantics based on *timestamps* [16, 13, 17, 29],
107 which enables one to encode each thread's view and define how these views are updated. In
108 this paper, we characterise the release-acquire-relaxed subset of C11 [12] (C11 RAR) using
109 timestamps, which has a restriction prohibiting the so-called *load-buffering* litmus test [20].

110 The main contribution of our paper is an assertion language that enables one to reason
111 about thread views in a Hoare-style proof calculus, resulting in the proof outline given in
112 Figure 3. As already noted, the key advantage of these assertions is the fact that standard
113 rules of Hoare and Owicki-Gries logic remain unchanged. For message passing, we require
114 three main types of assertions (see Section 5):

115 **Possible value.** A possible value assertion (denoted $x \approx_t n$) states that thread t can read
116 value n of global variable x , i.e., there is a write to x with value n beyond or including
117 the *viewfront*² of thread t . Note that there may be more than one such write, and hence
118 there may be several possible values for a given variable. Moreover, the last write to each
119 variable is always viewable as a possible value.

120 **Definite value.** A definite value assertion (denoted $x =_t n$) states that thread t 's viewfront
121 is up-to-date with the writes to x (i.e., there is a single write to x beyond or including
122 the viewfront of thread t), and this write updates x 's value to n . Thus, t definitely knows
123 the variable x to have value n .

² We borrow the term viewfront from Popkadaev et al. [29].

$$\begin{array}{c}
\text{Init: } d := 0; f := 0; \\
\{f =_1 0 \wedge f =_2 0 \wedge d =_1 0 \wedge d =_2 0\} \\
\begin{array}{c}
\text{Thread 1} \\
\{f \not\approx_2 1 \wedge d =_1 0\} \\
1 : d := 5; \\
\{f \not\approx_2 1 \wedge d =_1 5\} \\
2 : f :=^R 1; \\
\{true\}
\end{array}
\parallel
\begin{array}{c}
\text{Thread 2} \\
\{[f = 1](d =_2 5)\} \\
3 : \text{do } r1 \leftarrow^A f \text{ until } r1 = 1; \\
\{d =_2 5\} \\
4 : r2 \leftarrow d; \\
\{r2 = 5\}
\end{array} \\
\{r2 = 5\}
\end{array}$$

■ **Figure 3** Proof outline for message passing

124 **Conditional value.** A conditional value assertion (denoted $[x = n](y =_t m)$) captures the
125 message passing idiom for variable y via variable x . It guarantees that when thread t
126 reads x to be n via an acquiring read, a release-acquire synchronisation is induced and
127 thereby t learns the definite value of y to be m . In particular, after reading $x = n$ via
128 an acquiring read, the viewfront for t is updated so that the only write to y beyond or
129 including this viewfront is a write with value m .

130 For the example in Figure 3, after initialisation, both threads 1 and 2 have definite value 0
131 for both d and f . The precondition of $d := 5$ states that thread 2 cannot possibly observe 1
132 for f (i.e., $f \approx_2 1$) and thread 1 definitely observes 0 for d (i.e., $d =_1 0$). These assertions can
133 be proven *locally correct* and *interference free* since thread 2 neither modifies d nor f . The
134 precondition of $f :=^R 1$ is similar but with $d =_1 5$ in place of $d =_1 0$. The precondition of the
135 **until** loop in thread 2 contains a conditional value assertion, which ensures that if thread 2
136 reads $f = 1$ then it will definitely read $d = 5$. This conditional value assertion enables one to
137 establish local correctness of the precondition (i.e., $d =_2 5$) of the statement $r2 \leftarrow d$, which
138 leads to the postcondition of the program. Each of the assertions in thread 2 can be proven
139 to be interference free against thread 1.

140 3 Program Syntax

141 We start by defining the syntax of concurrent programs, starting with the structure of
142 sequential programs (single threads). A thread may use *global* shared variables (from Var_G)
143 and local registers (from Var_L). We let $Var = Var_G \cup Var_L$ and assume $Var_G \cap Var_L = \emptyset$.
144 Global variables can be accessed in three different *synchronisation modes*: acquire (A, for
145 reads), release (R, for writes) and relaxed (no annotation). The annotation RA is employed
146 for *update* operations, which read and write to a shared variable in a single atomic step. We
147 use x, y, z to range over global variables and $r1, r2, \dots$ to range over local variables. We
148 assume that \ominus is a unary operator (e.g., \neg), \oplus is a binary operator (e.g., $\wedge, +, =$) and n
149 is a value (of type Val). Expressions may only involve local variables. For a treatment of
150 expressions with global variables in the semantics see [12]. The syntax of sequential programs,
151 Com , is given by the following grammar (with $r \in Var_L, x \in Var_G$):

$$\begin{array}{l}
152 \quad Exp_L ::= Val \mid r \mid \ominus Exp_L \mid Exp_L \oplus Exp_L \\
\quad ACom ::= \mathbf{skip} \mid x.\mathbf{swap}(n)^{RA} \mid r := Exp_L \mid x :=^{[R]} Exp_L \mid r \leftarrow^{[A]} x \\
\quad Com ::= ACom \mid Com; Com \mid \mathbf{if } B \mathbf{ then } Com \mathbf{ else } Com \mid \mathbf{while } B \mathbf{ do } Com
\end{array}$$

153 where we assume B to be an expression of type Exp_L that evaluates to a boolean. The
154 statement $x.\mathbf{swap}(n)^{RA}$ atomically reads the variable x (using an acquiring read) and updates

155 x to value n (using a releasing write) in a single atomic step. Its execution therefore gives
 156 rise to an atomic read-modify-write update event. We have not included a **CAS** operation
 157 here; it could similarly be implemented by an update event (see e.g. [35]).

158 The notation $[X]$ denotes that the annotation X is optional, where $X \in \{A, R\}$, enabling
 159 one to distinguish relaxed, acquiring and releasing accesses. Loops will be used in other
 160 forms, like **do-until** or **do-while**, which are straightforward to define in terms of the command
 161 syntax above.

162 As is standard in Owicki-Gries proofs, we make use of *auxiliary variables*, which are
 163 variables that do not affect the meaning of a program, but appear in proof assertions. We
 164 require that each auxiliary variable is *local* to the thread in which it occurs. Auxiliary
 165 variables may only occur in assignments, not in conditional statements, and only in the form
 166 $a := E$, where $E \in Exp_L$ and a is an auxiliary variable³. Finally, we require that writes
 167 to auxiliary variables occur atomically in conjunction with another (non-auxiliary) atomic
 168 program step. Such atomic operations are written as $\langle A, a := E \rangle$, where $A \in ACom$. This
 169 is more of a technical requirement which could also easily be relaxed. It guarantees that
 170 the programs without and with auxiliary variables have the same number of transitions (no
 171 stuttering steps).

172 For simplicity, we assume concurrency at the top level only. We let Tid be the set of
 173 all thread identifiers and use a function $Prog : Tid \rightarrow Com$ to model a program comprising
 174 multiple threads. In examples, we typically write concurrent programs as $C_1 || \dots || C_n$, where
 175 $C_i \in Com$. We further assume some initialisation of variables. The structure of our programs
 176 thus is **Init**; $(C_1 || \dots || C_n)$.

177 4 Semantics

178 The operational semantics for this language is defined in two parts. The *program semantics*
 179 fixes the steps that the concurrent program can take. This gives rise to transitions $(P, lst) \xrightarrow{a}_t$
 180 (P', lst') of a thread t where P and P' are programs, lst and lst' is the state of local variables
 181 and a is an action (possibly the silent action τ , see below). The program semantics is
 182 combined with a *memory semantics* which reflects the C11 state (denoted by σ), and in
 183 particular the write actions from which a read action can read.

184 We start by fixing the actions, where $x \in Var_G$ and $m, n \in Val$:

$$185 \text{Act} = \{rd(x, n), rd^A(x, n), wr(x, n), wr^R(x, n), upd^{RA}(x, n, m)\}$$

186 containing actions for (releasing) reads, (acquiring) writes and updates (reading value n and
 187 writing m). We furthermore employ a silent τ action and let $\text{Act}_\tau = \text{Act} \cup \{\tau\}$. For an action
 188 $a \in \text{Act}$, we let $var(a) \in Var_G$ be the variable read (or written to), $rdval(a) \in Val$ be the
 189 value read and $wrval(a) \in Val$ be the value written. We let U denote the update actions, and
 190 distinguish the sets $W_R \supseteq U$ (write release), $R_A \supseteq U$ (read acquire), W_X (write relaxed) and
 191 R_X (read relaxed). Finally, we define $R = R_A \cup R_X$ (all reads) and $W = W_R \cup W_X$ (all writes).
 192 Typically, we refer to the elements of W as *writes*, but note that this set also includes update
 193 actions.

³ The locality requirement is the only difference to “normal” Owicki-Gries auxiliary variables.

$$\begin{array}{c}
\frac{r \in \text{Var}_L \quad n = \llbracket E \rrbracket_{ls}}{(r := E, ls) \xrightarrow{\tau} (\mathbf{skip}, ls[r := n])} \qquad \frac{x \in \text{Var}_G \quad a = wr^{[R]}(x, \llbracket E \rrbracket_{ls})}{(x :=^{[R]} E, ls) \xrightarrow{a} (\mathbf{skip}, ls)} \\
\\
\frac{a = rd^{[A]}(x, n) \quad n \in \text{Val}}{(r \leftarrow^{[A]} x, ls) \xrightarrow{a} (\mathbf{skip}, ls[r := n])} \qquad \frac{a = upd^{\text{RA}}(x, m, n) \quad m \in \text{Val}}{(x.\mathbf{swap}(n)^{\text{RA}}, ls) \xrightarrow{a} (\mathbf{skip}, ls)} \\
\\
\frac{(C_1, ls) \xrightarrow{a} (C'_1, ls')}{(C_1; C_2, ls) \xrightarrow{a} (C'_1; C_2, ls')} \qquad \frac{}{(\mathbf{skip}; C_2, ls) \xrightarrow{\tau} (C_2, ls)} \\
\\
\frac{\llbracket B \rrbracket_{ls}}{(\mathbf{if } B \mathbf{ then } C_1 \mathbf{ else } C_2, ls) \xrightarrow{\tau} (C_1, ls)} \qquad \frac{\neg \llbracket B \rrbracket_{ls}}{(\mathbf{if } B \mathbf{ then } C_1 \mathbf{ else } C_2, ls) \xrightarrow{\tau} (C_2, ls)} \\
\\
\frac{\llbracket B \rrbracket_{ls}}{(\mathbf{while } B \mathbf{ do } C, ls) \xrightarrow{\tau} (C; \mathbf{while } B \mathbf{ do } C, ls)} \qquad \frac{\neg \llbracket B \rrbracket_{ls}}{(\mathbf{while } B \mathbf{ do } C, ls) \xrightarrow{\tau} (\mathbf{skip}, ls)} \\
\\
\text{AUX} \frac{(A, ls) \xrightarrow{a} (\mathbf{skip}, ls') \quad (a := E, ls') \xrightarrow{\tau} (\mathbf{skip}, ls'')}{\langle\langle A; a := E \rangle\rangle, ls) \xrightarrow{a} (\mathbf{skip}, ls'')} \qquad \text{PROG} \frac{(P(t), lst(t)) \xrightarrow{a} (C, ls) \quad a \in \text{Act}_\tau}{(P, lst) \xrightarrow{a}_t (P[t := C], lst[t := ls])}
\end{array}$$

■ **Figure 4** Program semantics

194 4.1 Program Semantics

195 In the program semantics, we assume a function $lst \in \text{Tid} \rightarrow (\text{Var}_L \rightarrow \text{Val})$, which returns
196 the local state for the given thread. We assume that the local variables of threads are disjoint,
197 i.e., if $t \neq t'$, then $\text{dom}(lst(t)) \cap \text{dom}(lst(t')) = \emptyset$. For an expression E over local variables,
198 we write $\llbracket E \rrbracket_{ls}$ for the value of E in local state ls ; we write $ls[r := n]$ to state that ls remains
199 unchanged except for the value of local variable r which becomes n .

200 Figure 4 gives the transition rules of the program semantics. The last rule, **Prog**, lifts
201 the transitions of threads to a transition for a concurrent program. The other rules concern
202 the sequential part of the language. The rules in a sense ignore the fact that the language
203 allows for global variables; the program semantics just details the values of local variables in
204 component ls . When global variables are read, the program semantics allows for *all* possible
205 values to be read. This is combined with the memory semantics (formalised by \xrightarrow{a}_t) as
206 follows:

$$\begin{array}{c}
\frac{(P, lst) \xrightarrow{\tau}_t (P', lst')}{(P, lst, \sigma) \Longrightarrow (P', lst', \sigma)} \qquad \frac{(P, lst) \xrightarrow{a}_t (P', lst') \quad \sigma \xrightarrow{a}_t \sigma'}{(P, lst, \sigma) \Longrightarrow (P', lst', \sigma')}
\end{array}$$

209 The transitions defined by $\sigma \xrightarrow{a}_t \sigma'$ ensure that read actions only return a value allowed
210 by the C11 semantics and are defined in Section 4.2. The rules for all imperative program
211 constructs (sequential composition, **if** and **while**) are standard.

212 4.2 Memory Semantics

213 Next, we detail the memory semantics, which is equivalent to an earlier operational reformu-
214 lation [12] of the RAR fragment from [20].

215 **C11 State.** Table 1 summarises the components of a C11 state. Each global write is
216 represented by a pair $(a, q) \in \mathbb{W} \times \mathbb{Q}$, where a is a write action, and q is a rational number

■ **Table 1** Components of a C11 state

Component	Informal meaning	Initial value
$writes \subseteq W \times \mathbb{Q}$	The writes which have happened so far	$writes_{\mathbf{Init}}$
$tview_t \in Var_G \rightarrow writes$	The view of a thread t	$tview_{\mathbf{Init}}$
$mview_w \in Var_G \rightarrow writes$	The view of a thread when writing w	$mview_{\mathbf{Init}}$
$covered \subseteq writes$	The covered writes	\emptyset

217 that we use as a *timestamp* (c.f., [16, 13, 29]). The timestamps totally order the writes to
 218 each variable; the ordering induced by timestamps is also referred to as the *modification*
 219 *order* [20, 12] or *coherence order* [3]. For each write $w = (a, q)$, we denote w 's timestamp
 220 by $tst(w) = q$. We also lift the functions *var* and *wrval* to timestamped writes, e.g.,
 221 $var((a, q)) = var(a)$. The set of all writes that have occurred in the execution thus far is
 222 recorded in the state component $writes \subseteq W \times \mathbb{Q}$.

223 As described in Section 2, each state must record the writes that are observable to each
 224 read. To achieve this, we use two families of functions from global variables to writes, both
 225 of which record the *viewfronts* (c.f., [29, 17]).

