

Verifying that a compiler preserves concurrent value-dependent information-flow security

Robert Sison 

Data61, CSIRO, Australia

UNSW Sydney, Australia

Robert.Sison@data61.csiro.au

Toby Murray

University of Melbourne, Australia

toby.murray@unimelb.edu.au

Abstract

It is common to prove by reasoning over source code that programs do not leak sensitive data. But doing so leaves a gap between reasoning and reality that can only be filled by accounting for the behaviour of the compiler. This task is complicated when programs enforce *value-dependent* information-flow security properties—in which the sensitivity of values is itself value-dependent—and furthermore especially complicated when programs exploit shared-variable concurrency.

Prior work has formally defined a notion of concurrency-aware refinement for preserving value-dependent security properties. However, that notion is considerably more complex than standard refinement definitions typically applied in the verification of semantics preservation by compilers. To date it remains unclear whether it can be applied to a realistic compiler, because there exist no general decomposition principles for separating it into smaller, more familiar, proof obligations.

In this work, we provide such a decomposition principle, which we show dramatically reduces the complexity of proving secure refinement. Further, we demonstrate its applicability to secure compilation, by proving in Isabelle/HOL the preservation of value-dependent security by a proof-of-concept compiler from an imperative While language to a generic RISC-style assembly language, for programs with shared-memory concurrency mediated by locking primitives. Finally, we execute our compiler in Isabelle on a While language model of the Cross Domain Desktop Compositor, demonstrating to our knowledge the first use of a compiler verification result to carry an information-flow security property down to the assembly-level model of a non-trivial concurrent program.

2012 ACM Subject Classification Security and privacy → Logic and verification

Keywords and phrases Secure compilation, Information flow security, Concurrency, Verification

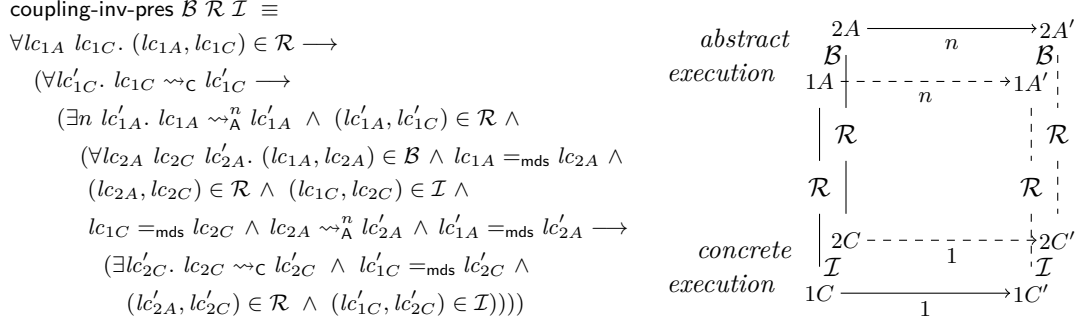
Funding *Robert Sison*: Australian Government RTP Scholarship & Data61 Research Project Award

Acknowledgements Many thanks to Carroll Morgan, Kai Engelhardt, Gerwin Klein, Christine Rizkallah, Matthew Brecknell, and Johannes Åman Pohjola, for their very helpful feedback on earlier versions of this paper.

1 Introduction

It is well known that program translations of the kind carried out by compilers can in principle break security properties like confidentiality [8, 1]. Yet source level reasoning about confidentiality remains common [14, 13, 12]. Existing verified compilers like CompCert [10] and CakeML [9] preserve semantics, but semantics preservation alone may be insufficient to preserve confidentiality, especially for shared memory concurrent programs whose threads must guard against timing leaks in order to prevent them manifesting as storage leaks [16].

Supporting secure compilation of programs that must enforce *value-dependent* security policies poses an additional challenge, because in such policies the sensitivity of a memory location can depend on the values held in other memory locations. Thus, unlike prior work



■ **Figure 1** Definition, graphical depiction of refinement preservation for secure-refinement (Def. 6)

45 on secure compilation [3], preserving security under refinement requires a refinement relation
 46 that is strong enough to preserve those memory contents on which the policy depends.

47 In prior work [16], we presented a definition for a notion of value-dependent security-
 48 preserving refinement that is compositional for concurrent programs: by applying it to each
 49 thread individually, one can derive a secure refinement of the concurrent composition.

50 The essence of this notion of security-preserving refinement (presented fully in Section 2.2)
 51 is in its refinement preservation obligation (coupling-inv-pres in Figure 1). Here, the usual
 52 square-shaped commuting diagram that is commonly used to depict (semantics-preserving)
 53 refinement (Figure 2a) has been replaced by a *cube* (Figure 1). The additional dimension of
 54 this cube reflects that it preserves a 2-safety hyperproperty [5] that compares two executions
 55 rather than examining a single one. As such, it is significantly more complicated to prove than
 56 standard notions of semantics-preserving refinement typical in verified compilation [10, 9].

57 To date there exist no verified compilers for shared-variable concurrent programs proved to
 58 preserve value-dependent information-flow security. We argue that without a *decomposition*
 59 *principle* the cube-shaped refinement notion is too cumbersome to prove for realistic compilers.

60 In this paper, we tackle the central problem of making our notion of secure refinement
 61 applicable to verified secure compilation. Firstly, we present a decomposition principle that
 62 makes the cube-shaped notion more tractable. Secondly, we demonstrate its tractability
 63 with our major contribution: a machine-checked formal proof of concurrent value-dependent
 64 security preservation, for a proof-of-concept compiler.

65 In Section 3 we present our decomposition principle, which decomposes the cube (Figure 1)
 66 into three separate obligations (Figure 2). The first of these is akin to semantics-preserving
 67 refinement, while the second and third essentially ensure together that the refinement has
 68 not introduced any termination- and timing-leaks.

69 In Section 4 we show how the decomposition principle allows the effort to prove secure
 70 refinement for individual programs to be dramatically reduced. There, we present a side-by-
 71 side comparison of the proofs of secure refinement for a program that branches on secrets (and
 72 is thereby especially prone to introduced timing leaks), with and without the decomposition
 73 principle. We find that the decomposition reduces proof complexity by almost 50%.

74 In Section 5, we present our compiler and its formal verification, as an application of
 75 the decomposition principle. This compiler translates concurrent programs written in an
 76 imperative While language, with locking primitives for mediating access to shared memory,
 77 into a RISC-style assembly language. It does so by compiling each thread individually, and
 78 in doing so preserves a formal security property that remains compositional between threads.
 79 Furthermore, our compiler demonstrates a way of formalising and proving when it is safe

80 for a compiler to perform optimisations in the presence of concurrency. To ensure that
 81 the contents of shared memory locations are preserved under compilation despite potential
 82 interference from other threads, our compiler tracks which shared memory locations are
 83 *stable* (free from any such interference). It then leverages this tracking to avoid redundant
 84 loads from stable shared variables safely, that would otherwise be considered unsafe to omit.

85 All results are mechanised in Isabelle/HOL,¹ and in Section 6 we explain how, in order
 86 to validate our theory, we instantiated it so that we could execute our compiler in Isabelle.
 87 This enabled us to execute it over a While language model of the Cross Domain Desktop
 88 Compositor [4] (CDDC), a concurrent program that enforces information flow control over
 89 value-dependently classified input. To our knowledge this is the first proof of information
 90 flow security for an assembly-level model of a non-trivial concurrent program, demonstrating
 91 the power of verified secure compilation for deriving security properties of compiled code.

92 **2 Background**

93 To begin with, we recall relevant details of the theory presented in our previous work [16].

94 **2.1 Concurrent value-dependent noninterference (CVDNI)**

95 The information-flow security property CVDNI is defined in terms of two main elements:

- 96 1. a binary *strong low bisimulation modulo modes* relation \mathcal{B} between program configurations,
 97 that establishes the required information-flow security property. In a manner typical of
 98 techniques for Goguen & Meseguer-style noninterference [6], it relates states that agree on
 99 their “low” portions, and demands that lock-step execution preserve that correspondence.
 100 This section will explain how it is specialised further for shared-variable concurrency.
- 101 2. a *classification* function \mathcal{L} on a program configuration that determines its “low” portion,
 102 thus affecting \mathcal{B} ’s requirements. Atypically however, \mathcal{L} here can be *value-dependent* and
 103 depend on the program configuration itself – i.e. it is dynamic rather than static.