226 ■ A function $tview_t$ that returns the *viewfront* of thread t . The thread t can read from any
 227 write to variable x whose timestamp is not earlier than $tview_t(x)$. Accordingly, we define,
 228 for each state σ , thread t and global variable x , the set of *observable writes*:

$$229 \quad \sigma.OW(t, x) = \{(a, q) \in \sigma.writes \mid var(a) = x \wedge tst(\sigma.tview_t(x)) \leq q\} \quad (1)$$

231 ■ A function $mview_w$ that records the *viewfront* of write w , which is set to be the viewfront
 232 of the thread that executed w at the time of w 's execution. We use $mview_w$ to compute
 233 a new value for $tview_t$ if a thread t *synchronizes* with w , i.e., if $w \in W_R$ and another
 234 thread executes an $e \in R_A$ that reads from w .

235 Finally, our semantics maintains a variable $covered \subseteq writes$. In C11 RAR, each update
 236 action occurs in modification order immediately after the write that it reads from [12]. This
 237 property constitutes the atomicity of updates. In order to preserve this property, we must
 238 prevent any newer write from intervening between any update and the write that it reads
 239 from. As we explain below, *covered* writes are those that are immediately prior to an update
 240 in modification order, and new write actions never interact with a covered write.

241 **Initialisation.** Table 1 also states how these components are initialised by **Init**. If $Var_G =$
 242 $\{x_1, \dots, x_n\}$, $Var_L = \{r_1, \dots, r_m\}$ and $k_1, \dots, k_n, l_1, \dots, l_m \in Val$, we assume **Init** = $x_1 :=$
 243 $k_1; \dots; x_n := k_n; [r_1 := l_1;] \dots [r_m := l_m;]$, where we use the notation $[r_i := l_i;]$ to mean
 244 that the assignment $r_i := l_i$ may optionally appear in **Init**. Thus each shared variable is
 245 initialised exactly once and each local variable is initialised at most once. The initial values
 246 of the state components are then as follows, where we assume that 0 is the initial timestamp.

$$247 \quad writes_{\mathbf{Init}} = \{(wr(x_1, k_1), 0), \dots, (wr(x_n, k_n), 0)\}$$

$$248 \quad tview_{\mathbf{Init}}(x_i) = (wr(x_i, k_i), 0) \quad \text{for each thread } x_i \in Var_G$$

$$249 \quad mview_{\mathbf{Init}} = tview_{\mathbf{Init}}$$

251 The local state component of each thread must also be compatible with **Init**, i.e., for each t
 252 if $r_i \in \mathbf{dom}(lst(t))$ we have that $(lst(t))(r_i) = l_i$ provided $r_i := l_i$ appears in **Init**.

253 We let $lst_{\mathbf{Init}}$ be the local state compatible with **Init**, let $\sigma_{\mathbf{Init}}$ denote the initial state
 254 defined by **Init**, and define $\Gamma_{\mathbf{Init}} = (lst_{\mathbf{Init}}, \sigma_{\mathbf{Init}})$.

$$\begin{array}{c}
a \in \{rd(x, n), rd^A(x, n)\} \quad (w, q) \in \sigma.OW(t, x) \quad wrval(w) = n \\
tview'_t = \begin{cases} \sigma.tview_t \otimes \sigma.mview_{(w, q)} & \text{if } (w, a) \in \mathbb{W}_R \times \mathbb{R}_A \\ \sigma.tview_t[x := (w, q)] & \text{otherwise} \end{cases} \\
\text{READ} \text{-----} \\
\sigma \xrightarrow{a}_t \sigma[tview_t := tview'_t] \\
\\
a \in \{wr(x, n), wr^R(x, n)\} \quad (w, q) \in \sigma.OW(t, x) \setminus \sigma.covered \quad \sigma.fresh(q, q') \\
writes' = \sigma.writes \cup \{(a, q')\} \quad tview'_t = \sigma.tview_t[x := (a, q')] \\
\text{WRITE} \text{-----} \\
\sigma \xrightarrow{a}_t \sigma[tview_t := tview'_t, mview_{(a, q')} := tview'_t, writes := writes'] \\
\\
a = upd^{RA}(x, m, n) \quad (w, q) \in \sigma.OW(t, x) \setminus \sigma.covered \\
wrval(w) = m \quad \sigma.fresh(q, q') \\
writes' = \sigma.writes \cup \{(a, q')\} \quad covered' = \sigma.covered \cup \{(w, q)\} \\
tview'_t = \begin{cases} \sigma.tview_t[x := (a, q')] \otimes \sigma.mview_{(w, q)} & \text{if } w \in \mathbb{W}_R \\ \sigma.tview_t[x := (a, q')] & \text{otherwise} \end{cases} \\
\text{UPDATE} \text{-----} \\
\sigma \xrightarrow{a}_t \sigma[tview_t := tview'_t, mview_{(a, q')} := tview'_t, \\
writes := writes', covered := covered']
\end{array}$$

■ **Figure 5** Transition relation of the memory semantics

255 **Transition semantics.** The transition relation of our semantics for global reads and writes
256 is given in Figure 5. Each transition $\sigma \xrightarrow{a}_t \sigma'$ is labelled by an action a and thread t . The
257 premise of each rule must identify the write w that the action interacts with. This is made
258 more precise below.

259 **READ transition by thread t .** Here we assume that

- 260 ■ a is either a relaxed or acquiring read to variable x ,
- 261 ■ w is a write to x that t can observe (i.e., $(w, q) \in \sigma.OW(t, x)$), and
- 262 ■ the value read by a is the value written by w .

263 Each read causes the viewfront of t to be updated. This is computed as follows. If the read
264 synchronises with the write, then the thread's new view will be a combination of its existing
265 view, and the view of that write. In particular, for each variable x the new view of x will
266 be the later of either $tview_t(x)$ or $mview_w(x)$, in timestamp order. To express this, we use
267 an operation that combines two views v_1 and v_2 , by constructing a new view that takes the
268 later of the writes at each variable:


$$269 \quad (v_1 \otimes v_2)(x) = \begin{cases} v_1(x) & \text{if } tst(v_2(x)) \leq tst(v_1(x)) \\ v_2(x) & \text{otherwise} \end{cases}$$

270

271 If w and a do not synchronise, then $tview_t$ is simply updated to include the new write.

272 For illustration, consider the picture in Figure 6. The x-axis depicts the timestamps of
273 the writes, the y-axis the variables x, y and z , which we assume are initialised by writes x_0 ,
274 y_0 and z_0 , respectively. The orange line shows the view of a thread, say t_1 , and the blue line
275 depicts the view of another thread that executes $w = (wr^R(y, 42), 3)$. If thread t_1 performs
276 an acquiring read of y and reads from w (i.e., it performs a synchronising read), thread t_1 's
277 view changes to the diagram on the right, whereby its current viewfront is combined with
278 the viewfront of w .

279 **WRITE transition by thread t .** A write transition must identify the write (w, q) after
280 which a occurs. This w must be observable and must *not* be covered — the second condition



■ **Figure 6** Illustration of views and view updates: pre-state (left) and post-state (right)

281 is required to preserve the read-modify-write atomicity of updates. We must choose a fresh
 282 timestamp $q' \in \mathbb{Q}$ for a , which is formalised by $fresh(q, q')$:

$$283 \quad \sigma.fresh(q, q') = q < q' \wedge \forall w' \in \sigma.writes. q < tst(w') \Rightarrow q' < tst(w')$$

284 The predicate $fresh(q, q')$ ensures that q' is a new timestamp for the variable x , such that
 285 (a, q') occurs immediately after (w, q) . The new write is added to the set $writes$. We update
 286 $tview_t$ to include the new write, which means t can no longer observe any writes prior to (a, q') .
 287 Finally, we set the viewfront of (a, q') to be the new viewfront of t , i.e., $mview_{(a, q')} := tview'_t$.
 288 Now, if some other thread synchronises with this new write in some later transition, that
 289 thread's view will become at least as recent as t 's view at this transition.

290 **UPDATE transition by thread t .** These transitions are best understood as a combination
 291 of the read and write transitions. As with a write transition, we must choose a valid fresh
 292 q' , and the state components $writes$ and $mview$ are updated in the same way. As discussed
 293 earlier, in UPDATE transitions it is necessary to record that the write that the update interacts
 294 with is now covered, which is achieved by adding that write to $covered$. Finally, we must
 295 compute a new thread view, which is similar to a READ transition, except that the thread's
 296 new view always includes the new write introduced by the update.

297 4.3 Relationship to the Axiomatic Semantics

298 We prove that the timestamp-based semantics presented here is equivalent to an earlier
 299 operational semantics [12] that is already known to be equivalent to the C11 RAR fragment.
 300 Here, we just roughly sketch how this proof proceeds, the appendix contains more details.

301 The semantics in [12] describes C11 states in the form $E = (X, sb, rf, mo)$, where X is
 302 a set of read and write events (roughly equivalent to actions) and sb , rf and mo describe
 303 the sequenced-before and reads-from relation as well as the modification order of the C11
 304 axiomatic semantics. A number of further relations are derived from these, in particular the
 305 extended coherence order eco and the happens-before order hb . The proof of equivalence of
 306 the semantics shows the two semantics to *simulate* each other. For this, we need to define a
 307 correspondence between C11 states of form E and of form σ such that: (1) For $\sigma.writes$, we
 308 take $X \cap W$; (2) For $\sigma.covered$, we take the writes w in $X \cap W$ such that there is an update u
 309 with $(w, u) \in rf$; and (3) For $mview$ and $tview$, we use a downward closure operator, $cclose$,
 310 which for a given set of events S determines the set of events prior to S in the relation $eco^? \circ hb^?$.
 311 Then $\sigma.tview_t = \mathbf{max}_{mo}(X.cclose(X_t))$ and $\sigma.mview_w = \mathbf{max}_{mo}(X.cclose(\{w\}))$, where
 312 \mathbf{max}_{mo} selects writes being maximal wrt. mo and X_t are all actions of t in X . In all these
 313 cases, timestamps for writes have to be selected consistent with mo .

314 Given such a correspondence, the proof proceeds by showing this correspondence is
 315 preserved by the read, write and update transitions.

316 4.4 Well Formedness

317 Our proofs in subsequent sections require that the state under consideration is *well-formed*.
 318 This is formalised by predicate wfs over a C11 state σ , where

$$\begin{aligned}
 319 \quad wfs(\sigma) \iff & \text{ran}((\bigcup_t \sigma.tview_t) \cup (\bigcup_w \sigma.mview_w)) \subseteq \sigma.writes \wedge \\
 320 & \text{finite}(\sigma.writes) \wedge \sigma.covered \subseteq \sigma.writes \wedge \\
 321 & (\forall w. w \in \sigma.writes \Rightarrow \sigma.mview_w(\text{var}(w)) = w)
 \end{aligned}$$

323 The first conjunct ensures that each viewable write is in $\sigma.writes$. The second conjunct
 324 ensures there are only a finite number of writes, and the third ensures that every covered
 325 write is an actual write. The final conjunct ensures that for each write in $\sigma.writes$, the
 326 viewfront of w for $\text{var}(w)$ is w itself.

327 Well-formedness is invariant for any program, i.e., every initialisation establishes well-
 328 formedness and every program transition preserves well-formedness.

329 ► **Lemma 1.** *For any program C constructed using the syntax described in Section 3, $wfs(\sigma)$*
 330 *is invariant.*

331 **Proof.** In Isabelle. We show that every initialisation establishes $wfs(\sigma)$. Furthermore, if
 332 $wfs(\sigma)$ and $\sigma \xrightarrow{a}_t \sigma'$, then $wfs(\sigma')$ for any action a and thread t . ◀

333 5 Hoare Logic and Owicki-Gries Reasoning for C11

334 In this section, we present a Hoare logic [15] for C11 RAR that enables Owicki-Gries
 335 reasoning [27]. For compound statements (including concurrent composition) we use the
 336 standard rules of Hoare logic as well as the standard interference freedom proof obligations
 337 described by Owicki and Gries. Our contribution is a novel set of high-level predicates
 338 that describe the *observations* of each thread for a C11 state, together with a set of *basic*
 339 *axioms* that describe how these predicates interact with read, write and update transitions.
 340 Soundness of these axioms has been checked using Isabelle.

341 In Section 5.1, we link our operational semantics to the proof outlines of Hoare logic
 342 and Owicki-Gries' notion of interference freedom. Section 5.2 provides an overview of our
 343 assertion language and briefly discusses the main categories of assertions, i.e., assertions
 344 describing observability, ordering and occurrences of writes. We present the basic axioms
 345 in stages, using specific litmus tests (in Sections 5.3, 5.4, 5.5) to motivate each group of
 346 assertions. The proof outlines of all litmus tests have been verified using Isabelle.

347 5.1 Soundness and Classical Verification Rules

348 We first define the meaning of a Hoare triple under partial correctness and present the
 349 classical proofs rules for compound statements. Unlike Hoare logic, where a state is modelled
 350 by a mapping from variables to values, as we have seen in Section 4.1, states of a C11
 351 program contain two components: a local state lst and a global state σ . We let Σ_G
 352 be the set of all possible global state configurations (as described in Table 1) and let
 353 $\Sigma_{C11} = (\text{Var}_L \rightarrow \text{Val}) \times \Sigma_G$ be the set of all possible C11 states. Predicates over Σ_{C11} are
 354 therefore of type $\Sigma_{C11} \rightarrow \mathbb{B}$. This leads to the following definition of a Hoare triple, which we
 355 note is the same as the standard definition — the only difference is that the state component
 356 is of type Σ_{C11} .

$$\begin{array}{c}
\text{SKIP} \frac{}{\{p\} \mathbf{skip}\{p\}} \quad \text{SEQ} \frac{\{p\}C_1\{r\} \quad \{r\}C_2\{q\}}{\{p\}C_1; C_2\{q\}} \\
\\
\text{IF} \frac{\{p \wedge B\}C_1\{q\} \quad \{p \wedge \neg B\}C_2\{q\}}{\{p\} \mathbf{if } B \mathbf{ then } C_1 \mathbf{ else } C_2\{q\}} \quad \text{WHILE} \frac{\{p \wedge B\}C\{p\}}{\{p\} \mathbf{while } B \mathbf{ do } C\{p \wedge \neg B\}} \\
\\
\text{UNTIL} \frac{\{p\}C\{r\} \quad \{r\} \mathbf{while } \neg B \mathbf{ do } C\{r \wedge B\}}{\{p\} \mathbf{do } C \mathbf{ until } B\{r \wedge B\}} \quad \text{CONS} \frac{p \Rightarrow p' \quad \{p'\}C\{q'\} \quad q' \Rightarrow q}{\{p\}C\{q\}}
\end{array}$$

■ **Figure 7** Classical proof rules for sequential programs

357 ► **Definition 2.** Suppose $p, q \in \Sigma_{C11} \rightarrow \mathbb{B}$, $P \in \text{Prog}$ and $\mathbf{E} = \lambda t : \text{Tid. skip}$. The semantics
358 of a Hoare triple under partial correctness is given by:

$$\begin{array}{c}
359 \quad \{p\} \mathbf{Init}\{q\} = q(\Gamma_{\mathbf{Init}}) \\
360 \quad \{p\}P\{q\} = \forall lst, \sigma, lst', \sigma'. p(lst, \sigma) \wedge (P, lst, \sigma) \Longrightarrow^* (\mathbf{E}, lst', \sigma') \Rightarrow q(lst', \sigma') \\
361 \quad \{p\} \mathbf{Init}; P\{q\} = \exists r. \{p\} \mathbf{Init}\{r\} \wedge \{r\}P\{q\} \\
362
\end{array}$$

363 The classical rules of sequential Hoare logic for compound (i.e., non-atomic) statements are
364 given in Figure 7. Soundness of these proof rules (with respect to Definition 2) holds for
365 exactly the same reason as soundness of Hoare logic [15].