104 The definitions now presented are adaptations from [16] Section III-2b, with simplifications
 105 for presentation as noted. The theory is parameterised over the type of values Val , a finite
 106 set of shared variables Var , and a *deterministic evaluation step semantics* \rightsquigarrow between *local*
 107 *configurations* (of a thread in a concurrent program) each denoted by a triple $\langle tps, mds, mem \rangle$:

- 108 ■ tps is the *thread-private state*, including the program text, current program position, and
 109 any statically allocated memory or storage space considered permanently thread-private
 110 by both of the attacker-model and the concurrency-model.
- 111 ■ $mds \in Mode \Rightarrow Var$ set is the (*access*) *mode state*, which is ghost state associating each
 112 of $Mode = \{\mathbf{AsmNoW}, \mathbf{AsmNoRW}, \mathbf{GuarNoW}, \mathbf{GuarNoRW}\}$ with (a set of) shared
 113 variables. Most pertinent to this paper are the **Asm** modes, which track assumptions made
 114 by the thread over time, about which variables will not be modified (**NoW**) or neither
 115 modified nor read (**NoRW**) by other threads. (The **Guar** modes specify corresponding
 116 guarantees that the thread makes to other threads.) For instance, if a variable $x \in Var$ is
 117 present in mds **AsmNoRW**, then the thread is assuming that no other threads will read
 118 nor modify x . Mode states facilitate compositional, assume-guarantee [7] style reasoning.
- 119 ■ $mem \in Var \Rightarrow Val$ is *shared memory* considered potentially accessible to all the threads,
 120 in the form of a total map from shared variable names to their values.

¹ The Isabelle/HOL theories are available at <https://covern.org>. The *wr-compiler* totals ~ 7.6 k lines of Isabelle proof script, and the verification + compilation of the 2-thread CDDC model totals ~ 1.8 k lines.

121 The theory is then further parameterised by the value-dependent classification function
 122 $\mathcal{L} :: (Var \Rightarrow Val) \Rightarrow Var \Rightarrow \{\text{High}, \text{Low}\}$, where $\mathcal{L}_{mem} x = \text{Low}$ indicates that the value of
 123 variable x is considered observable to the attacker when the memory is mem . This allows the
 124 instantiator of the theory to specify an attacker model that can make observations of a subset of
 125 the shared memory whose scope changes dynamically over the lifetime of the program. The set
 126 $\mathcal{C} = \{y \mid \exists x. y \in \mathcal{Cvars} x\}$ (where $\forall y \in \mathcal{Cvars} x. mem_1 y = mem_2 y \implies \mathcal{L}_{mem_1} x = \mathcal{L}_{mem_2} x$)
 127 contains all *control variables* $y \in Var$ on whose values the classification \mathcal{L} of any variable
 128 $x \in Var$ are dependent. Consequently, when two memories mem_1 and mem_2 agree on the
 129 values of all control variables in \mathcal{C} , we have that $\mathcal{L}_{mem_1} x = \mathcal{L}_{mem_2} x$ for all variables $x \in Var$.

130 The notion of observational indistinguishability used for the noninterference property is
 131 then defined over memories as follows. To support compositionality for concurrent programs,
 132 observability is relative not only to the attacker, but also to the other threads in the system,
 133 as designated by x being assumed to be *readable*: $readable\ mds\ x \equiv x \notin mds\ \mathbf{AsmNoRW}$.

► **Definition 1** (Low-equivalent memories modulo modes).

$$134 \quad mem_1 =_{m\!ds}^{\text{Low}} mem_2 \equiv$$

$$135 \quad \forall x. x \in \mathcal{C} \vee \mathcal{L}_{mem_1} x = \text{Low} \wedge readable\ mds\ x \longrightarrow mem_1 x = mem_2 x$$

137 The guard $x \in \mathcal{C}$ reflects that control variables are always considered observable, regardless of
 138 whether or not they are assumed to be readable. For the rest of this paper, we will use $=_{m\!ds}^{\text{Low}}$
 139 to lift $=_{m\!ds}^{\text{Low}}$ to local program configurations, asserting also that both local configurations
 140 have the same mode state. Additionally, we will introduce the notation $lc_1 =_{m\!ds} lc_2$ to denote
 141 when local program configurations lc_1 and lc_2 are *modes-equal* (have the same mode state).

142 The per-thread compositional security property **com-secure** asserts the existence of a
 143 witness relation \mathcal{B} for every possible observationally equivalent pair of starting configurations:

► **Definition 2** (Per-thread compositional CVDNI property).

$$144 \quad \text{com-secure} (tps, mds) \equiv$$

$$145 \quad \forall mem_1 mem_2. mem_1 =_{m\!ds}^{\text{Low}} mem_2 \longrightarrow$$

$$146 \quad (\exists \mathcal{B}. \text{strong-low-bisim-mm } \mathcal{B} \wedge (\langle tps, mds, mem_1 \rangle, \langle tps, mds, mem_2 \rangle) \in \mathcal{B})$$

148 where all such witness relations \mathcal{B} must be a *strong low bisimulation modulo modes*:

$$149 \quad \text{strong-low-bisim-mm } \mathcal{B} \equiv \text{cg-consistent } \mathcal{B} \wedge \text{sym } \mathcal{B} \wedge$$

$$150 \quad (\forall lc_1 lc_2. (lc_1, lc_2) \in \mathcal{B} \wedge lc_1 =_{m\!ds} lc_2 \longrightarrow lc_1 =_{m\!ds}^{\text{Low}} lc_2 \wedge$$

$$151 \quad (\forall lc'_1. lc_1 \rightsquigarrow lc'_1 \longrightarrow (\exists lc'_2. lc_2 \rightsquigarrow lc'_2 \wedge lc'_1 =_{m\!ds} lc'_2 \wedge (lc'_1, lc'_2) \in \mathcal{B})))$$

153 That is, \mathcal{B} must maintain observational indistinguishability by requiring that all configu-
 154 ration pairs it relates that have the same mode state, are low-equivalent modulo modes.

155 Furthermore, it must be a *bisimulation* by being symmetric and *progressing to itself*: any
 156 step taken by one of the configurations must be able to be matched by a step taken by the
 157 configuration related to it, such that the destinations remain related by \mathcal{B} (and modes-equal).

158 Finally—and the most crucial element ensuring the property’s compositionality for con-
 159 current programs—is the condition that \mathcal{B} must be **cg-consistent**: *closed under globally*
 160 *consistent changes* made to memory by other threads, which is to say, changes that preserve
 161 low-equivalence and are permitted by the current mode state mds . Specifically, the environ-
 162 ment (of other threads) is permitted to change either of variable x ’s value or its classification
 163 only when x is *writable*: $writable\ mds\ x \equiv x \notin mds\ \mathbf{AsmNoW} \wedge x \notin mds\ \mathbf{AsmNoRW}$.

► **Definition 3** (Closedness under globally consistent changes).

$$\begin{aligned}
164 \quad & \text{cg-consistent } \mathcal{B} \equiv \\
165 \quad & \forall tps_1 \ mem_1 \ tps_2 \ mem_2 \ mds. (\langle tps_1, mds, mem_1 \rangle, \langle tps_2, mds, mem_2 \rangle) \in \mathcal{B} \longrightarrow \\
166 \quad & (\forall mem'_1 \ mem'_2. (\forall x. (mem_1 \ x \neq mem'_1 \ x \ \vee \ mem_2 \ x \neq mem'_2 \ x \ \vee \\
167 \quad & \quad \mathcal{L}_{mem_1} \ x \neq \mathcal{L}_{mem'_1} \ x) \longrightarrow \text{writable } mds \ x) \wedge mem'_1 \stackrel{\text{Low}}{=}_{mds} mem'_2 \longrightarrow \\
168 \quad & (\langle tps_1, mds, mem'_1 \rangle, \langle tps_2, mds, mem'_2 \rangle) \in \mathcal{B})
\end{aligned}$$

170 Theorem 3.1 of our prior work [16] then gives us that the parallel composition of com-secure
171 programs is itself a program that enforces a system-wide value-dependent noninterference
172 property (sys-secure, for whose details we refer the reader to Section III-2(a) of [16]).

173 2.2 CVDNI-preserving refinement

174 The proof of CVDNI-preserving refinement for a thread of a concurrent program relies on
175 two binary relations to be nominated by the instantiator of the theory:

- 176 1. a *refinement relation* \mathcal{R} relating local configurations of the abstract program to local
177 configurations of the concrete program: abstract must simulate concrete, in a sense typical
178 of much other work on program refinement, including compiler verification efforts.
- 179 2. a *coupling invariant* \mathcal{I} that serves as a means of discarding unwanted pairs of local
180 configurations of the concrete program *after the refinement*, when deriving a strong low
181 bisimulation modulo modes for the concrete program from a given \mathcal{B} and \mathcal{R} .

182 The definitions in the remainder of this section are adaptations of those from Murray et
183 al. [16] Section V. For better readability, we present a simplified version of the requirements,
184 in which no new shared variables are added by the refinement. Consequently we will introduce
185 the notation $\stackrel{\text{mem}}{=}_{mds}$ to denote that two local configurations have equal mode state and memory,
186 regardless of whether relating configurations of the same language, or of differing languages.

187 The essence of the proof technique is to require that a number of conditions—analogueous
188 to those for strong-low-bisim-mm—be imposed on the nominated \mathcal{R} and \mathcal{I} in relation to a
189 given witness relation \mathcal{B} establishing CVDNI for the abstract program.