366 The sequential part is combined with the Owicki-Gries rule for concurrent composition
367 in the standard way [27, 7]. First, we construct *proof outlines* for every component of the
368 concurrent program in isolation. A proof outline inserts assertions (in $\{ \}$ brackets) into a
369 program. In a so-called *standard* proof outline every statement R of the program has exactly
370 one assertion before it. This assertion is its *precondition*, $pre(R)$. Next, all assertions in one
371 component have to be checked for non-interference with all statements in other components.

372 ► **Definition 3.** A statement $R \in ACom$ with precondition $pre(R)$ (in the standard proof
373 outline) does not interfere with an assertion p if

$$374 \quad \{p \wedge pre(R)\} R \{p\} .$$

375 Interference freedom guarantees that proof outlines in each thread are stable under the
376 execution of other threads. This is formalised in the Owicki-Gries proof rule for concurrent
377 composition:

$$378 \quad \text{PARALLEL} \frac{\text{Proof outlines } \{p_i\}C_i\{q_i\} \text{ are interference free}}{\{\bigwedge_{i=1}^n p_i\} C_1 \parallel \dots \parallel C_n \{\bigwedge_{i=1}^n q_i\}}$$

380 We say a proof outline is *valid* if it is both sequentially valid (or locally correct) and
381 interference free.

382 Finally, there is a standard proof rule for auxiliary variables in parallel programs [7]. Let
383 V be a set of auxiliary variables of a parallel program P and q be a predicate that does not
384 mention auxiliary variables. Then we can prove that a Hoare triple holds for a program
385 extended with auxiliary variables and transfer this proof to the original program:

$$386 \quad \text{AUXVAR} \frac{\{true\} \mathbf{Init}; P \{q\}}{\{true\} \mathbf{Init}_0; P_0 \{q\}} \text{ provided } vars(q) \cap V = \emptyset$$

387 where \mathbf{Init}_0 is obtained from \mathbf{Init} by removing all auxiliary assignments and P_0 is obtained
388 by replacing all statements $\langle A, a := E \rangle$ in P (for $a \in V$) by A .

389 **5.2 An Assertion Language**

390 We studied a number of well-known litmus tests and examples and discovered three main
 391 categories of assertions required for specification and verification of a wide range of problems.
 392 These three main categories are dealing with (values of) writes to variables and the order in
 393 which they occur.

394 ■ **Observability.** Observability assertions describe if or when a thread may observe or
 395 has encountered a write to a variable. As described in Section 2, these assertions are
 396 thread-specific and deal with the thread's view. We repeat the main ideas here to simplify
 397 comparison with the other types of assertions. The main observability assertions are as
 398 follows:

- 399 1. **Possible observation** which is denoted by $x \approx_t u$ means that thread t *may* observe
 400 value u for x . The formal definition and an example motivating this assertion is given
 401 in Section 5.4.
- 402 2. **Definite observation** which is denoted by $x =_t u$ means that thread t *must* observe
 403 the value u for x . The formal definition and an example motivating this assertion is
 404 given in Section 5.3.
- 405 3. **Conditional observation** which is denoted by $[x = u](y =_t v)$ means that if thread
 406 t synchronises with a write to variable x with value u , it *must* observe value v for y .
 407 The formal definition and an example motivating this assertion is given in Section 5.4.
- 408 4. **Encountered value** which is denoted by $x \stackrel{enc}{=}_t v$ means that thread t has encountered
 409 (had the opportunity to observe) a write to variable x with value v . The formal definition
 410 and three examples motivating this assertion are given in Section 5.5.

411 ■ **Ordering.** Ordering assertions specify the order of values written to a variable by
 412 different writes. These assertions are thread-independent and specify an order over the
 413 timestamp of various writes with specific values:

- 414 1. **Possible value order** which is denoted by $m \prec_x n$ means that there exists two writes
 415 w and w' to variable x where the timestamp of w' is larger than the timestamp of w
 416 and the value of w and w' is m and n , respectively.
- 417 2. **Definite value order** which is denoted by $m \ll_x n$ means that for all writes w and
 418 w' to x where the value of w is m and the value of w' is n , the timestamp of w' is
 419 larger than the timestamp of w and $m \prec_x n$.

420 Both the above assertions are formally defined in Section 5.5 and examples showing their
 421 usage are provided.

422 ■ **Occurrence.** Occurrence assertions specify the occurrence of a write with a specific
 423 value to a variable (regardless of observability). Similar to the previous category, these
 424 assertions are thread-independent:

- 425 1. **Value occurrence** assertions specify the limit of occurrence of writes to a variable
 426 with a specific value. For instance, $\mathbb{0}_x n$ means that no write with value n to variable
 427 x has occurred or $\mathbb{1}_x n$ means that there is *at most* one write with value n to x in
 428 the current state. The formal definition and examples of these assertions are given in
 429 Section 5.5.
- 430 2. **Initial value** which is denoted by $x_{\text{Init}} = n$ means that the initial value written to x
 431 is n . The formal definition and examples of this assertion are also given in Section 5.5.
- 432 3. **Covered write** assertions, denoted by \mathbf{C}_x^n , state that all writes to variable x except
 433 the last write are covered by an update (see Section 4.2), and that the last write to x
 434 has value n . This assertion is formally defined in Section 6 and is used in verification
 435 of Peterson's mutual exclusion algorithm.

$$\begin{array}{c}
\mathbf{Init}: x := 0; y := 0; r1 := 0; r2 := 0; \\
\{x =_1 0 \wedge y =_2 0 \wedge r1 = 0 \wedge r2 = 0\} \\
\mathbf{Thread 1} \quad \parallel \quad \mathbf{Thread 2} \\
\{y =_2 0 \wedge r2 = 0\} \quad \parallel \quad \{x =_1 0 \wedge r1 = 0\} \\
1 : r1 \leftarrow x; \quad \parallel \quad 3 : r2 \leftarrow y; \\
\{y =_2 0 \wedge r2 = 0\} \quad \parallel \quad \{x =_1 0 \wedge r1 = 0\} \\
2 : y := 1; \quad \parallel \quad 4 : x := 1; \\
\{r1 = 0 \vee r2 = 0\} \quad \parallel \quad \{r1 = 0 \vee r2 = 0\} \\
\{r1 = 0 \vee r2 = 0\}
\end{array}$$

■ **Figure 8** Proof outline for load buffering

5.3 Load Buffering

Our first example is the load buffering litmus test (see Figure 8), which we can show satisfies the postcondition $r1 = 0 \vee r2 = 0$ since our semantics assumes absence of cycles in the sequence-before relation combined with reads-from [20, 12]. The assertions about the C11 state capture properties about *definite observations* (i.e., observability assertions), which we formalise below.

For a set of writes W and variable $x \in \text{Var}_G$, let $W_x = \{w \in W \mid \text{var}(w) = x\}$ be the set of writes in W that write to x . We define the *last write* to x in W as:

$$\text{last}(W, x) = w \iff w \in W_x \wedge (\forall w' \in W_x. \text{tst}(w') \leq \text{tst}(w))$$

Moreover, we define the definite observation of a view function, *view* with respect to a set of writes as follows:

$$\text{dview}(\text{view}, W, x) = n \iff \text{view}(x) = \text{last}(W, x) \wedge \text{wrval}(\text{last}(W, x)) = n$$

The first conjunct ensures that the viewfront of *view* for x is the last write to x in W , and the second conjunct ensures that the value written by the last write to x in W is n .

Definite observation. For a variable x , thread t and value n , we define:

$$x =_t n = \lambda \sigma. \text{dview}(\sigma.\text{tview}_t, \sigma.\text{writes}, x) = n$$

Expanding this out, we obtain:

$$\sigma.\text{tview}_t(x) = \text{last}(\sigma.\text{writes}, x) \wedge \text{wrval}(\text{last}(\sigma.\text{writes}, x)) = n$$

The first conjunct ensures that the viewfront of t for x is the last write to x in σ (thus t can only read this last write to x). The second conjunct ensures that the value written by the last write is n . The function *dview* is also used in the definition of conditional observation in Section 5.4.

The proof of load buffering relies on the basic axioms in the following lemma. We assume *atoms(Init)* returns the set of assignments contained within **Init**.

► **Lemma 4.** *Each of the basic axioms below is sound (as per Definition 2), where the statements are decorated with the thread identifier of the executing thread.*

$$\text{INIT} \frac{x := n \in \text{atoms}(\mathbf{Init})}{\{true\} \mathbf{Init} \{x =_t n\}} \quad \text{DOPRES-RD} \frac{}{\{x =_{t'} m\} r \leftarrow_t^{[A]} y \{x =_{t'} m\}}$$

$$\text{DOPRES-WR} \frac{x \neq y}{\{x =_{t'} n\} y :=_t m \{x =_{t'} n\}}$$

466 **Proof.** In Isabelle. ◀

467 Thus by rule INIT an assignment $x := n$ in INIT ensures that $x =_t n$ for all threads t
 468 holds at program start. Note that such an initial assertion for the entire program is not
 469 subject to non-interference checks. The rule DOPRES-RD states that a definite observation
 470 $x =_{t'} m$ is invariant over a read step executed by thread t . Note that pre/post conditions
 471 for DOPRES-RD refer to thread t' , while the read statement refers to thread t . Also note
 472 that there is no additional restriction on t and t' , thus the rule applies regardless of whether
 473 $t = t'$, or not. Similarly, there are two global variables x and y mentioned in the rule, but
 474 there are no further restrictions on their values. Rule DOPRES-WR gives a condition for
 475 invariance of a definite observation assertion over a write. It requires that the variable being
 476 observed is different from the variable that is updated.

477 ▶ **Theorem 5.** *The proof outline for load buffering in Figure 8 is valid.*

478 **Proof.** The proof has been established in Isabelle. We outline the main steps below as it is
 479 instructive to understand the high-level proof strategy. First we establish local correctness:

- 480 ■ The initial condition is established by rule INIT, which is in turn used to establish the
 481 initial assertions in both threads.
- 482 ■ In thread 1, local correctness of the postcondition of line 1 (precondition of line 2) follows
 483 from rule DOPRES-RD, and the postcondition of line 2 follows by weakening. The proof
 484 of local correctness in thread 2 is symmetric.

485 We now establish interference freedom. The precondition of line 1 is interference free wrt
 486 line 3 by DOPRES-RD, and wrt line 4 by DOPRES-WR. This argument also applies to the
 487 precondition of line 2. Interference freedom of the postcondition of line 2 is trivial. The
 488 proof of interference freedom of the assertions in thread 2 is symmetric. ◀

489 5.4 Message Passing

490 Next we return to the message passing example from Section 2. Its verification requires the
 491 usage of the other two observability assertions.

492 **Possible observation.** For a variable x , thread t and value n , we define:

$$493 \quad x \approx_t n \quad = \quad \lambda\sigma. \exists w \in \sigma.OW(t, x). \text{wrval}(w) = n$$

494 Thus, there is a write to x that is observable to thread t with a value n .

Conditional observation. For variables x, y , thread t and values m, n , we define:

$$[x = n](y =_t m) \quad = \quad \lambda\sigma. \forall w \in \sigma.OW(t, x). \text{wrval}(w) = n \Rightarrow \\ \text{act}(w) \in W_R \wedge \text{dview}(\sigma.\text{mview}_w, \sigma.\text{writes}, y) = m$$

495 The antecedent assumes that the value read for x is n , and the consequent ensures that w
 496 is a releasing write such that the definite view of this write for variable y returns m . As
 497 we shall see, one useful way of establishing this condition is by falsifying the antecedent by
 498 ensuring that thread t cannot observe n for x (see (4) below).

499 Some useful relationships between the assertions above are given by the lemma below.

500 ► **Lemma 6.** For variables $x, y \in \text{Var}_G$, thread t and values $m, n \in \text{Val}$, each of the following
501 holds:

$$502 \quad \text{wfs} \wedge x =_t n \Rightarrow x \approx_t n \quad (2)$$

$$503 \quad \text{wfs} \wedge x =_t n \wedge x \approx_t m \Rightarrow n = m \quad (3)$$

$$504 \quad x \not\approx_t n \Rightarrow [x = n](y =_t m) \quad (4)$$

$$505 \quad x =_t n \wedge x =_{t'} m \Rightarrow n = m \quad (5)$$

507 **Proof.** In Isabelle. ◀

508 By (2), given a well-formed state any definite observation implies a possible observation,
509 and by (3) a definite observation must agree with a possible observation. By (4) if it is
510 not possible to observe the antecedent of a conditional observation, then the conditional
511 observation must hold. By (5) any two definite value observations must agree (since they
512 both observe the last write to x).

513 The next lemma lists the basic axioms that are used to prove correctness of the message
514 passing example.

515 ► **Lemma 7.** Each of the rules next is sound (as per Definition 2), where the statements are
516 decorated with the thread identifier of the executing thread.

$$517 \quad \text{MODLAST} \frac{}{\{x =_t n\} x :=_t m \{x =_t m\}} \quad \text{MODSOME} \frac{}{\{\text{true}\} x :=_t m \{x \approx_t m\}}$$

$$518 \quad \text{NPOPRES} \frac{}{\{x \not\approx_t m\} r \leftarrow_{t'}^{[A]} y \{x \not\approx_t m\}} \quad \text{NOOW} \frac{x \neq y}{\{x \not\approx_t n\} y :=_{t'} m \{x \not\approx_t n\}}$$

$$519 \quad \text{READLAST} \frac{}{\{x =_t m\} r \leftarrow_t x \{r = m\}}$$

$$520 \quad \text{CO-INTRO} \frac{x \neq y}{\{y =_t m \wedge x \not\approx_{t'} n\} x :=_t^R n \{[x = n](y =_{t'} m)\}}$$

$$521 \quad \text{TRANSFER} \frac{}{\{[x = n](y =_t m)\} r \leftarrow_t^A x \{r = n \Rightarrow y =_t m\}}$$

523 **Proof.** In Isabelle. ◀

524 ► **Theorem 8.** The proof outline of message passing in Figure 3 is valid.