190 Regarding the closedness under changes by other threads that ensures compositionality
191 for concurrency, here there are two requirements: closed-others (Definition 4) pertaining to
192 \mathcal{R} , and cg-consistent (Definition 3 from Section 2.1) pertaining to \mathcal{I} .

► **Definition 4** (Closedness of refinements under changes by others).

$$\begin{aligned}
193 \quad & \text{closed-others } \mathcal{R} \equiv \\
194 \quad & \forall tps_A \ tps_C \ mds \ mem \ mem'. (\langle tps_A, mds, mem \rangle_A, \langle tps_C, mds, mem \rangle_C) \in \mathcal{R} \wedge \\
195 \quad & (\forall x. mem \ x \neq mem' \ x \longrightarrow \text{writable } mds \ x) \wedge \\
196 \quad & (\forall x. \mathcal{L}_{mem} \ x \neq \mathcal{L}_{mem'} \ x \longrightarrow \text{writable } mds \ x) \longrightarrow \\
197 \quad & (\langle tps_A, mds, mem' \rangle_A, \langle tps_C, mds, mem' \rangle_C) \in \mathcal{R}
\end{aligned}$$

199 Note that closed-others could be viewed as a simplification of cg-consistent considering only
200 environmental actions that affect the memories on both sides of the relation identically, and
201 that furthermore ensure equality of *all* shared variables, not just those judged observable.

202 Then, regarding the maintenance of modes- and observational-equivalence across the
203 relation, the restrictions on refinement are somewhat tighter than those that applied to
204 strong-low-bisim-mm. The refinement relation \mathcal{R} is required to preserve the shared memory
205 in its entirety:

206 ► **Definition 5** (Preservation of modes and memory).

207 preserves-modes-mem $\mathcal{R} \equiv \forall lc_A lc_C. (lc_A, lc_C) \in \mathcal{R} \longrightarrow lc_A \stackrel{\text{mem}}{=}_{\text{mds}} lc_C$

208 In the general case, the final major requirement for CVDNI-preservation is to prove
209 \mathcal{R} closed under the pairwise executions of the concrete and abstract programs with the
210 aforementioned cube-shaped diagram (coupling-inv-pres, Figure 1) whose edges are given by
211 relations in \mathcal{B} , \mathcal{R} , and \mathcal{I} . All then that remains is for the nominated coupling invariant \mathcal{I} to
212 be symmetric, and the predicate secure-refinement puts together all the requirements:

213 ► **Definition 6** (Requirements for secure refinement of the per-thread CVDNI property).

214 secure-refinement $\mathcal{B} \mathcal{R} \mathcal{I} \equiv$

215 preserves-modes-mem $\mathcal{R} \wedge$ closed-others $\mathcal{R} \wedge$ cg-consistent $\mathcal{I} \wedge$ sym $\mathcal{I} \wedge$ coupling-inv-pres $\mathcal{B} \mathcal{R} \mathcal{I}$

216 Theorem 5.1 of our prior work [16] gives us that under the aforementioned conditions,

$$217 \quad \mathcal{B}_{\text{Cof}} \mathcal{B} \mathcal{R} \mathcal{I} \equiv \{(lc_{1C}, lc_{2C}) \mid \exists lc_{1A} lc_{2A}. (lc_{1A}, lc_{1C}) \in \mathcal{R} \wedge (lc_{2A}, lc_{2C}) \in \mathcal{R} \wedge$$

$$218 \quad \quad \quad (lc_{1A}, lc_{2A}) \in \mathcal{B} \wedge lc_{1C} \stackrel{\text{Low}}{=}_{\text{mds}} lc_{2C} \wedge (lc_{1C}, lc_{2C}) \in \mathcal{I}\}$$

220 is a witness strong-low-bisim-mm for the concrete program:

221 strong-low-bisim-mm $\mathcal{B} \wedge$ secure-refinement $\mathcal{B} \mathcal{R} \mathcal{I} \implies$ strong-low-bisim-mm $(\mathcal{B}_{\text{Cof}} \mathcal{B} \mathcal{R} \mathcal{I})$.

222 3 Decomposition principle for CVDNI-preserving refinement

223 We now present our first contribution: an alternative way of proving secure-refinement
224 (Definition 6) that does away with the use of the cube-shaped, two-sided refinement obligation
225 coupling-inv-pres $\mathcal{B} \mathcal{R} \mathcal{I}$ (depicted by Figure 1), by decomposing its concerns into (1) proving
226 \mathcal{R} closed under the pairwise executions of the concrete and abstract programs alone using a
227 square-shaped diagram (depicted by Figure 2a, which is akin to ordinary semantics-preserving
228 refinement), and (2) a number of smaller and more separable obligations gathered together
229 under the side-condition predicate decomp-refinement-safe.

► **Definition 7** (Decomposed requirements for CVDNI-preserving secure refinement).

230 secure-refinement-decomp $\mathcal{B} \mathcal{R} \mathcal{I} \text{ abs-steps} \equiv$

231 preserves-modes-mem $\mathcal{R} \wedge$ closed-others $\mathcal{R} \wedge$ cg-consistent $\mathcal{I} \wedge$ sym $\mathcal{I} \wedge$

232 decomp-refinement-safe $\mathcal{B} \mathcal{R} \mathcal{I} \text{ abs-steps} \wedge (\forall lc_A lc_C. (lc_A, lc_C) \in \mathcal{R} \longrightarrow$

$$233 \quad (\forall lc'_C. lc_C \rightsquigarrow_C lc'_C \longrightarrow (\exists lc'_A. lc_A \rightsquigarrow_A^{(\text{abs-steps } lc_A lc_C)} lc'_A \wedge (lc'_A, lc'_C) \in \mathcal{R})))$$

235 The decomposition requires the provision of a new refinement parameter that we will call
236 *abs-steps* or the *pacings function*, whose role is to dictate the pace of the refinement by
237 returning the number of abstract steps that ought to be taken for a single concrete step, for
238 a given abstract-concrete local configuration pair related by \mathcal{R} . The side-conditions on all of
239 the refinement parameters (depicted by Figures 2b, 2c) are then defined as follows:

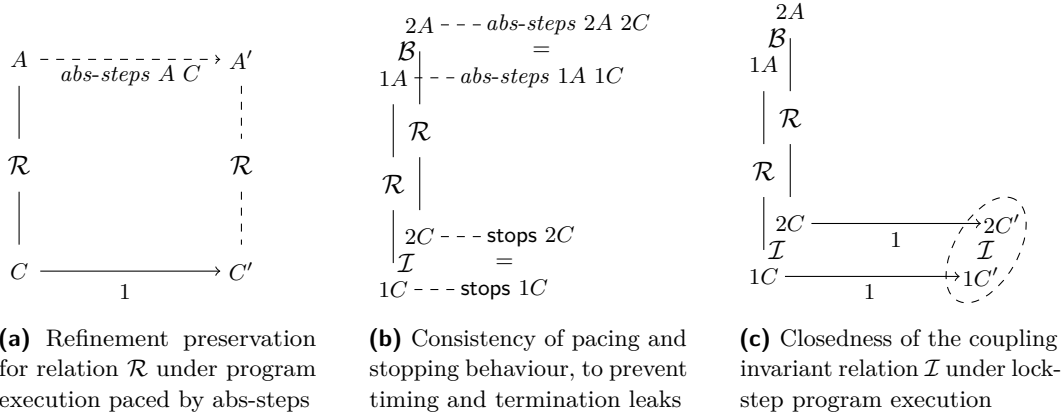
► **Definition 8** (Side-conditions for CVDNI-preserving refinement decomposition).

240 decomp-refinement-safe $\mathcal{B} \mathcal{R} \mathcal{I} \text{ abs-steps} \equiv \forall lc_{1A} lc_{2A} lc_{1C} lc_{2C}. (lc_{1A}, lc_{2A}) \in \mathcal{B} \wedge$

241 $lc_{1A} \stackrel{\text{mds}}{=} lc_{2A} \wedge (lc_{1A}, lc_{1C}) \in \mathcal{R} \wedge (lc_{2A}, lc_{2C}) \in \mathcal{R} \wedge (lc_{1C}, lc_{2C}) \in \mathcal{I} \wedge lc_{1C} \stackrel{\text{mds}}{=} lc_{2C}$

242 \longrightarrow stops $lc_{1C} =$ stops $lc_{2C} \wedge$ abs-steps $lc_{1A} lc_{1C} =$ abs-steps $lc_{2A} lc_{2C} \wedge$

$$243 \quad (\forall lc'_{1C} lc'_{2C}. lc_{1C} \rightsquigarrow_C lc'_{1C} \wedge lc_{2C} \rightsquigarrow_C lc'_{2C} \longrightarrow (lc'_{1C}, lc'_{2C}) \in \mathcal{I} \wedge lc'_{1C} \stackrel{\text{mds}}{=} lc'_{2C})$$



■ **Figure 2** Graphical depictions of refinement decomposition obligations

245 On the intuitive meaning of the side-conditions in Definition 8:

- 246 ■ **stops** $lc_{1C} = \text{stops } lc_{2C}$ ensures that the refinement has not introduced any termina-
247 tion leaks, by asserting *consistent stopping behaviour* for \mathcal{I} -related concrete program
248 configurations, which we know to be observationally indistinguishable.
 - 249 ■ *abs-steps* $lc_{1A} lc_{1C} = \text{abs-steps } lc_{2A} lc_{2C}$ ensures that the refinement has not introduced
250 any timing leaks, by asserting *consistency of the pace of the refinement* for \mathcal{R} -related
251 program configurations, which we again know to be observationally indistinguishable.
 - 252 ■ The final \forall -quantified clause asserts \mathcal{I} 's suitability as a coupling invariant, in that it must
253 remain *closed under lockstep evaluation* of the concrete program configurations it relates.
254 Furthermore it must *maintain mode state equality* with each lockstep evaluation, which
255 ensures that the refinement has not introduced any inconsistencies in the memory access
256 assumptions and guarantees needed for the concurrent compositionality of the property.
- 257 Note that the availability of all of \mathcal{B} , \mathcal{R} , and \mathcal{I} grants the ability to leverage facts known
258 about a particular program verification technique or compiler (as captured by \mathcal{B} and \mathcal{R} ,
259 respectively), in the proofs of the side conditions even just on \mathcal{I} alone.