525 **Proof.** The proof has been established in Isabelle. We outline the main steps below. First
526 we show local correctness.

- 527 ■ Using INIT we establish the precondition $f =_1 0 \wedge f =_2 0 \wedge d =_1 0 \wedge d =_2 0$.
- 528 ■ The precondition of the program implies the initial assertions of both threads. In thread 1,
529 we use (3) to establish $f \not\approx_2 1$ since (3) is logically equivalent to

$$530 \quad \text{wfs} \wedge x =_t n \wedge n \neq m \Rightarrow x \not\approx_t m$$

531 In thread 2, we use (3) in combination with (4).

- 532 ■ In thread 1, the post condition of line 1 (precondition of line 2) follows by application of
533 NOOW and MODLAST. The post condition of line 2 is trivial.
- 534 ■ In thread 2, the postcondition of line 3 follows by application of TRANSFER, while the
535 postcondition of line 4 follows by application of READLAST.

599 ► **Lemma 9.** For $x \in \text{Var}_G$ and $m, n \in \text{Val}$, we have:

$$600 \quad m \prec_x n \wedge \mathbb{1}_x m \wedge \mathbb{1}_x n \Rightarrow m \ll_x n \quad (6)$$

$$601 \quad m \ll_x n \Rightarrow n \not\prec_x m \quad (7)$$

603 **Proof.** In Isabelle. ◀

604 We discuss the proof of RRC2 in detail. Its proof relies on the following lemma which
605 captures some basic properties about value assertions.

606 ► **Lemma 10.** Each of the rules below is sound (as per Definition 2), where the statements
607 are decorated with the thread identifier of the executing thread.

$$608 \quad \text{ZWR} \frac{m \neq n}{\{\mathbb{0}_x m\} y :=_t^{[R]} n \{\mathbb{0}_x m\}} \quad \text{DVPRES} \frac{}{\{m \ll_x n\} r \leftarrow_t^{[A]} y \{m \ll_x n\}}$$

$$609 \quad \text{1INTRO} \frac{i \neq m}{\{x_{\text{Init}} = i \wedge \mathbb{0}_x m\} x :=_t^{[R]} m \{\mathbb{1}_x m\}} \quad \text{ENCWR} \frac{}{\{\text{true}\} x :=_t^{[R]} m \{x \stackrel{\text{enc}}{=} m\}}$$

$$610 \quad \text{ENCRD} \frac{}{\{\text{true}\} r \leftarrow_t^{[A]} x \{x \stackrel{\text{enc}}{=} r\}} \quad \text{EPO} \frac{}{\{x \stackrel{\text{enc}}{=} m\} r \leftarrow_t^{[A]} x \{r \neq m \Rightarrow m \prec_x r\}}$$

$$611 \quad \text{DVINTRO} \frac{i \neq n}{\{x_{\text{Init}} = i \wedge \mathbb{0}_x n \wedge \mathbb{1}_x m \wedge x \stackrel{\text{enc}}{=} m\} x :=_t^{[R]} n \{m \ll_x n\}}$$

$$612 \quad \text{1PRESR} \frac{}{\{\mathbb{1}_x m\} r \leftarrow_t^{[A]} y \{\mathbb{1}_x m\}} \quad \text{PORD} \frac{}{\{m \prec_x n\} C \{m \prec_x n\}}$$

614 **Proof.** In Isabelle. ◀

615 ► **Theorem 11.** The proof outline for RRC2 in Figure 9 is valid.

616 **Proof.** This proof has been mechanised in Isabelle. Once again, we describe the proof outline
617 to give an overview of how our proofs are used. For local correctness we have the following.

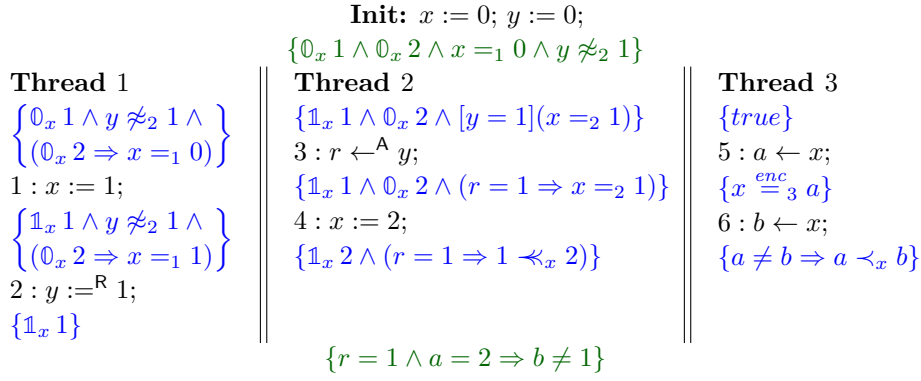
- 618 ■ The initialisation clearly satisfies the precondition of the program, and this implies the
619 precondition of thread 1. The precondition of thread 2 is trivial.
- 620 ■ Next we consider the postcondition of line 1. The first conjunct holds by ZWR, the
621 second conjunct holds by 1INTRO and the third by rule ENCWR.
- 622 ■ The postcondition of line 2 holds by rule DVINTRO.
- 623 ■ In thread 2, the postcondition of line 3 holds by rule ENCRD, and the postcondition of
624 line 4 holds by rule EPO.

625 Next we check interference freedom.

- 626 ■ The precondition of line 1 is stable with respect to lines 3 and 4 by ZWR.
- 627 ■ Next consider the precondition of line 2. The first and second conjuncts are stable with
628 respect to lines 3 and 4 by ZWR and 1PRESR, respectively. The third conjunct is trivially
629 preserved (see Isabelle).
- 630 ■ The postcondition of line 2 holds by DVPRES.
- 631 ■ The precondition of line 3 is trivial and the postcondition of line 3 holds by PORD.

632 ◀

633 Correctness of RRC and RRC3 is established by the following theorem.



■ **Figure 11** Proof outline for RRC3, where $x, y \in Var_G$ and $a, b \in Var_L$

634 ► **Theorem 12.** *The proof outlines for RRC and RRC3 in Figure 10 and Figure 11, respectively*
635 *are valid.*

636 **Proof.** In Isabelle. ◀

637 For RRC (Figure 10), the precondition of line 4 records the fact that thread 3 has
638 encountered a (whatever the value of a may be). Moreover, it guarantees that there is
639 at most one write of x with values 1 and 2. The first conjunct (i.e., $x \stackrel{enc}{=}_3 a$) allows us
640 to conclude that after x is read at line 4, if a and b are different, then the value for a is
641 possibly ordered before the value for b . The second and third conditions are used to establish
642 the postconditions $1_x 1$ and $1_x 2$. This argument also applies to the assertions in thread 4.
643 Finally, we show that the postcondition of the program holds as follows, where we assume
644 *post* is the conjunction of the postcondition of each thread.

$$\begin{array}{ll}
645 & \text{post} \Rightarrow (a = 1 \wedge b = 2 \wedge c = 2 \Rightarrow d \neq 1) \\
646 & \iff \text{post} \wedge a = 1 \wedge b = 2 \wedge c = 2 \wedge d = 1 \Rightarrow \text{false} \quad (\text{logic}) \\
647 & \iff 1_x 1 \wedge 1_x 2 \wedge 1 \prec_x 2 \wedge 2 \prec_x 1 \Rightarrow \text{false} \quad (\text{logic}) \\
648 & \iff 1 \prec_x 2 \wedge 2 \prec_x 1 \Rightarrow \text{false} \quad (6) \\
649 & \iff \text{true} \quad (7) \\
650 &
\end{array}$$

651 The calculation above has been verified with Isabelle, but we recall the proof here as it
652 provides insight into the interactions between different value assertions.

653 RRC3 (Figure 11) combines message passing on y with RRC on x . Namely, knowledge
654 of $x := 1$ in thread 1 is transferred to thread 2 using a release-acquire synchronisation on
655 y . Thus, if thread 2 reads 1 for y it must also have encountered 1 for x . Thus, if $r = 1$,
656 then the write on line 4 must have happened *after* the write on line 1. This means that it
657 should be impossible for thread 3 to read 2 for x (at line 5) then read 1 for x (at line 6).
658 Unlike message passing, in RRC3, the “data” variable x is updated both before and after
659 synchronisation. Thus, the assertions on definite values (e.g., $x =_1 1$) become conditional
660 on whether line 4 has already been executed. In particular, the antecedent $0_x 2$ allows us
661 to assume that line 4 has not yet been executed. As with RRC, we must separately prove
662 that the conjunction of the postconditions of the threads implies the postcondition of the
663 program. This proof is mechanised in Isabelle, and is elided here.

Init: $flag_1 := 0; flag_2 := 0; turn := 0 \wedge after_1 := false; after_2 := false$

Thread 1

$$\left\{ \begin{array}{l} \neg after_1 \wedge flag_1 =_1 0 \wedge turn \not\approx_2 2 \wedge (C_{turn}^0 \vee [turn = 1](flag_2 =_1 1)) \\ \wedge (after_2 \Rightarrow C_{turn}^1 \wedge [turn = 1](flag_2 =_1 1)) \end{array} \right\}$$

1: $flag_1 := 1;$
 $\left\{ \neg after_1 \wedge flag_1 =_1 1 \wedge turn \not\approx_2 2 \wedge (after_2 \Rightarrow C_{turn}^1 \wedge [turn = 1](flag_2 =_1 1)) \right\}$

2: $\langle turn.swap(2)^{RA}; after_1 := true \rangle$
 $\left\{ after_1 \wedge (after_2 \wedge (flag_2 \approx_1 0 \vee turn \approx_1 1) \Rightarrow turn =_2 1) \right\}$

do

3: $r_1 \leftarrow^A flag_2$
 $\left\{ after_1 \wedge (after_2 \wedge (r_1 = 0 \vee turn \approx_1 1 \vee flag_2 \approx_1 0) \Rightarrow turn =_2 1) \right\}$

4: $r_2 \leftarrow turn$
 $\left\{ after_1 \wedge (after_2 \wedge (r_1 = 0 \vee r_2 = 1 \vee turn \approx_1 1 \vee flag_2 \approx_1 0) \Rightarrow turn =_2 1) \right\}$

5: **until** $(r_1 = 0 \vee r_2 = 1)$
 $\left\{ after_1 \wedge (after_2 \Rightarrow turn =_2 1) \right\}$

6: Critical section ;

7: $\langle flag_1 :=^R 0; after_1 := false \rangle$

■ **Figure 12** Peterson’s algorithm (adapted from [36]) and its proof outline. **Thread 2** (not shown) is symmetric.

664 6 Case study: Peterson’s algorithm

665 We turn to our final case study, the verification of the mutual exclusion property of a
 666 version of Peterson’s algorithm. The complexity of this case study is much greater than
 667 our earlier examples. This program contains a loop, features a careful mixture of relaxed
 668 and release/acquire operations to the same variable, and an RMW operation whose precise
 669 semantics is critical to the correctness of the algorithm.

670 Our version of Peterson’s algorithm⁴, presented in Figure 12 is a mutual exclusion
 671 algorithm for two threads implemented for C11 using release-acquire annotations [36]. As
 672 with the original algorithm, variable $flag_i$, for $i \in \{1, 2\}$ is used to indicate whether thread i
 673 intends to enter its critical section. In this version of the algorithm, we let $flag_i$ range over
 674 $\{0, 1\}$, where 0 is used for the boolean value “false”, and 1 is used for the boolean value “true”.
 675 The shared variable $turn$ is used to cause a thread to “give way” when both threads intend
 676 to enter their critical sections at the same time. Our verification uses auxiliary variables
 677 $after_i$ for each thread i (as does the proof for a sequentially consistent setting in [7]), the
 678 purpose of which we describe below.

679 We describe the algorithm for thread 1; the other thread is symmetric. For now, we
 680 ignore the assertions. The flag variable is set to 1 (line 1) using a relaxed write (which cannot
 681 induce any synchronisation), but is set to 0 (line 7) using a release annotation. The intention
 682 of the latter is to synchronise this write (of 0 to $flag_1$) with the read of $flag_1$ at line 3 in
 683 thread 2. The value of $turn$ is set using a **swap** command. The **swap** is implemented using
 684 an C11 RMW operation that has both the release and acquire annotations. When the **swap**
 685 is executed, as part of the same transition, the auxiliary variable $after_1$ is also set, indicating
 686 that thread 1 is ready to enter the busy wait loop beginning at line 3, and then to enter the
 687 critical section.

688 The busy wait loop forces thread 0 to wait until either $flag_2$ is 0 (indicating that thread
 689 2 is not trying to enter the critical section) or $turn = 1$ (indicating that it is thread 1’s turn

⁴ For simplicity our version of the algorithm does not have an outermost loop.

690 to enter the critical section). Note that the read of $turn$ within the guard of the busy wait
691 loop (line 5) is relaxed.

692 We turn now to the proof that this version of Peterson's algorithm has the mutual
693 exclusion property. We prove mutual exclusion in two steps. First, we show that the given
694 proof outline is valid, and second, that the conjunction of the precondition of thread 1's
695 critical section (line 6) and thread 2's must be false. Therefore, the two threads cannot
696 simultaneously be in their critical sections.

697 We deal with the second step first by showing that the formula below is *false*:

$$698 \text{after}_1 \wedge (\text{after}_2 \Rightarrow \text{turn} =_2 1) \wedge \text{after}_2 \wedge (\text{after}_1 \Rightarrow \text{turn} =_1 2)$$

700 It is easy to see that this implies $\text{turn} =_1 2 \wedge \text{turn} =_2 1$. However, by (5) this situation is
701 impossible.

702 The first step is more elaborate and we only describe certain aspects. The precondition of
703 line 3 is also an invariant of the busy wait loop. This assertion ensures that if thread 1 is able
704 to exit the busy wait loop, then the precondition of the critical section will be satisfied. Note
705 that thread 1 exits the loop if it reads 0 from $flag_2$ (which is only possible when $flag_2 \approx_1 0$)
706 or it reads 1 from $turn$ (which is only possible when $turn \approx_1 1$). The invariant states that if
707 one of these conditions holds in a state where thread 2 is waiting to enter the critical section
708 (that is, after_2), we can conclude $\text{turn} =_2 1$ as required.

709 Proving that the precondition of line 3 is satisfied in the post-state of line 2 requires using
710 a feature of our assertion language, closely related to the semantics of RMW operations,
711 that we now introduce. Recall from the UPDATE rule in Figure 5 that whenever a write w
712 is read-from by an RMW operation, w becomes *covered*, so that no later write (or RMW)
713 operation can be inserted between w and the RMW. This feature of C11 is critical to the
714 correctness of Peterson's algorithm. Observe that the $turn$ variable is only modified by RMW
715 operations, and therefore every write to $turn$ is covered, except the last. To formally state
716 this, we need the third *occurrence* assertion \mathbf{C}_x^n , defined as follows.

$$717 \mathbf{C}_x^n = \lambda\sigma. \forall w \in \sigma.\text{writes}_x. w \notin \sigma.\text{covered} \Rightarrow \text{wrval}(w) = n \wedge w = \text{last}(W, x)$$

719 So \mathbf{C}_x^n means that every write to x except the last is covered and the value written by that
720 last write is n .

721 We use the following lemma on covered.

► **Lemma 13.**

$$722 \text{CVD-UPD} \frac{}{\{\mathbf{C}_x^n\} \text{upd}^{\text{RA}}(x, m, l) \{m = n \wedge \mathbf{C}_x^l\}} \quad \text{CVD-WR} \frac{x \neq y}{\{\mathbf{C}_x^n\} y :=^{\text{[R]}} m \{\mathbf{C}_x^n\}}$$

$$723 \text{CVD-RD} \frac{}{\{\mathbf{C}_x^n\} r \leftarrow^{\text{[A]}} y \{\mathbf{C}_x^n\}} \quad \text{CVD-DOBS} \frac{}{\{\mathbf{C}_x^n\} \text{upd}^{\text{RA}}(x, m, n) \{x =_t n\}}$$

725 Rule CVD-UPD states that if \mathbf{C}_x^n holds in the pre-state, then after executing $\text{upd}^{\text{RA}}(x, m, l)$,
726 we must have that $m = n$ (since the only value available for the RMW to read is n), and
727 furthermore we obtain a new covered predicate \mathbf{C}_x^l . Thus, it is possible to maintain a covered
728 predicate in a program (with possibly different return values) by ensuring each modification to
729 the covered variable is via an RMW. This is a property that is true of Peterson's algorithm
730 as given in Figure 12. Rules CVD-WR and CVD-RD give preservation properties for the
731 covered assertion for a read and a write, respectively. Finally, CVD-DOBS is used to establish
732 a definite observation of a covered assertion after an update command.

```

lemma d_obs_Wr_set:
  assumes "wfs  $\sigma$ "
    and "wr_val a = Some n"
    and "avar a = x"
    and "[x =t m]  $\sigma$ "
    and "step t a  $\sigma$   $\sigma'$ "
  shows "[x =t n]  $\sigma'$ "

corollary d_obs_WrX_set:
  "wfs  $\sigma \implies [x =t m] \sigma \implies \sigma [x := n]_t \sigma' \implies [x =t n] \sigma'$ "

corollary d_obs_WrR_set :
  "wfs  $\sigma \implies [x =t m] \sigma \implies \sigma [x :=^R n]_t \sigma' \implies [x =t n] \sigma'$ "

corollary d_obs_RMW_set :
  "wfs  $\sigma \implies [x =t m] \sigma \implies \sigma \text{RMW}[x,w,n]_t \sigma' \implies [x =t n] \sigma'$ "

```

■ **Figure 13** Isabelle encoding of basic axioms over C11 assertions

733 The precondition of line 2 asserts that if thread 2 is ready to enter the critical section
 734 (that is, $after_2$) then the RMW to be executed at line 2 must read from the last write which
 735 has value 1 (that is, C_{turn}^1) and when this RMW occurs then thread 1 will definitely see
 736 $flag_2$ set (that is, $[turn = 1](flag_2 =_1 1)$). This is enough to show that if $after_2$ then in the
 737 post-state of the RMW, $flag_2 \not\approx_1 0$ which is sufficient to prove the postcondition of line 2.

738 Of course, the sequential reasoning above must be combined with an interference freedom
 739 check, which is supported by a set of basic lemmas describing how C_x^n is updated. This leads
 740 to the following theorem, which establishes validity of the proof outline.

741 ► **Theorem 14.** *The proof outline of Peterson's algorithm (Figure 12) is valid.*

742 **Proof.** In Isabelle. ◀

743 We note that Peterson's algorithm represents a challenge in deductive verification. Unlike
 744 the litmus tests presented above, there is sufficient complexity in the algorithm and the
 745 resulting proof outline so that pen-and-paper proofs cannot be trusted. Using our mech-
 746 anisation, we explored several variations of the proof outline in Figure 12, and discovered
 747 simplifications to our original pen-and-paper proofs.

748 **7 Mechanisation**

749 As already mentioned, the operational semantics as well as all lemmas and theorems presented
 750 in this paper have been mechanised in Isabelle. In this section, we discuss our mechanisation
 751 effort.

752 To prove the lemmas about basic assertions, we typically prove a more general result
 753 relating to reads and writes, which are then specialised so that they can be used in the
 754 verification of the algorithms. For example, we first prove the lemma in Figure 13, which
 755 describes changes to definite values and applies to any writing transition. This is then
 756 specialised to the corollaries on the right, which are easier for Isabelle to find when performing
 757 the verification of the proof outlines.

758 The generic lemmas require some amount of interactive work. However, once verified, it is
 759 straightforward to use them to prove the corollaries. For example, corollary `d_obs_WrX_set` in
 760 Figure 13 is verified with the command “`by (metis WrX_def avar.simps(2) d_obs_Wr_set`
 761 `wr_val.simps(1))`”, which is found automatically by Isabelle’s built in `sledgehammer`
 762 tool [10].

763 Such lemmas and corollaries are in turn used in the proofs of programs. First the program
 764 state (i.e., Σ_{C11}) is encoded as a `record` type with a special variable that models the C11
 765 state. The programs themselves are encoded as a relation over these records with program
 766 counters modelling control flow. This allows the proof outlines to be encoded as predicates
 767 mapping program counters to the assertions at that control point. We then verify a set
 768 of lemmas that guarantee local correctness and interference freedom, where we decompose
 769 proofs and apply case analysis over the individual program steps (e.g., reads, writes for
 770 each thread). Once a proof has been decomposed, `sledgehammer` is able to find the relevant
 771 corollaries (e.g., those in Figure 13) to discharge proofs automatically.

772 8 Related Work

773 The semantics and verification of programs running on weak memory models has recently
 774 received a lot of attention. Lahav [21] gives a brief survey for C11.

775 Our timestamp based operational semantics is motivated by ideas in [13] and is similar
 776 to the semantics of Kaiser et al. [16, 17]. We note there are differences in coverage of the
 777 memory models in [13, 16, 17]. Dolan et al. [13] cover a sequentially consistent (SC) and
 778 relaxed accesses for OCAML, where the SC operations behave like Java volatiles. Kaiser et
 779 al [16] covers non-atomics and release-acquire, while Kang et al. [17] support a much larger
 780 fragment of C11, including so-called load-buffering cycles.

781 Abdulla et al. have shown the reachability problem for release-acquire to be *undecidable* [2].
 782 A number of works target *model checking* for weak memory, e.g., by explicitly encoding
 783 architectural structures leading to weak behaviour, like store buffers [33, 5]. Ponce de León
 784 et al. [30, 14] have developed a bounded model checker for weak memory models, taking the
 785 axiomatic description of a memory model as input. (Bounded) model checkers for specific
 786 weak memory models are furthermore the tools CBMC [6] (for TSO), NIDHUGG [1] (for TSO
 787 and PSO), RCMC [18] (for C11) and GENMC [19] (again, parametric in memory model).

788 A (non-automatic) reasoning technique for proving invariants – parameterised by a
 789 weak memory model – has been proposed by Alglave and Cousot [4]. They propose a new
 790 semantics, different from an operational one without any coherence order (or modification
 791 order) constraining the order of writes to memory. Their assertions contain so-called pythia
 792 variables to uniquely identify values of read events, and require a separate communication
 793 proof (differentiating their method from standard Owicki-Gries reasoning). They say “In
 794 addition to the initialisation, sequential, and non-interference proof, the main difference
 795 with Owicki and Gries [27] (and Lamport 1977) is the use of pythia variables and the
 796 read-from relation in assertions and the communication proof showing that reads-from is
 797 well-formed.” [4]. Our method in contrast only requires the initialisation, sequential, and
 798 non-interference proofs as with the original technique.

799 Another manual method for the RC11 memory model has been developed by Doherty
 800 et al. [12], who cover the message passing example and Peterson’s algorithm. Our work is
 801 inspired by this existing work, however, there are several differences. They use a classical
 802 model of the C11 state (expressed in terms of a set of relations, e.g., reads-from, sequenced-
 803 before etc), develop assertions over these relations and a small proof calculus for these

804 assertions. However, their methods are at a lower level of abstraction than the techniques
 805 presented in this paper since the assertions are stated in terms of individual relations that
 806 make up each state. Thus, it is not possible to directly develop a Hoare logic for their
 807 assertions and mechanisation itself is more difficult.

808 Also close to our work is that of Lahav and Vafeiadis [22] who also develop an Owicki-Gries
 809 style proof calculus. We consider all their examples except RCU — our logic can handle the
 810 RCU example, but this proof has thus far not been mechanised. Moreover, we include several
 811 other case studies such as litmus tests that combine read-read coherence with message passing
 812 and the non-trivial Peterson’s algorithm. There are several additional differences to note. (1)
 813 Lahav and Vafeiadis’ proof calculus is developed in the absence of an operational semantics,
 814 and hence, their definition of a valid Hoare triple is non standard (see [22, Definition 9]). A
 815 consequence of this is that they must be careful about the introduction of auxiliary variables,
 816 resorting to the more restricted notion of a *ghost* variable. In contrast, we use traditional
 817 auxiliary variables — an auxiliary variable must not affect the control flow of a program
 818 nor be assigned to any program variable. Note however, that to simplify the presentation, we
 819 use auxiliary variables in a more restricted manner (see Section 3). (2) They do not handle
 820 relaxed accesses — as stated in their conclusion: “While OGRA’s non-interference condition
 821 appears to be restrictive, it is unsound for weaker memory models, such as C11’s relaxed
 822 accesses . . .”. (3) They do not provide a mechanisation.

823 A frequently employed starting point for program logic is separation logic, for which
 824 a number of extensions to weak memory exist (GPS [34], RSL [16]). Svendsen et al. [32]
 825 propose a separation logic based on the promising semantics of Kang et al. [17]. The principle
 826 of ownership transfer used therein naturally fits to message passing using release acquire.
 827 Prover support for such separation logic based proofs — like ours with Isabelle — has been
 828 developed for the Iris proof system [16]. Tool support has also been developed by Summers
 829 and Müller [31], where the RSL logic has been encoded in the Viper tool, offering a level
 830 of proof automation. Their encoding is proved sound and complete with respect to RSL.
 831 However, such efforts do not provide a clear link between C11 semantics and traditional
 832 reasoning using Hoare logics.

833 9 Conclusion

834 In this paper, we have introduced an assertion language for C11 RAR which allows to re-use
 835 the entire Owicki-Gries proof calculus except for the axiom of assignment. The assertion
 836 language is based on an operational semantics for C11 RAR which we have shown to be sound
 837 wrt. standard axiomatic semantics. We have exemplified reasoning on a number of standard
 838 C11 RAR litmus tests as well as a C11 RAR annotated version of Peterson’s algorithm. All
 839 proofs ranging are mechanised within Isabelle — this includes soundness of the basic axioms
 840 for weak memory reads, writes and updates, and validity of proof outlines for the examples
 841 presented.

842 We are currently integrating this work⁵ into the standard Owicki-Gries library that
 843 is included in the Isabelle distribution [26]. As future work, we aim to tackle fragments
 844 of C11 larger than C11 RAR, e.g., fragments that allow the load buffering example to
 845 terminate with postcondition $r1 = 1 \wedge r2 = 1$ [9, 17], SC annotations [20], as well as
 846 release sequences and fences [8]. Extending our operational semantics to handle the final

⁵ A set of preliminary results is available <https://www.dropbox.com/sh/4yr2w7792qwyw09/AACsWUXtZbK3PvyfJkqjyDYa> within the file OG-Isabelle.zip.

847 two features is straightforward, but is not considered in this paper as it complicates the
848 semantics and detracts from our main contribution, i.e., a simple extension to Hoare logic to
849 enable reasoning about C11 programs. Hoare-style reasoning that incorporates the other two
850 features is currently being investigated.

851 ——— References ———

- 852 1 Parosh Aziz Abdulla, Stavros Aronis, Mohamed Faouzi Atig, Bengt Jonsson, Carl Leonardsson,
853 and Konstantinos Sagonas. Stateless model checking for TSO and PSO. *Acta Inf.*, 54(8):789–
854 818, 2017.
- 855 2 Parosh Aziz Abdulla, Jatin Arora, Mohamed Faouzi Atig, and Shankara Narayanan Krishna.
856 Verification of programs under the release-acquire semantics. In McKinley and Fisher [25],
857 pages 1117–1132.
- 858 3 J. Alglave, L. Maranget, and M. Tautschnig. Herding cats: Modelling, simulation, testing,
859 and data mining for weak memory. *ACM Trans. Program. Lang. Syst.*, 36(2):7:1–7:74, 2014.
- 860 4 Jade Alglave and Patrick Cousot. Ogre and Pythia: an invariance proof method for weak
861 consistency models. In Castagna and Gordon [11], pages 3–18.
- 862 5 Jade Alglave, Daniel Kroening, Vincent Nimal, and Michael Tautschnig. Software verification
863 for weak memory via program transformation. In M. Felleisen and P. Gardner, editors, *ESOP*,
864 volume 7792 of *LNCS*, pages 512–532. Springer, 2013.
- 865 6 Jade Alglave, Daniel Kroening, and Michael Tautschnig. Partial orders for efficient bounded
866 model checking of concurrent software. In Natasha Sharygina and Helmut Veith, editors, *CAV*,
867 volume 8044 of *LNCS*, pages 141–157. Springer, 2013.
- 868 7 Krzysztof R. Apt, Frank S. de Boer, and Ernst-Rüdiger Olderog. *Verification of Sequential
869 and Concurrent Programs*. Texts in Computer Science. Springer, 2009.
- 870 8 M. Batty, S. Owens, S. Sarkar, P. Sewell, and T. Weber. Mathematizing C++ concurrency.
871 In T. Ball and M. Sagiv, editors, *POPL*, pages 55–66. ACM, 2011.
- 872 9 Mark Batty, Alastair F. Donaldson, and John Wickerson. Overhauling SC atomics in C11 and
873 OpenCL. In *POPL*, pages 634–648. ACM, 2016.
- 874 10 Sascha Böhme and Tobias Nipkow. Sledgehammer: Judgement day. In *IJCAR*, volume 6173
875 of *Lecture Notes in Computer Science*, pages 107–121. Springer, 2010.
- 876 11 Giuseppe Castagna and Andrew D. Gordon, editors. *POPL*. ACM, 2017.
- 877 12 Simon Doherty, Brijesh Dongol, Heike Wehrheim, and John Derrick. Verifying C11 programs
878 operationally. In Jeffrey K. Hollingsworth and Idit Keidar, editors, *PPoPP*, pages 355–365.
879 ACM, 2019.
- 880 13 Stephen Dolan, KC Sivaramakrishnan, and Anil Madhavapeddy. Bounding data races in space
881 and time. In *PLDI*, PLDI 2018, pages 242–255, New York, NY, USA, 2018. ACM.
- 882 14 Natalia Gavrilenko, Hern’an Ponce de Le’on, Florian Furbach, Keijo Heljanko, and Roland
883 Meyer. BMC for weak memory models: Relation analysis for compact SMT encodings. In
884 Isil Dillig and Serdar Tasiran, editors, *CAV*, volume 11561 of *LNCS*, pages 355–365. Springer,
885 2019. doi:10.1007/978-3-030-25540-4_19.
- 886 15 C. A. R. Hoare. An axiomatic basis for computer programming. *Commun. ACM*, 12(10):576–
887 580, 1969.
- 888 16 Jan-Oliver Kaiser, Hoang-Hai Dang, Derek Dreyer, Ori Lahav, and Viktor Vafeiadis. Strong
889 logic for weak memory: Reasoning about release-acquire consistency in Iris. In Peter Müller,
890 editor, *ECOOP*, volume 74 of *LIPICs*, pages 17:1–17:29. Schloss Dagstuhl - Leibniz-Zentrum
891 fuer Informatik, 2017.
- 892 17 Jeehoon Kang, Chung-Kil Hur, Ori Lahav, Viktor Vafeiadis, and Derek Dreyer. A promising
893 semantics for relaxed-memory concurrency. In Castagna and Gordon [11], pages 175–189.
- 894 18 Michalis Kokologiannakis, Ori Lahav, Konstantinos Sagonas, and Viktor Vafeiadis. Effective
895 stateless model checking for C/C++ concurrency. *PACMPL*, 2(POPL):17:1–17:32, 2018.