260 Assuming the fulfilment of all of the decomposed requirements, we obtain that they are a
261 sound method for establishing secure refinement of the per-thread CVDNI property:

► **Theorem 9** (Soundness of secure-refinement-decomp).

$$262 \text{ secure-refinement-decomp } \mathcal{B} \mathcal{R} \mathcal{I} \text{ abs-steps} \implies \text{secure-refinement } \mathcal{B} \mathcal{R} \mathcal{I}$$

264 In the interests of brevity we relegate proof sketches for all results to Appendix B, and for
265 fuller details we refer the reader to our Isabelle/HOL formalisation.

266 We now devote our attention to two applications of this new decomposition principle.
267 See Section 4 for an instantiation for the proof of CVDNI-preservation for the refinement of
268 a program that branches on a secret. See Section 5.5 for an instantiation for the proof of
269 CVDNI-preservation by a compiler.

270 4 Proof effort comparison

271 To demonstrate how the decomposition principle reduces proof complexity and effort, we
272 returned to the example refinement discussed in Section V-E of our previous work [16], of a
273 program that branches on a sensitive value. We use proof size as a proxy for proof effort,

274 since the former is known to be strongly linearly correlated with the latter [20]. Formalised
 275 in Isabelle/HOL as `EgHighBranchRevC.thy` [15], the proof line count for that theory stood
 276 at about 4.8K lines of definitions and proof, of which approx. 3.6K line were proofs.

277 Adapting the proof instead to use the decomposition principle `secure-refinement-decomp`
 278 (Definition 7), the proof line count drops from 3.6K to approx. 2K, a 46% reduction. As
 279 would be expected, the bulk of the deletions are from the full cube-shaped refinement diagram
 280 proof (Figure 1) of `secure-refinement` (Definition 6) for that example refinement, the surviving
 281 parts of which just become the square-shaped refinement diagram proof (Figure 2a) of
 282 `secure-refinement-decomp` without much modification. The deletions are replaced by the
 283 newly added proofs of the three sub-obligations of `decomp-refinement-safe` (Definition 8).

284 **5 The COVERN `wr`-compiler**

285 Having presented our new decomposition principle for CVDNI-preserving refinement, we now
 286 turn to our compiler, whose most notable features for formal proof of secure refinement are:

- 287 1. Its implementation tracks variable stability (Section 5.4) responsive to use of locking
 288 primitives, to know when accesses to shared variables are safe to optimise, and when
 289 register contents can be still be considered consistent with shared variable contents.
- 290 2. Its verification uses a pacing function (Section 5.5.2) and coupling invariant (Section 5.5.3)
 291 as the decomposition demands, to ensure it does not introduce timing leaks.

292 First, we describe its source and target languages, and parameters to the compilation.

293 **5.1 Source language**

294 The COVERN `wr`-compiler—short for *While-to-RISC compiler*—takes the simple imperative
 295 language with while-looping and lock-based synchronisation targeted by the COVERN program
 296 logic [14], which we will refer to as `While`, consisting of the commands `cmd`:

$$\begin{aligned}
 &exp ::= n \mid v \mid exp \oplus exp \\
 &cmd ::= \mathbf{skip} \mid cmd ; cmd \mid \mathbf{if} \ exp \ \mathbf{then} \ cmd \ \mathbf{else} \ cmd \mid \\
 &\quad \mathbf{while} \ exp \ \mathbf{do} \ cmd \mid v := exp \mid \\
 &\quad \mathbf{lock}(k) \mid \mathbf{unlock}(k)
 \end{aligned}$$

298 The language is parameterised over a type of values Val , and binary operators $\oplus :: Val \Rightarrow$
 299 $Val \Rightarrow Val$. Constants $n \in Val$; $v \in Var$ and $k \in Lock$ are (resp.) shared program- and
 300 lock-variables. The semantics of the locking primitives `lock(k)` and `unlock(k)` is informed
 301 by a locking discipline provided by the user of the theory as a parameter (see Section 5.3.2).
 302 We leave for future work adding support for pointers and arrays, which we believe will be
 303 straightforward because our assume-guarantee framework already provides the means to
 304 encode the memory footprint of a command in a way that depends on values in memory.

305 We assume that the underlying concurrent execution model (e.g. OS, scheduler) for the
 306 `While` language prevents threads from seeing each others' current program location, and
 307 thus (as in previous work [16, 13]) the `While` program command $c \in cmd$ still remaining
 308 to be executed we model as thread-private state: $\langle c, mds, mem \rangle_w$. In contrast, all program
 309 variables $v \in Var$ and lock variables $k \in Lock$ reside in the shared memory.

310 5.2 Target language

311 The `wr`-compiler's target is a generic RISC-style assembly language like that of [21] except
 312 with lock-based synchronisation primitives added, which we will refer to as RISC:

$$\begin{aligned}
 I &::= [l:]B \\
 B &::= \mathbf{Load} \ r \ v \mid \mathbf{Store} \ v \ r \mid \mathbf{Jump} \ l \mid \mathbf{Jz} \ l \ r \mid \mathbf{Nop} \\
 &\quad \mathbf{MoveK} \ r \ n \mid \mathbf{MoveR} \ r \ r' \mid \mathbf{Op} \ \oplus \ r \ r' \\
 &\quad \mathbf{LockAcq} \ k \mid \mathbf{LockRel} \ k
 \end{aligned}$$

314 The language is parameterised over the same value type Val and binary operators \oplus ,
 315 shared program variables $v \in Var$ and shared lock variables $k \in Lock$ as the `While` language.
 316 For the time being, direct-addressing **Load** and **Store** instructions are adequate for RISC to
 317 implement all existing `While` features, and we expect adding indirect addressing to RISC to
 318 be as straightforward as adding pointer and array support to `While`.

319 RISC program texts (typically named) P are just lists of binary instructions I , each
 320 optionally associated with a label $l \in Lab$. We assume that the underlying concurrency
 321 model for the RISC language (e.g. OS, scheduler etc.) prevents one thread from reading the
 322 program code (instructions) of another, as well as another's registers (including the program
 323 counter). Thus, we model the distinguished program counter register pc , program text P ,
 324 and register bank $regs$ as thread-private state: $\langle (pc, P), regs \rangle_r$.

325 Apart from its adaptation in this manner to the format of our local configuration triples,
 326 its evaluation semantics follows that of the RISC target of [21]. This semantics assumes
 327 sequentially consistent execution.

328 5.3 Compilation parameters

329 We leave certain concerns underspecified as parameters to the theory, but rely on the
 330 user-provided parameters to fulfil a few modest requirements.

331 5.3.1 Register allocation scheme

332 We generalise over the register allocation scheme, which we model by two functions reg_alloc
 333 and reg_alloc_cached on the *register record* Φ (see Section 5.4) and the set A of registers
 334 currently in use by the program. In order to avoid loading from memory unnecessarily,
 335 the compiler may first call $reg_alloc_cached \ \Phi \ A \ v$ to identify whether there is an unused
 336 register (i.e. not in A) that Φ records as already containing the variable v . When the compiler
 337 needs a fresh register, it will call $reg_alloc \ \Phi \ A$ to allocate an unused register. Since neither
 338 function may return a register that is in A , the allocator is permitted to fail if it runs out
 339 of registers. Spilling of register contents to memory is not addressed here, and we leave a
 340 treatment of register spilling to future work.

341 5.3.2 Locking discipline

342 Like the `COVERN` logic [14], we assume that the `While` language program being compiled
 343 follows a certain locking discipline, about which the compiler has knowledge, so as to ensure
 344 that the RISC program it produces follows the same discipline.

345 The user of the theory provides the necessary details of the locking discipline in the form
 346 of a *lock interpretation* parameter: $lock_interp :: Lock \Rightarrow (Var \ set \times Var \ set)$, which for each
 347 lock gives the two non-overlapping sets of program variables over which acquiring the lock
 348 grants exclusive permission to write, (resp.) read and write. These permissions are then

349 reflected in the semantics of the `While` and RISC locking primitives such that they act on the
 350 mode state to achieve the following: a thread that does not possess a lock k guarantees it
 351 does not write (resp. read or write) program variables in the `fst` (resp. `snd`) of *lock-interp* k ,
 352 and conversely, a thread that possesses lock k assumes that no other threads access those
 353 program variables in that same manner.

354 Regarding lock interpretations and the way they interact with the user-provided value-
 355 dependent classification function \mathcal{L} (see Section 2.1), we inherit a few cleanliness conditions
 356 from that earlier work [14], chief of which are that lock variables k cannot be control variables,
 357 a lock variable k governing access to a program variable v must govern the same kind of
 358 access to all of v 's control variables, and \mathcal{L} must classify all lock variables as `Low`.

359 Note finally that not all program variables are required to have their access governed by
 360 a lock, and that in fact the choice should be made carefully with respect to the intended
 361 security outcomes. Intuitively, the user of the theory should typically model the permanent
 362 untrusted output sinks of the whole concurrent program as program variables in shared
 363 memory for which \mathcal{L} *always returns* `Low`. Furthermore, they should ensure that access to
 364 such variables is *not governed by any lock*, i.e. such variables are absent from *lock-interp*,
 365 since in general the attacker cannot be assumed to follow the locking protocol.

366 5.4 Compiler implementation and tracking of shared variable stability

367 We chose as a starting point the *fault-tolerant noninterference*-preserving compilation scheme
 368 of [21], on the basis of their preserving a property similar to CVDNI with respect to
 369 resilience to changes made by an environment. Aiming to repurpose that for concurrent
 370 compositionality, we adapt it to Isabelle, implementing it as a primitive recursive function:

371 $\text{compile-cmd} :: \text{CompRec} \Rightarrow \text{Lab option} \Rightarrow \text{Lab} \Rightarrow \text{cmd} \Rightarrow$
 372 $(I \times \text{CompRec}) \text{ list} \times \text{Lab option} \times \text{Lab} \times \text{CompRec} \times \text{bool}$
 373

374 where we choose $\text{Lab} = \text{nat}$ for RISC instruction labels, and the *compilation record* type
 375 CompRec is bookkeeping maintained by the compiler that we will describe further below.

376 A typical invocation to compile a `While` program $c \in \text{cmd}$ takes the form:

377 $(\text{PCs}, l', nl', C', \text{failed}) = \text{compile-cmd } C \ l \ nl \ c$ (1)

378 Here, `compile-cmd` takes an *initial compilation record* C , an optional *entry label* l , and the
 379 *next available label* nl , and for the benefit of the next invocation returns an optional *exit*
 380 *label* l' if one is used by the program just compiled, the *new next available label* nl' , and a
 381 *final compilation record* C' . We leave discussion of the label allocation scheme and its impact
 382 on achieving sequential composability for the compiled RISC programs to Appendix B.2.

383 In addition to the output RISC program $P \in I \text{ list}$ itself, a call to `compile-cmd` also
 384 outputs every CompRec associated with the state of the program just before executing every
 385 instruction in P . These are returned zipped up together with P as the *CompRec-annotated*
 386 *RISC program* $\text{PCs} \in (I \times \text{CompRec}) \text{ list}$. (P can trivially be recovered as `map fst PCs`.)

387 Finally, `compile-cmd` may return `True` for *failed* to reject the input program in various
 388 circumstances, such as propagating a failure by the register allocator (see Section 5.3.1).

389 In the style of the compilation scheme on which it was based [21], the *wr-compiler*
 390 maintains a *register record* (typically named) $\Phi \in \text{reg} \rightarrow \text{exp}$, i.e. a partial map of registers
 391 to expressions on shared variables. In addition to using it to compile away any unnecessary
 392 loads from variables in shared memory, we also use it to ensure that an expression calculated
 393 by RISC in registers is equal to the value of the expression as if it had all been calculated by

394 **While** in one step. This is especially important when writing the result of an expression back
 395 to shared memory, because the refinement is required to maintain all shared memory values.

396 Additionally, we have added to the **wr-compiler** the responsibility of maintaining an
 397 *assumption record*, which it uses primarily to assert and maintain the absence of data races
 398 on shared memory. Each assumption record (typically named) $\mathcal{S} \in (Var\ set \times Var\ set)$ is a
 399 pair tracking the set of variables on which (resp.) **AsmNoW**, **AsmNoRW** assumptions
 400 are currently active at a given point in the program being compiled. As a secondary concern
 401 we also use it to assert that the two sides of any if-conditional branches act consistently on
 402 the mode state, and that while-loops restore the original mode state on termination.

403 A compilation record e.g. $C = (\Phi, \mathcal{S}) \in CompRec$ is then just a pair of a register record
 404 and an assumption record. For readability, we will use **regrec** and **asmrec** respectively to
 405 denote a *CompRec*'s **fst** and **snd** projections.

406 To explain how the compilation record is used to prevent data races, and to ensure
 407 consistency of expression evaluation between source and target program, firstly we must
 408 introduce the concept of *stability* of a variable v according to an assumption record \mathcal{S} :

$$409 \quad \text{var-stable } \mathcal{S} \ v \equiv v \in (\text{fst } \mathcal{S} \cup \text{snd } \mathcal{S}) \wedge (\forall v' \in Cvars \ v. \ v' \in (\text{fst } \mathcal{S} \cup \text{snd } \mathcal{S}))$$

410 In short, this means that the variable and all its control variables ($Cvars \ v$) are recorded as
 411 having either of **AsmNoW** or **AsmNoRW** active on them.

412 For register record entries to be of any help in ensuring consistency of **While** and **RISC**
 413 expression evaluation, we exclude expression evaluation on data race-prone variables by
 414 lifting the concept of stability to register records. The following predicate asserts internal
 415 consistency of the compilation record C created by **compile-cmd**, in the sense that the register
 416 record may only map to expressions that mention variables that are recorded as **stable** by
 417 the assumption record accompanying it. (Here, **ran** denotes the *range* of a map.)

$$418 \quad \text{regrec-stable } C \equiv \forall e \in \text{ran} \ (\text{regrec } C). \ (\forall v \in \text{exp-vars } e. \ \text{var-stable} \ (\text{asmrec } C) \ v)$$

419 To ensure that an input **While** program maintains register record stability, we define
 420 the predicate **no-unstable-exprs** $c \ C$ to capture the requirement that a program c if started
 421 with a configuration consistent with compilation record C will never access a lock-protected
 422 variable without holding the relevant lock. (On the side, it also uses the assumption record
 423 to check the secondary concerns that if-conditional branches and while-loops keep consistent,
 424 resp. restore the mode state on termination.) We implement this predicate as a simple static
 425 check carried out by a primitive recursive function on the structure of **While** programs.

426 Together, **regrec-stable** and **no-unstable-exprs** make up the main two requirements of a
 427 predicate **compile-cmd-input-reqs** $C \ l \ nl \ c$ imposed on the input arguments to **compile-cmd**,
 428 which gives us enough information to prove a lemma that **compile-cmd** only ever outputs
 429 stable register records. Full details of these we leave to Appendix B.

430 5.5 Proof of CVDNI-preserving compilation

431 We can now present the refinement relation \mathcal{R}_{wr} (Section 5.5.1), pacing function **abs-steps**_{wr}
 432 (Section 5.5.2), and coupling invariant \mathcal{I}_{wr} (Section 5.5.3) that we use with our new decomposi-
 433 tion principle (of Section 3) to prove that the **wr-compiler** preserves CVDNI (Section 5.5.4).

434 5.5.1 Refinement relation \mathcal{R}_{wr} and its invariants

435 For full details of \mathcal{R}_{wr} , we refer the reader to the Isabelle formalisation. We provide an
 436 informal description of all of its cases, with emphasis on their purpose and on the invariants
 437 they maintain, in Appendix A. Here, we will focus on its most notable characteristics for the
 438 purpose of understanding why it is suitable to describe a CVDNI-preserving compilation.

439 To serve as an illustrative example, we present for comparison the case of the implementa-
 440 tion of `compile-cmd` for the `While` program `if e then c1 else c2`, alongside the corresponding
 441 part of the refinement relation \mathcal{R}_{wr} , namely the introduction rule $\mathcal{R}_{wr}.if_expr$. Both have
 442 been adapted slightly to improve the clarity of the presentation. In the code listing, $\Phi \sqcap_R \Phi'$
 443 denotes the subset of mappings on which Φ and Φ' agree.