- 896 **19** Michalis Kokologiannakis, Azalea Raad, and Viktor Vafeiadis. Model checking for weakly
897 consistent libraries. In McKinley and Fisher [25], pages 96–110.
- 898 **20** O. Lahav, V. Vafeiadis, J. Kang, C.-K. Hur, and D. Dreyer. Repairing sequential consistency
899 in C/C++11. In *PLDI*, pages 618–632. ACM, 2017.
- 900 **21** Ori Lahav. Verification under causally consistent shared memory. *SIGLOG News*, 6(2):43–56,
901 2019.
- 902 **22** Ori Lahav and Viktor Vafeiadis. Owicki-Gries reasoning for weak memory models. In
903 Magnús M. Halldórsson, Kazuo Iwama, Naoki Kobayashi, and Bettina Speckmann, editors,
904 *ICALP*, volume 9135 of *LNCS*, pages 311–323. Springer, 2015.
- 905 **23** L. Lamport. How to make a multiprocessor computer that correctly executes multiprocess
906 programs. *IEEE Trans. Computers*, 28(9):690–691, 1979.
- 907 **24** Nancy A Lynch. *Distributed algorithms*. Elsevier, 1996.
- 908 **25** Kathryn S. McKinley and Kathleen Fisher, editors. *PLDI*. ACM, 2019.
- 909 **26** Tobias Nipkow and Leonor Prensa Nieto. Owicki/Gries in Isabelle/HOL. In *FASE*, volume
910 1577 of *Lecture Notes in Computer Science*, pages 188–203. Springer, 1999.
- 911 **27** Susan S. Owicki and David Gries. An axiomatic proof technique for parallel programs I. *Acta*
912 *Inf.*, 6:319–340, 1976.
- 913 **28** Lawrence C. Paulson. *Isabelle - A Generic Theorem Prover (with a contribution by T. Nipkow)*,
914 volume 828 of *LNCS*. Springer, 1994.
- 915 **29** Anton Podkopaev, Ilya Sergey, and Aleksandar Nanevski. Operational aspects of C/C++
916 concurrency. *CoRR*, abs/1606.01400, 2016. [arXiv:1606.01400](https://arxiv.org/abs/1606.01400).
- 917 **30** Hernán Ponce de León, Florian Furbach, Keijo Heljanko, and Roland Meyer. BMC with
918 memory models as modules. In Nikolaj Bjørner and Arie Gurfinkel, editors, *FMCAD*, pages
919 1–9. IEEE, 2018.
- 920 **31** Alexander J. Summers and Peter Müller. Automating deductive verification for weak-memory
921 programs. In Dirk Beyer and Marieke Huisman, editors, *TACAS*, volume 10805 of *LNCS*,
922 pages 190–209. Springer, 2018.
- 923 **32** Kasper Svendsen, Jean Pichon-Pharabod, Marko Doko, Ori Lahav, and Viktor Vafeiadis. A
924 separation logic for a promising semantics. In Amal Ahmed, editor, *ESOP*, volume 10801 of
925 *LNCS*, pages 357–384. Springer, 2018.
- 926 **33** Oleg Travkin, Annika Mütze, and Heike Wehrheim. SPIN as a linearizability checker under
927 weak memory models. In Valeria Bertacco and Axel Legay, editors, *HVC*, volume 8244 of
928 *LNCS*, pages 311–326. Springer, 2013.
- 929 **34** Aaron Turon, Viktor Vafeiadis, and Derek Dreyer. GPS: navigating weak memory with ghosts,
930 protocols, and separation. In Andrew P. Black and Todd D. Millstein, editors, *OOPSLA*,
931 pages 691–707. ACM, 2014.
- 932 **35** John Wickerson, Mark Batty, Tyler Sorensen, and George A. Constantinides. Automatically
933 comparing memory consistency models. In Castagna and Gordon [11], pages 190–204. URL:
934 <http://dl.acm.org/citation.cfm?id=3009838>.
- 935 **36** A. Williams. [https://www.justsoftwaresolutions.co.uk/threading/petersons_lock_](https://www.justsoftwaresolutions.co.uk/threading/petersons_lock_with_C++0x_atomics.html)
936 [with_C++0x_atomics.html](https://www.justsoftwaresolutions.co.uk/threading/petersons_lock_with_C++0x_atomics.html), 2018. Accessed: 2018-06-20.

A Correctness of the Operational Semantics

In this section, we prove that the operational semantics presented in this paper is sound w.r.t. a standard version of the C11 RAR axiomatic semantics.

We base this proof on the operational semantics for C11 RAR developed by Doherty et al. [12] which they have proved to be equivalent to a suitable fragment of the C11 RAR axiomatic semantics of [9]. In what follows, we refer to the semantics of [12] as the *Explicit semantics* as the states of its memory semantics are simply consistent C11 RAR executions. We refer to the semantics presented in this paper as the *View semantics*.

In Section A.1, we describe the operational semantics of [12]. In Section A.2, we present a simulation relation from the View memory semantics to the Explicit memory semantics. This is sufficient to show that all sequences of operations accepted by the View semantics are also accepted by the Explicit semantics.

In what follows, we refer to states of the Explicit semantics using the variables E, E' and states of the View semantics as V, V' .

A.1 The Explicit Memory Semantics

We refer the reader to [12] for a full discussion of the Explicit semantics. In this paper, we present only the semantics that relates to memory operations, and we do so only briefly.

States of the Explicit semantics are simply C11 RAR executions, of the form $E = (X, \text{sb}, \text{rf}, \text{mo})$. Where $\text{sb}, \text{rf}, \text{mo}$ are relations on the set of operations X with their usual meanings. That is $\text{sb} \subseteq X \times X$ is the *sequenced-before* relation; $\text{rf} \subseteq W \times R \cap X \times X$ is the *reads-from* relation; and $\text{mo} \subseteq W \times W \cap X \times X$ is the *modification order*.

In an Explicit state $E = (X, \text{sb}, \text{rf}, \text{mo})$ we let $X \subseteq \text{Act} \times \mathbb{Q} \times T$, where Act is the set of operations and T is the set of threads. In the original presentation, [12] we specify $X \subseteq \text{Act} \times G \times T$ where G is a set of tags that are not further specified, which serves to distinguish repeated occurrences of the same operation. Here, we let $G = \mathbb{Q}$, for uniformity with the view semantics.

The Explicit semantics uses additional *synchronized-with* (denoted sw) and *happens-before* (denoted hb) relations, defined as follows:

$$\text{sw} = \text{rf} \cap (W_R \times R_A) \tag{8}$$

$$\text{hb} = (\text{sb} \cup \text{sw})^+ \tag{9}$$

In the Explicit semantics, all variables are initialised by a special *initialising thread* $0 \in T$. Define the set of *initialising writes* to be $IWr = \{w \in W \mid \text{tid}(w) = 0\}$. The initial states of our operational model are those of the form $E_0 = ((I, \emptyset), \emptyset, \emptyset)$ where $I \subseteq IWr$, and for each variable x , there is exactly one write $w \in I$ such that $\text{var}(w) = x$. For a state $E = ((X, _), _, _)$, let $I_E = X \cap IWr$.

The relation $\text{fr} = (\text{rf}^{-1}; \text{mo}) \setminus Id$ (where $;$ is relational composition) is the “from-read” relation, that relates each read to all writes that are mo -after the write the read has read from. We must subtract Id (identity) edges from $\text{rf}^{-1}; \text{mo}$ to cope with update events, which have the potential to induce reflexivity in fr [9, 20].

In addition, the semantics uses the *extended coherence order* [20], denoted eco , which fixes the order of reads and writes to each variable. Formally we define:

$$\text{eco} = (\text{fr} \cup \text{mo} \cup \text{rf})^+ \tag{10}$$

$$\begin{array}{c}
\text{READ} \frac{a \in \{rd(x, n), rd^A(x, n)\} \quad wrval(w) = n}{w \in E.OW(t, x) \quad rf' = rf \cup \{(w, e)\} \quad mo' = mo} \\
\hline
((X, sb), rf, mo) \xrightarrow{e} ((X, sb) + e, rf', mo') \\
\\
\text{WRITE} \frac{a \in \{wr(x, n), wr^R(x, n)\} \quad w \in E.OW(t, x) \setminus E.covered}{rf' = rf \quad mo' = mo[w, e]} \\
\hline
((X, sb), rf, mo) \xrightarrow{e} ((X, sb) + e, rf', mo') \\
\\
\text{RMW} \frac{a = upd^{RA}(x, m, n) \quad w \in E.OW(t, x) \setminus E.covered}{wrval(w) = m \quad rf' = rf \cup \{(w, e)\} \quad mo' = mo[w, e]} \\
\hline
((D, sb), rf, mo) \xrightarrow{e} ((X, sb) + e, rf', mo')
\end{array}$$

■ **Figure 14** Event semantics assuming $E = ((X, sb), rf, mo)$, $e = (g, a, t)$ and $g \notin tags(X)$

981 The set of *encountered writes* are the writes that thread t is aware of (either directly or
982 indirectly) in state $E = ((X, sb), rf, mo)$, and are given by:

$$\begin{array}{c}
983 \quad E.EW = \{w \in W \cap D \mid \exists e \in D. tid(e) = t \wedge \\
984 \quad \quad \quad (w, e) \in eco^?; hb^?\} \cup I_E
\end{array}$$

985 where $R^?$ is the reflexive closure of relation R . Thus, for each $w \in E.EW$, there must exist
986 an event e of thread t such that w is either *eco*- or *hb*- or *eco;hb*-prior to e .

987 From these, we determine the *observable writes*, which are the writes that thread t can
988 observe in its next read. These are defined as:

$$\begin{array}{c}
989 \quad E.OW(t, x) = \{w \in W \cap E.X \mid loc(w) = x \wedge \forall w' \in E.EW(t) \wedge (w, w') \notin E.mo\} \\
990
\end{array}$$

991 Thus, observable writes are not succeeded by any encountered write in modification order,
992 i.e., the thread has not seen another write overwriting the value being read.

993 Finally, to guarantee *atomicity* of the update events, there cannot be any write operations
994 (in modification order) between the write that an update reads from and the write of the
995 update itself. We therefore define the set of *covered writes* as follows:

$$\begin{array}{c}
996 \quad E.covered = \{w \in W \cap E.X \mid \exists u \in U. (w, u) \in rf\} \\
997
\end{array}$$

998 The transition relation of the Explicit semantics is given in Figure 14.

999 ► **Lemma 15** (Invariants of the Explicit Semantics). *If $E = (X, sb, rf, mo)$ and E is reachable
1000 from an initial state via a sequence of transitions of the Explicit semantics, then*

1001 ■ *mo totally orders the writes to each variable x . That is, for all $w, w' \in X \cap W$, such that
1002 $loc(w) = loc(w')$.*

$$\begin{array}{c}
1003 \quad (w, w') \in mo \vee (w', w) \in mo \\
1004
\end{array} \tag{11}$$

1005 ■ *For each variable x , there is at least one write. That is, for each x*

$$\begin{array}{c}
1006 \quad \exists w \in X \cap W. loc(w) = x \\
1007
\end{array} \tag{12}$$

1008 **Proof.** These are simple inductive invariants of the semantics. ◀

1009 A.2 Soundness of the View Semantics

1010 We turn now to the soundness of the View semantics. Note that in the semantics as presented
 1011 in the body of this paper, writes are recorded in the state as a pair $(w, q) \in W \times Q$. For
 1012 uniformity with the Explicit state semantics, in this proof we record writes in the state as
 1013 a triple $(w, q, t) \in W \times Q \times T$. The transition rule for writes for the View semantics now
 1014 becomes

$$1015 \quad \text{WRITE} \frac{a \in \{wr(x, n), wr^R(x, n)\} \quad (w, q) \in \sigma.OW(t, x) \setminus \sigma.covered \quad \text{fresh}(q, q') \\ \text{writes}' = \sigma.writes \cup \{(a, q', t)\} \quad \text{tview}'_t = \sigma.tview_t[x := (a, q', t)]}{\sigma \xrightarrow{a}_t \sigma[tview_t := \text{tview}'_t, \text{mview}_{(a, q')} := \text{tview}'_t, \text{writes} := \text{writes}']}$$

1016 The other rules are changed similarly. This transformation has no effect on the observable
 1017 behaviour of the semantics. But now, the set of writes $V.writes$ in some View state now has
 1018 the same type (or structure) as the set of events in some Explicit state $E.X$, which will prove
 1019 to be convenient later.

1020 In this section, we prove the following theorem, which states that every behaviour of the
 1021 View semantics is also a behaviour of the Explicit semantics.