■ **Listing 1** Implementation of `compile-cmd` case for `if e then c1 else c2`

```

444 compile_cmd C l nl (If e c1 c2) =
445   (let (Pe, r, C1, faile) = (compile_expr C {} l e);
446       (br, nl') = (nl, Suc nl); (ex, nl'') = (nl', Suc nl');
447       (P1, l1, nl1, C2, fail1) = (compile_cmd C1 None nl'' c1);
448       (P2, l2, nl2, C3, fail2) = (compile_cmd C1 (Some br) nl1 c2);
449       (* Pre-compilation check ensures asmrec C2 = asmrec C3 *)
450       C' = (regrec C2  $\sqcap_R$  regrec C3, asmrec C2)
451   in (Pe @ [((if Pe = [] then l else None, Jz br r), C1)] @
452       P1 @ [((l1, Jmp ex), C2)] @ P2 @ [((l2, Nop'), C3)] ,
453       Some ex, nl2, C', faile  $\vee$  fail1  $\vee$  fail2))
454

```

456 The `if_expr` case of \mathcal{R}_{wr} relates the expression evaluation part of `if e then c1 else c2`
 457 with the corresponding part (including the conditional jump `Jz` after expression evaluation)
 458 of the RISC program obtained by running `compile-cmd` on it. (Variables ignored are in gray.)

► **Example 10** (Introduction rule for case `if_expr` of \mathcal{R}_{wr}).

$$\begin{aligned}
 & c = \text{if } e \text{ then } c_1 \text{ else } c_2 \wedge \text{compile-cmd-input-reqs } C \ l \ nl \ c \wedge \\
 & (PCs, l', nl_2, C', \text{False}) = \text{compile-cmd } C \ l \ nl \ c \wedge (P_e, r, C_1, \text{False}) = \text{compile-expr } C \ \emptyset \ l \ e \wedge \\
 & (P_1, l_1, nl_1, C_2, \text{False}) = \text{compile-cmd } C_1 \ \text{None} \ (\text{Suc} \ (\text{Suc} \ nl)) \ c_1 \wedge pc \leq \text{length } P_e \wedge \\
 & (P_2, l_2, nl_2, C_3, \text{False}) = \text{compile-cmd } C_1 \ (\text{Some } nl) \ nl_1 \ c_2 \wedge C_{pc} = (\text{map } \text{snd } PCs \ ! \ pc) \wedge \\
 & \text{compiled-cmd-config-consistent } C_{pc} \ \text{regs} \ \text{mds} \ \text{mem} \wedge \text{regrec-stable } C_{pc} \wedge \\
 & (\forall \text{mds}' \ \text{mem}' \ \text{regs}'. \text{compiled-cmd-config-consistent } C_1 \ \text{regs}' \ \text{mds}' \ \text{mem}' \wedge \text{regrec-stable } C_1 \\
 & \quad \longrightarrow (\langle c_1, \text{mds}', \text{mem}' \rangle_w, \langle ((0, \text{map } \text{fst } P_1), \text{regs}'), \text{mds}', \text{mem}' \rangle_r) \in \mathcal{R}_{wr} \wedge \\
 & \quad (\langle c_2, \text{mds}', \text{mem}' \rangle_w, \langle ((0, \text{map } \text{fst } P_2), \text{regs}'), \text{mds}', \text{mem}' \rangle_r) \in \mathcal{R}_{wr}) \implies \\
 & (\langle c, \text{mds}, \text{mem} \rangle_w, \langle ((pc, \text{map } \text{fst } PCs), \text{regs}), \text{mds}, \text{mem} \rangle_r) \in \mathcal{R}_{wr}
 \end{aligned}$$

469 This is a fairly typical case of \mathcal{R}_{wr} in a number of respects:

470 Firstly, there is a direct reference to the call to `compile-cmd` for the given `While` program.
 471 Secondly, various guards (`compiled-cmd-config-consistent` introduced below, and `regrec-stable`
 472 defined in Section 5.4) are asserted in order to restrict the scope of \mathcal{R}_{wr} only to consider
 473 wellformed local program configurations that line up with the conditions captured by the
 474 compilation record. Thirdly, the inductive references to \mathcal{R}_{wr} for P_1 and P_2 , the branches of
 475 the conditional *that have not been reached yet*, are quantified over all configurations that
 476 obey the guards `compiled-cmd-config-consistent` and `regrec-stable` relative to C_1 , the initial
 477 compilation record for each of the sub-calls to `compile-cmd` for those sub-programs.

478 The guard `compiled-cmd-config-consistent` mentioned above asserts that the compilation
479 record C is consistent with the registers $regs$, memory mem and mode state mds .

$$\begin{aligned}
480 \quad & \text{compiled-cmd-config-consistent } C \text{ } regs \text{ } mds \text{ } mem \equiv \\
481 \quad & (\forall r \ e. (\text{regrec } C) \ r = \text{Some } e \longrightarrow regs \ r = \text{ev}_{\text{exp}} \ mem \ e) \wedge \\
482 \quad & \text{asmrec } C = (mds \ \mathbf{AsmNoW}, \ mds \ \mathbf{AsmNoRW})
\end{aligned}$$

484 Firstly, for all entries in register record mapping some register r to some expression e , the
485 value held in r of the register bank $regs$ must match the value of e if evaluated under memory
486 mem . Secondly, the assumption record must consist exactly of the program variables the
487 mode state mds says have \mathbf{AsmNoW} , $\mathbf{AsmNoRW}$ on them respectively.

488 As we will see in Theorem 17, `compiled-cmd-config-consistent` also serves as *initial config-*
489 *uration requirements* for compiled programs: only configurations obeying them may be used
490 to initialise a RISC program compiled by the `wr-compiler` with initial compilation record C .

491 With \mathcal{R}_{wr} specified, we then prove the two requirements for `secure-refinement-decomp` that
492 pertain to \mathcal{R}_{wr} alone: `preserves-modes-mem` (Definition 5) and `closed-others` (Definition 4).

493 ▶ **Lemma 11** (\mathcal{R}_{wr} preserves modes and memory). `preserves-modes-mem` \mathcal{R}_{wr}

494 ▶ **Lemma 12** (\mathcal{R}_{wr} is closed under changes by others). `closed-others` \mathcal{R}_{wr}

495 5.5.2 Refinement pacing function abs-steps_{wr}

496 We now nominate an *abs-steps* function, determining the pace at which `While` programs
497 progress in comparison to the RISC programs that they are compiled to by the `wr-compiler`.

498 To assist here and elsewhere, we define a primitive recursive helper `leftmost-cmd` that
499 given a sequence of `;`-separated `While` commands, strips all but the first: given $c_1 ; c_2$ it
500 returns `leftmost-cmd` c_1 , and given any other `While` program c it returns c .

501 Our pacing function abs-steps_{wr} primarily looks at the form of the RISC program instruc-
502 tion about to be executed. The RISC instructions are divided into three categories:

- 503 ■ Instructions output by `compile-expr`: **Load**, **Op**, and **MoveK**. For those instructions,
504 abs-steps_{wr} returns 1 if the `leftmost-cmd` of the `While` program is some **while** exp **do** cmd ,
505 and 0 otherwise – the 0 indicates the `While` program standing still while the RISC program
506 takes new steps executing the expression evaluation. Returning 1 accounts for the se-
507 mantics of **while** exp **do** cmd taking 1 step to **if** exp **then** (cmd ;**while** exp **do** cmd) **else stop**
508 prior to evaluating its conditional expression, which abs-steps_{wr} allows to occur concur-
509 rently with the first RISC step of the compiled expression itself.
- 510 ■ “Epilogue” steps: **Jmp** and **Nop** when used for control flow at the end of a smaller
511 compiled program in the context of a larger one. For these, abs-steps_{wr} returns 0.
- 512 ■ All other RISC instructions are assumed to proceed at a lockstep pace with the `While`
513 command they were compiled from, and for these abs-steps_{wr} returns 1.

514 Having nominated abs-steps_{wr} and \mathcal{R}_{wr} , we now have the parameters over which we are
515 obliged to prove refinement preservation (Figure 2a) as demanded by `secure-refinement-decomp`
516 (Definition 7). To this end, we prove firstly (elided to Appendix B) that every step of
517 execution of a RISC program produced by the `wr-compiler` from a `While` program, maintains
518 the consistency demanded by `compiled-cmd-config-consistent` between configurations and
519 compilation records. Also, we must prove a correctness lemma for the expression compiler:

520 ▶ **Lemma 13.** $(PCs, r, C', \text{False}) = \text{compile-expr } C \ A \ l \ e \implies (\text{regrec } C') \ r = \text{Some } e$

521 Armed with these facts, we can now prove the main refinement preservation result:

► **Lemma 14** (\mathcal{R}_{wr} is a refinement paced by abs-steps_{wr}).

$$522 \quad \forall lc_w \ lc_r. (lc_w, lc_r) \in \mathcal{R}_{wr} \longrightarrow (\forall lc'_w \ lc'_r. lc_w \rightsquigarrow_r lc'_w \longrightarrow$$

$$523 \quad (\exists lc'_w. lc_w \rightsquigarrow_w^{(\text{abs-steps}_{wr} \ lc_w \ lc_r)} lc'_w \wedge (lc'_w, lc'_r) \in \mathcal{R}_{wr}))$$

525 5.5.3 Coupling invariant \mathcal{I}_{wr}

526 The coupling invariant \mathcal{I}_{wr} is defined as follows:

$$527 \quad \mathcal{I}_{wr} \equiv \{(((pc, P), regs), mds, mem)_r, (((pc', P'), regs'), mds', mem')_r \mid (pc, P) = (pc', P')\}$$

528 In other words, \mathcal{I}_{wr} asserts that we only need compare local configurations that are at
529 the same location $pc = pc'$ of the same RISC program $P = P'$. The effect of this is to assert
530 that the wr -compiler has not introduced any *new* branching on sensitive values.

531 5.5.4 Successful compilations are CVDNI-preserving refinements

532 We specify that a relation \mathcal{B} can describe only **While**-programs with no branching on sensitive
533 (**High-classified**) values (if used as a witness **strong-low-bisim-mm** for **com-secure**), as follows:

534 no-high-branching $\mathcal{B} \equiv$

$$535 \quad \forall c \ c' \ mds \ mem \ mem'. (\langle c, mds, mem \rangle_w, \langle c', mds, mem' \rangle_w) \in \mathcal{B} \longrightarrow c = c' \wedge$$

$$536 \quad (\forall e \ c_1 \ c_2. \text{leftmost-cmd } c = \text{if } e \ \text{then } c_1 \ \text{else } c_2 \longrightarrow \text{ev}_{\text{exp}} \ mem \ e = \text{ev}_{\text{exp}} \ mem' \ e)$$

538 That is, it refuses to relate configurations at different program locations, meaning (if imposed
539 on a **strong-low-bisim-mm**) that a branch on a sensitive value could not have occurred at
540 some point in the past. Furthermore if it is at a conditional branching point, the expression
541 e determining which branch will be taken evaluates to the same boolean value for both
542 configurations' memories. By imposing this condition on a relation that already ensures
543 **Low-equivalent** memory modulo modes, this effectively disallows any branching on values
544 that are **High-classified** or otherwise hidden to other threads. Then, for such programs:

545 ► **Lemma 15.** **strong-low-bisim-mm** $\mathcal{B} \wedge$ **no-high-branching** $\mathcal{B} \implies$
546 **decomp-refinement-safe** $\mathcal{B} \ \mathcal{R}_{wr} \ \mathcal{I}_{wr} \ \text{abs-steps}_{wr} \wedge$
547 **secure-refinement-decomp** $\mathcal{B} \ \mathcal{R}_{wr} \ \mathcal{I}_{wr} \ \text{abs-steps}_{wr}$

548 From this it follows immediately via Theorem 9 that \mathcal{R}_{wr} with the help of \mathcal{I}_{wr} describes a
549 CVDNI-preserving refinement for non-**High-branching** **While** programs:

► **Corollary 16** (\mathcal{R}_{wr} is a CVDNI-preserving refinement for non-**High-branching** programs).

$$550 \quad \text{strong-low-bisim-mm } \mathcal{B} \wedge \text{no-high-branching } \mathcal{B} \implies \text{secure-refinement } \mathcal{B} \ \mathcal{R}_{wr} \ \mathcal{I}_{wr}$$

552 Finally, we prove that successful compilation produces a RISC program related by \mathcal{R}_{wr} to its
553 input **While** program, when started with corresponding and reasonable initial configurations:

► **Theorem 17** (Successful compilations are refinements in \mathcal{R}_{wr}).

$$554 \quad (PCs, l', nl', C', failed) = \text{compile-cmd } C \ l \ nl \ c \wedge \text{compile-cmd-input-reqs } C \ l \ nl \ c \wedge$$

$$555 \quad failed = \text{False} \wedge \text{compiled-cmd-config-consistent } C \ regs \ mds \ mem \wedge P = \text{map fst } PCs$$

$$556 \quad \implies (\langle c, mds, mem \rangle_w, (((0, P), regs), mds, mem)_r) \in \mathcal{R}_{wr}$$

6 Case study: the `wr-compiler` in action

To test the theory, we instantiated it and applied the `wr-compiler` to a `While`-language model of the Cross Domain Desktop Compositor [4] (CDDC), a non-trivial concurrent program enforcing information-flow control on input classified value-dependently on other input—this facilitates a trusted user’s interaction with multiple desktop machines of differing clearance.

The CDDC model to which we applied the compiler is a 2-thread program that was a precursor to the 3-thread model that was verified using the COVERN program logic [14]. Each of the `While` threads of the CDDC program (together about 150 lines of `While`) we proved satisfy the compositional security property `com-secure` (Definition 2), using a precursor to the COVERN logic that yields CVDNI-witness bisimulations that are non-High-branching.

The resulting compiler is *executable* in Isabelle, meaning that `compile-cmd` can be executed on the `While` program text for each of the two threads to obtain their RISC compilations (together totalling about 250 RISC instructions) using the Isabelle tactic `eval`.

The secure compilation theorems (of Section 5.5.4), together with `strong-low-bisim-mm` preservation and compositionality for `com-secure` (Theorems 5.1, 3.1 of [16], mentioned earlier in Sections 2.2, 2.1 respectively) then allow us to derive that the compiled program is secure when its threads are run concurrently. (For full details, see the Isabelle/HOL formalisation.)

To our knowledge this is the first proof of source-level information-flow security being carried by a verified compiler to an assembly-level model of a non-trivial concurrent program.

7 Related work

A number of works, like ours, focus on compilation that preserves a noninterference property. Tedesco et al. [21] present a type-directed compilation scheme that preserves a *fault-resilient* noninterference property. The compilation scheme of our `wr-compiler` was inspired by theirs. Like our `com-secure` CVDNI security property that `wr-compiler` preserves, Tedesco et al.’s security property is also *strong bisimulation*-based [19]. However, whereas our CVDNI property accounts for *controlled interference* by other threads, via mode states, theirs instead quantifies over all possible interference by the environment with the memory contents. While doing so simplifies their task of proving that their security property is preserved under compilation—since it need not require that the compiler preserves the contents of memory—it makes their security property unable to capture notions like value-dependent noninterference, which ours can. In contrast, our `wr-compiler` must obey the requirement imposed by our notion of secure refinement that memory contents are preserved, and when adapting Tedesco et al.’s expression compiler, we found and fixed a bug (acknowledged in private correspondence) whereby registers in use would be incorrectly reallocated. This bug caused expressions like $v + (v + 1)$ to be compiled incorrectly to programs yielding the value of $(v + 1) + (v + 1)$ instead, causing a violation of memory contents preservation.

Barthe et al. [1] consider the problem of preserving *cryptographic constant-time policies* under compilation. These are a class of noninterference properties similar to CVDNI in its explicit consideration for capturing timing-sensitivity. Barthe et al. consider a number of common categories of compile-time optimisations, and mechanise proofs in Coq that such optimisations preserve various constant-time security properties. The scope of optimisations they consider is wider than those performed by our `wr-compiler`. Specifically, in our `While` language, all variables are shared. This severely limits the scope of optimisations to those that the compiler can perform when it knows that a shared variable is stable because it has been locked. At present, our `wr-compiler` avoids redundant loads during expression compilation, but other optimisations like loop hoisting and constant folding we are yet to implement. Their

604 proof technique, *constant-time simulation* (or *CT-simulation*) for preserving constant-time
 605 policies under refinement was developed independently to our original cube-shaped secure
 606 refinement definition [16]. Like ours, theirs is also a cube-shaped obligation and makes use of
 607 a pacing function analogous in their setting to our *abs-steps*. Unlike our work here, Barthe
 608 et al. do not give a general method for decomposing their cube-shaped simulation diagrams.

609 Neither of the works cited above consider per-thread compositional compilation of con-
 610 current, shared memory programs, nor value-dependent noninterference policies, which are
 611 the focus of our theory and compiler.

612 Patrignani et al. [17] prove that *trace-preserving compilation* preserves k -safety hyper-
 613 properties [5], including noninterference properties. However, it disallows compilers from
 614 eliminating any entries from traces. Thus for traces that need to reflect the passage of time
 615 as observable to other threads, as needed for security properties like our CVDNI com-secure
 616 to be compositional for concurrent programs, it is too restrictive to allow any optimising
 617 compilation that changes the pace of the program. This would exclude optimisations carried
 618 out by our compiler (which permits changes to pacing regulated by *abs-steps*) and studied
 619 by the two other works on timing-sensitive security-preserving compilation mentioned above.

620 Other work has aimed at preserving noninterference of multithreaded programs by
 621 compilation. Specifically, Barthe et al. [3] do so by extending a prior (security) *type-preserving*
 622 compilation approach [2]. This work differs from ours in two crucial ways. Firstly, unlike
 623 CVDNI, their noninterference property is termination-insensitive and timing-insensitive.
 624 Thus, their approach for preventing internal timing leaks introduced by compilation requires
 625 that the scheduler disallows certain interleavings between threads. Secondly, their security
 626 property like many mentioned above is not value-dependent. Nevertheless, they establish
 627 their security property for a compiled program using a type-preservation argument derived
 628 from a proof of (big-step) semantics preservation for their compiler. Here we also lean on a
 629 form of semantics preservation, in our case preservation of small-step semantics (specifically
 630 memory contents). This is necessary for us to avoid imposing non-standard requirements on
 631 the scheduler, as well as to preserve value-dependent security under compilation.