1022 ► **Theorem 16.** *For every sequence of steps of the View semantics*

$$1023 \quad V_0 \xrightarrow{a_1} V_1 \xrightarrow{a_2} \dots \xrightarrow{a_n} V_n$$

1025 *such that V_0 is an initial state of the View semantics, there is a sequence of steps of the*
 1026 *Explicit semantics*

$$1027 \quad V_0 \xrightarrow{a_1} V_1 \xrightarrow{a_2} \dots \xrightarrow{a_n} V_n$$

1029 *such that E_0 is an initial state of the Explicit semantics, having the same sequence of actions.*

1030 We prove this using the method of *simulation relations* (see e.g., [24]).

1031 First some preliminary definitions.

1032 ► **Definition 17.** *Given an Explicit state $E = (X, \text{sb}, \text{rf}, \text{mo})$*

- 1033 1. (*Events of a thread.*) Let $E.X_t = \{e \in X \mid \text{tid}(e) = t\}$
- 1034 2. (*mo-maximal.*) For $S \subseteq X$, let $E.\mathbf{max}_{\text{mo}}(S) = \{w \in S \cap W \mid \forall w' \in X. (w, w') \notin \text{mo}\}$
- 1035 3. (*Causal closure.*) For $S \subseteq X$, let $E.\mathbf{cclose}(S) = \{e \mid \exists e' \in S. (e, e') \in \text{eco}^? \circ \text{hb}^?\} \cup I_E$
- 1036 4. (*View modification order.*) Let $V.\mathbf{mo} = \{(w, w') \mid w, w' \in V.writes \wedge \text{tst}(w) < \text{tst}(w')\}$
- 1037 *where $P.\text{eco}$ and $P.\text{hb}$ are defined for an C11 RAR execution as usual.*

1038 We sometimes omit the Explicit state from these auxiliary variables, when the state in
 1039 question is clear from context. For example, $E.\mathbf{max}_{\text{mo}}$ sometimes becomes \mathbf{max}_{mo} .

1040 ► **Lemma 18.** *For any Explicit state $E = (X, \text{sb}, \text{rf}, \text{mo})$*

$$1041 \quad EW(t) = \mathbf{cclose}(X_t) \cap W \tag{13}$$

1043 *and for any set $S \subseteq X$,*

1. $\mathbf{max}_{\text{mo}}(S) \subseteq X$

$$1044 \quad \mathbf{max}_{\text{mo}}(S) \subseteq X \tag{14}$$

2. $\mathbf{max}_{\text{mo}}(S) \subseteq W$

$$1046 \quad \mathbf{max}_{\text{mo}}(S) \subseteq W \tag{15}$$

1048 3. for each variable x there is precisely one write $w \in \mathbf{max}_{\text{mo}}(S)$ such that $\text{loc}(w) = x$.

1049 **Proof.** We prove each property in turn.

1050 ■ (13) This follows immediately from the definitions of EW and \mathbf{cclose} :

$$\begin{aligned}
 1051 \quad w \in EW(t) &\iff w \in W \wedge \exists e' \in X_t. (e, e') \in \text{eco}^?; \text{hb} \\
 1052 &\iff w \in W \wedge w \in \mathbf{cclose}(X_t) \\
 1053 &\iff w \in \mathbf{cclose}(X_t) \cap W \\
 1054
 \end{aligned}$$

1055 ■ (14) By the definition of \mathbf{max}_{mo} , if $w \in \mathbf{max}_{\text{mo}}(S)$ then $w \in S \subseteq X$

1056 ■ (15) By the definition of \mathbf{max}_{mo} , if $w \in \mathbf{max}_{\text{mo}}(S)$ then $w \in W$.

1057 ■ (3) By Property 11 of the explicit semantics, $E.\text{mo}$ totally orders the writes to each x .
 1058 Thus, for any distinct writes v, w such that $\text{loc}(v) = \text{loc}(w) = x$, either $(v, w) \in E.\text{mo}$ or
 1059 $(w, v) \in E.\text{mo}$. In each case, the mo -earlier of the two writes cannot be mo -maximal, and
 1060 therefore v and w cannot both be in $E.\mathbf{max}_{\text{mo}}(S)$. This ensures uniqueness. The fact
 1061 that a write to x exists is immediate from Property 12 of the Explicit semantics.

1062 ◀

1063 Property 3 shows that $\mathbf{max}_{\text{mo}}(X)$ defines a function from variables to writes. We denote by
 1064 $\mathbf{max}_{\text{mo}}(X, x)$ the unique write in $\mathbf{max}_{\text{mo}}(X)$ such that $\text{loc}(w) = x$.

1065 For any Explicit state E , and any $S \subseteq E.X \cap W$, we say that S is *complete* if for every
 1066 location x there is some $w \in S$ with $x = \text{loc}(w)$. Note that if S is complete then $\mathbf{max}_{\text{mo}}(S, x)$
 1067 is defined.

1068 ▶ **Lemma 19.** For any Explicit state E , and any $S \subseteq E.X \cap W$, $E.\mathbf{cclose}(S)$ is complete.

1069 **Proof.** $I_E \subseteq E.\mathbf{cclose}(S)$, and I_E is clearly complete. ◀

1070 ▶ **Lemma 20.** For any Explicit state E , and any complete $S, S' \subseteq E.X$ where S is complete,
 1071 if there exists a $w \in S'$ such that $\text{loc}(w) = x$,

$$\begin{aligned}
 1072 \quad &E.\mathbf{max}_{\text{mo}}(S \cup S', x) \\
 1073 \quad &= \begin{cases} E.\mathbf{max}_{\text{mo}}(S, x) & \text{if } (E.\mathbf{max}_{\text{mo}}(S', x), E.\mathbf{max}_{\text{mo}}(S, x)) \in E.\text{mo} \\ E.\mathbf{max}_{\text{mo}}(S', x) & \text{otherwise} \end{cases} \quad (16) \\
 1074
 \end{aligned}$$

1075 otherwise

$$\begin{aligned}
 1076 \quad &E.\mathbf{max}_{\text{mo}}(S \cup S', x) = E.\mathbf{max}_{\text{mo}}(S, x) \quad (17) \\
 1077
 \end{aligned}$$

1078 **Proof.** If there exists a $w \in S'$ such that $\text{loc}(w) = x$ then both $E.\mathbf{max}_{\text{mo}}(S', x)$ and
 1079 $E.\mathbf{max}_{\text{mo}}(S, x)$ are defined and $E.\mathbf{max}_{\text{mo}}(S \cup S', x)$ is the maximum of these.

1080 If there is no such write then

$$1081 \quad (S \cup S') \cap \{w \mid \text{loc}(w) = x\} = S \cap \{w \mid \text{loc}(w) = x\}$$

1082 and thus $E.\mathbf{max}_{\text{mo}}(S \cup S', x) = E.\mathbf{max}_{\text{mo}}(S, x)$ as required. ◀

1083 Note that for every step of the Explicit semantics $E \xrightarrow{e} E'$ there is some unique write
 1084 that the event e interacts with, for example the write that e reads from if e is a read or
 1085 RMW. The write w mentioned in each of the Explicit semantics transition rules of Figure 14.
 1086 In what follows we exhibit this write as a label on the transition relation. Thus we write
 1087 $E \xrightarrow{w, e} E'$ iff $E \xrightarrow{e} E'$ and w is the write that e interacts with.

1088 ► **Definition 21.** Given an Explicit state E , we say that a thread t is before a write w at
 1089 location x , denoted $t \preceq_x w$ if

$$1090 \quad (E.\mathbf{max}_{\text{mo}}(E.EW(t), x), E.\mathbf{max}_{\text{mo}}(E.\mathbf{cclose}(\{w\}), x)) \in E.\text{mo}$$

1091 Likewise, we say that a write w is before a write t at location x , denoted $w \preceq_x t$ if

$$1092 \quad (E.\mathbf{max}_{\text{mo}}(E.\mathbf{cclose}(\{w\}), x), E.\mathbf{max}_{\text{mo}}(E.EW(t), x)) \in E.\text{mo}$$

1093 ► **Lemma 22 (Max Encountered Writes).** For any Explicit transition $E \xrightarrow{w,e} E'$ where w and
 1094 e synchronise, and for all x , if e is an (acquiring) read then

$$1095 \quad E'.\mathbf{max}_{\text{mo}}(E'.EW(t), x) = \begin{cases} E.\mathbf{max}_{\text{mo}}(E.EW(t), x) & \text{if } w \preceq_x t \\ E.\mathbf{max}_{\text{mo}}(E.\mathbf{cclose}(\{w\}), x) & \text{otherwise} \end{cases} \quad (18)$$

1096 and if e is an update then

$$1097 \quad E'.\mathbf{max}_{\text{mo}}(E'.EW(t), x) \quad (19)$$

$$1098 \quad = \begin{cases} e & \text{if } \text{loc}(w) = x \\ E.\mathbf{max}_{\text{mo}}(E.EW(t), x) & \text{if } \text{loc}(w) \neq x \wedge w \preceq_x t \\ E.\mathbf{max}_{\text{mo}}(E.\mathbf{cclose}(\{w\}), x) & \text{otherwise} \end{cases}$$

1099 Further, for any Explicit transition $E \xrightarrow{w,e} E'$ where w and e do not synchronise, if e is a
 1100 read then

$$1101 \quad E'.\mathbf{max}_{\text{mo}}(E'.EW(t), x) = \begin{cases} w & \text{if } \text{loc}(w) = x \\ E.\mathbf{max}_{\text{mo}}(E.EW(t), x) & \text{otherwise} \end{cases} \quad (20)$$

1102 and if e is a write or update then

$$1103 \quad E'.\mathbf{max}_{\text{mo}}(E'.EW(t), x) = \begin{cases} e & \text{if } \text{loc}(w) = x \\ E.\mathbf{max}_{\text{mo}}(E.EW(t), x) & \text{otherwise} \end{cases} \quad (21)$$

1104 **Proof.** We prove each in turn.

1105 (Equation 18) Consider

$$1106 \quad E'.\mathbf{max}_{\text{mo}}(E'.EW(t), x)$$

$$1107 \quad = \quad (\text{trans-rel})$$

$$1108 \quad E.\mathbf{max}_{\text{mo}}(E'.EW(t), x)$$

$$1109 \quad = \quad (\text{trans-rel and defn of } \mathbf{cclose})$$

$$1110 \quad E.\mathbf{max}_{\text{mo}}(E.EW(t) \cup E.\mathbf{cclose}(\{w\}), x)$$

$$1111 \quad = \quad (\text{see below})$$

$$1112 \quad \begin{cases} E.\mathbf{max}_{\text{mo}}(E.EW(t), x) & \text{if } w \preceq_x t \\ E.\mathbf{max}_{\text{mo}}(E.\mathbf{cclose}(\{w\}), x) & \text{otherwise} \end{cases}$$

1113 This last step is a consequence of Lemma 20 and the fact that both $E.EW(t)$ and $E.\mathbf{cclose}(\{w\})$
 1114 are complete.

(Equation 19) We first consider the case when $x = \text{loc}(e)$. In this case

$$E'.\mathbf{max}_{\text{mo}}(E'.EW(t), x) = e$$

0:32 Owicki-Gries Reasoning for C11 RAR

1120 as required. If $x \neq \text{loc}(e)$, then

$$\begin{aligned}
 1121 \quad & E'.\mathbf{max}_{\text{mo}}(E'.EW(t), x) \\
 1122 \quad & = \quad (\text{trans-rel}) \\
 1123 \quad & E.\mathbf{max}_{\text{mo}}(E'.EW(t), x) \\
 1124 \quad & = \quad (\text{trans-rel and defn of } \mathbf{cclose}) \\
 1125 \quad & E.\mathbf{max}_{\text{mo}}(E.EW(t) \cup E.\mathbf{cclose}(\{w\}) \cup \{e\}, x) \\
 1126 \quad & = \quad (x \neq \text{loc}(e) \text{ and Lemma 20}) \\
 1127 \quad & E.\mathbf{max}_{\text{mo}}(E.EW(t) \cup E.\mathbf{cclose}(\{w\}), x) \\
 1128 \quad & = \quad (\text{see below}) \\
 1129 \quad & \begin{cases} E.\mathbf{max}_{\text{mo}}(E.EW(t), x) & \text{if } w \preceq_x t \\ E.\mathbf{max}_{\text{mo}}(E.\mathbf{cclose}(\{w\}), x) & \text{otherwise} \end{cases} \\
 1130
 \end{aligned}$$

1131 Agin, this last step is a consequence of Lemma 20 and the fact that both $E.EW(t)$ and
 1132 $E.\mathbf{cclose}(\{w\})$ are complete.

(Equation 20) We first consider the case when $x = \text{loc}(e)$. In this case

$$E'.\mathbf{max}_{\text{mo}}(E'.EW(t), x) = e$$

1133 as required. If $x \neq \text{loc}(e)$, then

$$\begin{aligned}
 1134 \quad & E'.\mathbf{max}_{\text{mo}}(E'.EW(t), x) \\
 1135 \quad & = \quad (\text{trans-rel}) \\
 1136 \quad & E.\mathbf{max}_{\text{mo}}(E'.EW(t), x) \\
 1137 \quad & = \quad (\text{trans-rel and defn of } \mathbf{cclose}) \\
 1138 \quad & E.\mathbf{max}_{\text{mo}}(E.EW(t) \cup \{w\} \cup \{e\}, x) \\
 1139 \quad & = \quad (\text{because } x \neq \text{loc}(e) \text{ and Lemma 20}) \\
 1140 \quad & E.\mathbf{max}_{\text{mo}}(E.EW(t), x) \\
 1141
 \end{aligned}$$

1142 (Equation 21) The proof here is identical to that for Equation 20. ◀

1143 We define a simulation relation R between a View state $V = (\text{writes}, \text{tview}, \text{mview}, \text{covered})$
 1144 and an Explicit state $E = (X, \text{sb}, \text{rf}, \text{mo})$ to be the conjunction of the following properties.

1145 ■ Both executions contain the same set of writes:

$$1146 \quad V.\text{writes} = E.X \cap W \tag{22}$$

1148 ■ For all threads t and variables x ,

$$1149 \quad V.\text{tview}(t, x) = E.\mathbf{max}_{\text{mo}}(E.EW(t), x) \tag{23}$$

1151 ■ For all $w \in \text{writes}$ and variables x ,

$$1152 \quad V.\text{mview}(w, x) = E.\mathbf{max}_{\text{mo}}(E.\mathbf{cclose}(\{w\}), x) \tag{24}$$

1154 ■ Both executions agree on the order of writes.

$$1155 \quad V.\text{mo} = E.\text{mo} \tag{25}$$

1157 recall that $V.\text{mo} = \{(w, w') \mid w, w' \in V.\text{writes} \wedge \text{tst}(w) < \text{tst}(w')\}$.

1158 ■ The executions agree on the set of covered writes:

$$1159 \quad V.covered = E.covered \quad (26)$$

1161 Our first lemma relates the viewfront and the observable writes of R -related states.

1162 ► **Lemma 23.** *For any View state V and Explicit state E such that $(V, E) \in R$, and for all*
 1163 *threads t and locations x , $V.OW(t, x) = E.OW(t, x)$.*

1164 **Proof.** First, observe that for all w such that $loc(w) = x$

$$1165 \quad tst(V.tview_t(x)) \leq tst(w) \iff (E.max_{mo}(E.EW(t), x), w) \in E.mo \quad (27)$$

1167 But this is an immediate consequence of the fact that $V.tview(t, x) = E.max_{mo}(E.EW(t), x)$
 1168 and that $E.mo = V.mo$ (both of these are consequences of R).

1169 But now, because $E.max_{mo}(E.EW(t), x)$ is mo -maximal among the set of encountered
 1170 writes, we have

$$1171 \quad (E.max_{mo}(E.EW(t), x), w) \in E.mo \iff \forall w' \in E.EW(t), x. (w, w') \notin E.mo \quad (28)$$

1173 Now consider

$$\begin{aligned} 1174 \quad & w \in V.OW(t, x) \\ 1175 \quad & \iff \text{(definition)} \\ 1176 \quad & w \in V.writes \wedge loc(a) = x \wedge tst(V.tview_t(x)) \leq tst(w) \\ 1177 \quad & \iff \text{(by } R) \\ 1178 \quad & w \in E.X \cap W \wedge loc(a) = x \wedge tst(V.tview_t(x)) \leq tst(w) \\ 1179 \quad & \iff (28) \\ 1180 \quad & w \in E.X \cap W \wedge loc(a) = x \wedge \forall w' \in E.EW(t), x. (w, w') \notin E.mo \\ 1181 \quad & \iff w \in E.OW(t, x) \end{aligned}$$

1183 as required. ◀

1184 In the preservation proof, we use the following stability properties:

1185 ► **Lemma 24.** *For all $E \xrightarrow{w, e} E'$, every location x , and every thread $t' \neq tid(e)$*

$$1186 \quad E'.max_{mo}(E'.EW(t'), x) = E.max_{mo}(E.EW(t'), x) \quad (29)$$

1188 and for every write $w' \neq e$

$$1189 \quad E'.max_{mo}(E'.cclose(\{w'\}), x) = E.max_{mo}(E.cclose(\{w'\}), x) \quad (30)$$

1191 Further, for all $V \xrightarrow{w, e} V'$, every location x , and every thread $t' \neq tid(e)$

$$1192 \quad V'.tview(t', x) = V.tview(t', x) \quad (31)$$

1194 and for every write $w' \neq e$

$$1195 \quad V'.mview(w', x) = V.mview(w', x) \quad (32)$$

0:34 Owicki-Gries Reasoning for C11 RAR

1197 ► **Lemma 25** (Mod-order agreement). *For any View state*

$$1198 \quad V = (\text{writes}, \text{tview}, \text{mview}, \text{covered})$$

1199 *and Explicit state* $E = (X, \text{sb}, \text{rf}, \text{mo})$ *such that* $(V, E) \in R$, *and every thread* t , *write* w *and*
1200 *variable* x

$$1201 \quad \text{tst}(V.\text{tview}(t)(x)) < \text{tst}(V.\text{mview}(w)(x)) \iff t \preceq_x w$$

1203 **Proof.** We reason thusly

$$1204 \quad \text{tst}(V.\text{tview}(t)(x)) < \text{tst}(V.\text{mview}(w)(x)) \iff$$

$$1205 \quad \text{tst}(E.\mathbf{max}_{\text{mo}}(E.EW(t), x)) < \text{tst}(E.\mathbf{max}_{\text{mo}}(E.EW(t), x)) \iff \text{by 23}$$

$$1206 \quad t \preceq_x w \quad \text{by 25}$$

1208 ◀

1209 ► **Lemma 26** (View Mod Order). *For any* $V \xrightarrow{w,e}_t V'$ *where* e *is a write or update,* $V'.\text{mo} =$
1210 $V.\text{mo}[w, (e, q, t)]$ *where* q *is the fresh rational used to tag the operation in the View transition*
1211 *relation.*

1212 **Proof.** The new write is added into mo immediately after the write w and before all
1213 subsequent writes to the same variable. ◀

1214 ► **Lemma 27** (R -preservation). *For any View state*

$$1215 \quad V = (\text{writes}, \text{tview}, \text{mview}, \text{covered})$$

1216 *and Explicit state* $E = (X, \text{sb}, \text{rf}, \text{mo})$ *such that* $(V, E) \in R$, *and every View state* $V' =$
1217 $(\text{writes}', \text{tview}', \text{mview}', \text{covered}')$ *such that* $V \xrightarrow{w,e} V'$, *we have* $E \xrightarrow{w,e} E'$ *for some* E' *and*
1218 $(V', E') \in R$.

1219 **Proof.** We proceed by cases on the type of the operation e and whether or not the operation
1220 is synchronising. In what follows, let $t = \text{tid}(e)$.

1221 **Case 1.** e is of the form $rd(x, n)$. Let $E = (X', \text{sb}', \text{rf}', \text{mo}')$ where

$$1222 \quad (X', \text{sb}') = (X, \text{sb}) + e$$

$$1223 \quad \text{rf}' = \text{rf} \cup \{(w, e)\}$$

$$1224 \quad \text{mo}' = \text{mo}$$

1226 The precondition of the View transition ensures that $w \in V.OW(t, x)$, and thus by Lemma
1227 23, we have $w \in E.OW(t)$. Thus, we have $E \xrightarrow{w,e} E'$. It remains to show that $(V', E') \in R$,
1228 which we do by considering each of the equations in the definition of R in turn.

■

$$1229 \quad V'.\text{writes} = V.\text{writes} \quad \text{by View trans-rel}$$

$$1230 \quad = E.X \cap W \quad \text{by } R$$

$$1231 \quad = E'.X \cap W \quad \text{by Explicit trans-rel}$$

1233 ■ We need to show that $V'.\text{tview}(t, x) = E'.\mathbf{max}_{\text{mo}}(E'.EW(t), x)$. For all $t' \neq t$, it is easy
1234 to see that Equations 29 and 31 ensure that this property is preserved.

1235 Several cases remain. We discuss the first two. The remaining cases follow in a similar
1236 way. In the first case, we assume the following

$$1237 \quad w, e \text{ synchronise} \quad (33)$$

$$1238 \quad \text{loc}(e) = x \quad (34)$$

$$1239 \quad \text{tst}(V.\text{tview}(t)(x)) < \text{tst}(V.\text{mview}(w)(x)) \quad (35)$$

$$1240 \quad (E.\mathbf{max}_{\text{mo}}(E.EW(t), x), E.\mathbf{max}_{\text{mo}}(E.\mathbf{cclose}(\{w\}), x)) \in E.\text{mo} \quad (36)$$

1241

1242 Note that by 25 the last two assumptions are equivalent. Then,

$$1243 \quad V'.\text{tview}(t)(x) = (V.\text{tview}(t) \otimes V.\text{mview}(w))(x) \quad \text{by View trans-rel}$$

$$1244 \quad = V.\text{mview}(w)(x) \quad \text{by hyp. and def of } \otimes$$

$$1245 \quad = E.\mathbf{max}_{\text{mo}}(E.\mathbf{cclose}(\{w\}), x) \quad \text{by } R$$

$$1246 \quad = E'.\mathbf{max}_{\text{mo}}(E'.EW(t), x) \quad \text{see below}$$

1247

1248 The last step is a consequence of the first, second, and fourth assumptions together with
1249 Lemma 18.

1250 In the second case, we assume

$$1251 \quad w, e \text{ synchronise} \quad (37)$$

$$1252 \quad \text{loc}(e) = x \quad (38)$$

$$1253 \quad \text{tst}(V.\text{mview}(w)(x)) < \text{tst}(V.\text{tview}(t)(x)) \quad (39)$$

$$1254 \quad (E.\mathbf{max}_{\text{mo}}(E.\mathbf{cclose}(\{w\}), x), E.\mathbf{max}_{\text{mo}}(E.EW(t), x)) \in E.\text{mo} \quad (40)$$

1255

1256 (The new hypothesis differs from the previous in that we have changed the view that
1257 supplies the latest write.) Note that by 25 the last two assumptions are equivalent. Then,

$$1258 \quad V'.\text{tview}(t)(x) = (V.\text{tview}(t) \otimes V.\text{mview}(w))(x) \quad \text{by View trans-rel}$$

$$1259 \quad = V.\text{tview}(t)(x) \quad \text{by hyp. and def of } \otimes$$

$$1260 \quad = E.\mathbf{max}_{\text{mo}}(E.EW(t), x) \quad \text{by } R$$

$$1261 \quad = E'.\mathbf{max}_{\text{mo}}(E'.EW(t), x) \quad \text{see below}$$

1262

1263 The last step is a consequence of the first, second, and fourth assumptions together with
1264 Lemma 18.

1265 ■ We need to show that $V'.\text{mview}(w, x) = E'.\mathbf{max}_{\text{mo}}(E'.\mathbf{cclose}(\{w\}), x)$ for all w . But
1266 because e is a read we have

$$1267 \quad V'.\text{mview} = V.\text{mview}$$

$$1268 \quad E'.\mathbf{cclose} = E.\mathbf{cclose}$$

$$1269 \quad E'.\mathbf{max}_{\text{mo}} = E.\mathbf{max}_{\text{mo}}$$

1270

1271 so the preservation result follows immediately.

■

$$1272 \quad V'.\text{mo} = V.\text{mo} \quad \text{by View trans-rel}$$

$$1273 \quad = E.\text{mo} \quad \text{by } R$$

$$1274 \quad = E'.\text{mo} \quad \text{by Explicit trans-rel}$$

1275

■

$$\begin{array}{ll}
 1276 & V'.covered = V.covered & \text{by View trans-rel} \\
 1277 & = E.covered & \text{by } R \\
 1278 & = E'.covered & \text{by Explicit trans-rel} \\
 1279 & &
 \end{array}$$

1280 **Case 2.** e is of the form $wr(x, n)$. Let $E = (X', sb', rf', mo')$ where

$$\begin{array}{ll}
 1281 & (X', sb') = (X, sb) + e \\
 1282 & rf' = rf \\
 1283 & mo' = mo[w, e] \\
 1284 &
 \end{array}$$

1285 The precondition of the View transition ensures that $w \in V.OW(t, x)$, and thus by Lemma
 1286 23, we have $w \in E.OW(t)$. Thus, we have $E \overset{w, e}{\rightsquigarrow} E'$. It remains to show that $(V', E') \in R$,
 1287 which we do by considering each of the equations in the definition of R in turn.

■

$$\begin{array}{ll}
 1288 & V'.writes = V.writes \cup \{e\} & \text{by View trans-rel} \\
 1289 & = (E.X \cap W) \cup \{e\} & \text{by } R \\
 1290 & = E'.X \cap W & \text{by Explicit trans-rel} \\
 1291 &
 \end{array}$$

■ If $loc(e) \neq x$ then the thread view and mo restricted to x are unchanged. If $loc(e) = x$, then $E'.max_{mo}(E'.EW(t), x) = e$ and $V'.tview(t, x) = e$ so

$$E'.max_{mo}(E'.EW(t), x) = V'.tview(t, x) = e$$

1292 as required.

■ If $loc(e) \neq x$ then the thread view and mo restricted to x are unchanged. Assume $loc(e) = x$. For all $w' \in V'.writes$ where $w' \neq e$, then

$$E'.max_{mo}(E'.cclose(w), x) = E.max_{mo}(E.cclose(w), x)$$

so the property is preserved. In the final case,

$$E'.max_{mo}(E'.cclose(e), x) = E.max_{mo}(E.EW(t) \cup \{e\}) = e$$

1293 but $V'.mview(e, x) = e$ as required.

■

$$\begin{array}{ll}
 1294 & V'.mo = V.mo[w, e] & \text{by Lemma 26} \\
 1295 & = E.mo[w, e] & \text{by } R \\
 1296 & = E'.mo & \text{by Explicit trans-rel} \\
 1297 &
 \end{array}$$

■

$$\begin{array}{ll}
 1298 & V'.covered = V.covered & \text{by View trans-rel} \\
 1299 & = E.covered & \text{by } R \\
 1300 & = E'.covered & \text{by Explicit trans-rel} \\
 1301 &
 \end{array}$$

1302 **Case 3.** e is of the form $upd^{RA}(x, m, n)$. Let $E = (X', sb', rf', mo')$ where

1303 $(X', sb') = (X, sb) + e$

1304 $rf' = rf \cup \{(w, e)\}$

1305 $mo' = mo[w, e]$
1306

1307 The precondition of the View transition ensures that $w \in V.OW(t, x)$, and thus by Lemma
1308 23, we have $w \in E.OW(t)$. Thus, we have $E \xrightarrow{w, e} E'$. It remains to show that $(V', E') \in R$,
1309 which we do by considering each of the equations in the definition of R in turn.

1310 The proof that $(V', E') \in R$ is essentially a combination of the proof for reads and writes.

1311 $V'.writes = V.writes \cup \{e\}$ by View trans-rel
1312 $= (E.X \cap W) \cup \{e\}$ by R
1313 $= E'.X \cap W$ by Explicit trans-rel
1314

1315 ■ The proof here is the same as that for reads.

1316 ■ The proof here is the same as that for writes.

1317 $V'.mo = V.mo[w, e]$ by Lemma 26
1318 $= E.mo[w, e]$ by R
1319 $= E'.mo$ by Explicit trans-rel
1320

1321 $V'.covered = V.covered \cup \{w\}$ by View trans-rel
1322 $= E.covered \cup \{w\}$ by R
1323 $= E'.covered$ by Explicit trans-rel
1324

1325

1326 Finally, we must prove that for every initial View state, there is an R -related initial
1327 Explicit state.

1328 ► **Lemma 28.** *For every initial state of the View semantics V_0 , there is an initial state of*
1329 *the Explicit semantics E_0 such that $(V_0, E_0) \in R$*

1330 **Proof.** Let $V_0 = (writes, tview, mview, covered)$. We let $E_0 = (X, sb, rf, mo)$, where ⁶

1331 $X = writes$ (41)

1332 $sb = \emptyset$ (42)

1333 $rf = \emptyset$ (43)

1334 $mo = \emptyset$ (44)
1335

1336 We prove each property of the simulation relation in turn.

1337 ■ (22) This is immediate.

⁶ Recall that in our setting $writes$ and X have the same type.

0:38 Owicki-Gries Reasoning for C11 RAR

- 1338 ■ (23) For all threads t and variables x , $E.\mathbf{max}_{\text{mo}}(E.EW(t), x)$ is the initialising write to x ,
1339 which is also the value of $V.tview(t, x)$.
- 1340 ■ (24) Similarly to 23, for all writes w and variables x , $E.\mathbf{cclose}(w)$ and $V.mview(w)$ is the
1341 initialising write to x .
- 1342 ■ (25) This is immediate.
- 1343 ■ (26) This is immediate.
- 1344 ◀