632 Finally, there has been much recent work on large-scale verified compilation [10, 9] some
 633 of which has also treated compilation of shared-memory concurrent programs [11] including
 634 taking weak-memory consistency into account [18]. Our work here does not consider the
 635 effects of weak-memory models. However, it differs to prior work on verified concurrent
 636 compilation, in that it formalises and proves a compiler’s ability to use information about the
 637 application’s locking protocol, to exclude unsafe access to shared variables, and conversely to
 638 know when it is safe to allow optimisations that would typically be excluded (see Section 5.4).

639 **8 Conclusion**

640 To our knowledge, we have presented the first mechanised verification that a compiler
 641 preserves concurrent, value-dependent noninterference. To this end, we provided a general
 642 decomposition principle for compositional, secure refinement. Although our compiler is a
 643 proof-of-concept targeting simple source and target languages, we nevertheless applied it to
 644 produce a verified assembly-level model of the CDDC [4], a non-trivial concurrent program.

645 This work serves to demonstrate that verified security-preserving compilation for concur-
 646 rent programs is now within reach, by augmenting traditional proof obligations for verified
 647 compilation (e.g. square-shaped semantics preservation) with those specific to security (e.g.
 648 absence of termination- and timing-leaks) as depicted in Figure 2. We hope that this work
 649 paves the way for future large-scale verified security-preserving compilation efforts.

References

- 650 —
- 651 **1** G. Barthe, B. Grégoire, and V. Laporte. Secure compilation of side-channel countermeasures:
652 The case of cryptographic “constant-time”. In *2018 IEEE 31st Computer Security Foundations*
653 *Symposium (CSF)*, pages 328–343, July 2018.
- 654 **2** Gilles Barthe, Tamara Rezk, and Amitabh Basu. Security types preserving compilation.
655 *Comput. Lang. Syst. Struct.*, 33(2):35–59, July 2007. URL: [http://dx.doi.org/10.1016/j.](http://dx.doi.org/10.1016/j.csl.2005.05.002)
656 [csl.2005.05.002](http://dx.doi.org/10.1016/j.csl.2005.05.002).
- 657 **3** Gilles Barthe, Tamara Rezk, Alejandro Russo, and Andrei Sabelfeld. Security of multithreaded
658 programs by compilation. *ACM Trans. Inf. Syst. Secur.*, 13(3):21:1–21:32, July 2010. URL:
659 <http://doi.acm.org/10.1145/1805974.1805977>.
- 660 **4** Mark Beaumont, Jim McCarthy, and Toby Murray. The cross domain desktop compositor:
661 Using hardware-based video compositing for a multi-level secure user interface. In *Annual*
662 *Computer Security Applications Conference (ACSAC)*, pages 533–545, 2016.
- 663 **5** Michael R. Clarkson and Fred B. Schneider. Hyperproperties. *J. Comput. Secur.*, 18(6):1157–
664 1210, September 2010. URL: <http://dl.acm.org/citation.cfm?id=1891823.1891830>.
- 665 **6** Joseph Goguen and José Meseguer. Security policies and security models. In *Proceedings of*
666 *the IEEE Symposium on Security and Privacy*, pages 11–20, Oakland, California, USA, April
667 1982. IEEE Computer Society.
- 668 **7** Cliff B. Jones. *Development Methods for Computer Programs including a Notion of Interference*.
669 D.Phil. thesis, University of Oxford, June 1981.
- 670 **8** Thierry Kaufmann, Hervé Pelletier, Serge Vaudenay, and Karine Villegas. When constant-
671 time source yields variable-time binary: Exploiting curve25519-donna built with msvc 2015.
672 In *Cryptology and Network Security*, pages 573–582, Cham, 2016. Springer International
673 Publishing.
- 674 **9** Ramana Kumar, Magnus Myreen, Michael Norrish, and Scott Owens. CakeML: A verified
675 implementation of ML. In *ACM SIGPLAN-SIGACT Symposium on Principles of Programming*
676 *Languages*, pages 179–191, San Diego, January 2014. ACM Press.
- 677 **10** Xavier Leroy. A formally verified compiler back-end. *J. Autom. Reason.*, 43(4):363–446,
678 December 2009. URL: <http://dx.doi.org/10.1007/s10817-009-9155-4>.
- 679 **11** Andreas Lochbihler. Mechanising a type-safe model of multithreaded java with a verified
680 compiler. *Journal of Automated Reasoning*, 61(1):243–332, Jun 2018. URL: [https://doi.](https://doi.org/10.1007/s10817-018-9452-x)
681 [org/10.1007/s10817-018-9452-x](https://doi.org/10.1007/s10817-018-9452-x).
- 682 **12** Luísa Lourenço and Luís Caires. Dependent information flow types. In *ACM SIGPLAN-*
683 *SIGACT Symposium on Principles of Programming Languages*, pages 317–328, Mumbai, India,
684 January 2015. ACM.
- 685 **13** Heiko Mantel, David Sands, and Henning Sudbrock. Assumptions and guarantees for composi-
686 tional noninterference. In *IEEE Computer Security Foundations Symposium*, pages 218–232,
687 Cernay-la-Ville, France, June 2011. IEEE.
- 688 **14** Toby Murray, Robert Sison, and Kai Engelhardt. COVERN: A logic for compositional
689 verification of information flow control. In *European Symposium on Security and Privacy*,
690 London, United Kingdom, April 2018. IEEE.
- 691 **15** Toby Murray, Robert Sison, Edward Pierzchalski, and Christine Rizkallah. Compositional
692 security-preserving refinement for concurrent imperative programs. *Archive of Formal Proofs*,
693 June 2016. http://isa-afp.org/entries/Dependent_SIFUM_Refinement.shtml, Formal
694 proof development.
- 695 **16** Toby Murray, Robert Sison, Edward Pierzchalski, and Christine Rizkallah. Compositional
696 verification and refinement of concurrent value-dependent noninterference. In *IEEE Computer*
697 *Security Foundations Symposium*, pages 417–431, Lisbon, Portugal, June 2016.
- 698 **17** Marco Patrignani and Deepak Garg. Secure Compilation and Hyperproperty Preservation.
699 In *IEEE 30th Computer Security Foundations Symposium, CSF 2017, Santa Barbara, USA,*
700 *August 21 - 25, 2017, CSF’17, 2017.*

- 701 **18** Anton Podkopaev, Ori Lahav, and Viktor Vafeiadis. Bridging the gap between programming
702 languages and hardware weak memory models. *Proc. ACM Program. Lang.*, 3(POPL):69:1–
703 69:31, January 2019. URL: <http://doi.acm.org/10.1145/3290382>.
- 704 **19** Andrei Sabelfeld and David Sands. Probabilistic noninterference for multi-threaded programs.
705 In *Proceedings of the 13th IEEE Workshop on Computer Security Foundations*, CSFW '00,
706 pages 200–, Washington, DC, USA, 2000. IEEE Computer Society. URL: <http://dl.acm.org/citation.cfm?id=794200.795151>.
- 708 **20** Mark Staples, Ross Jeffery, June Andronick, Toby Murray, Gerwin Klein, and Rafal Kolanski.
709 Productivity for proof engineering. In *Empirical Software Engineering and Measurement*,
710 page 15, Turin, Italy, September 2014.
- 711 **21** F. Del Tesesco, D. Sands, and A. Russo. Fault-resilient non-interference. In *2016 IEEE 29th*
712 *Computer Security Foundations Symposium (CSF)*, pages 401–416, June 2016.

713 **A** Informal descriptions of the cases of refinement relation \mathcal{R}_{wr}

714 **A.1** Base cases

- 715 ■ **stop**: This case relates a terminated **While** program with a terminated RISC program
716 (i.e. one where the program counter is at the length of the program text).
- 717 ■ **skip_nop**: This case relates the **While** program **skip** with the configuration where the
718 program counter is at the start of the RISC program [**Nop**].
- 719 ■ **assign_expr**: This case relates the expression evaluation part (for the expression e) of
720 the **While** program $v := e$ with the corresponding part of the RISC program obtained by
721 compiling it with the wr -compiler.
- 722 ■ **assign_store**: As for **assign_expr**, but for the very last **Store** instruction that commits
723 the result of the expression evaluation back to shared memory variable v .
724 It asserts additionally that v must be stable if lock-governed, and non-lock-governed
725 otherwise. This prevents threads from violating the locking discipline (see Section 5.3.2).
- 726 ■ **lock_acq**: This case relates **lock**(k) with **LockAcq** k .
- 727 ■ **lock_rel**: This case relates **unlock**(k) with **LockRel** k .

728 **A.2** Inductive cases

- 729 ■ **seq**: This case relates the **While** program $c_1 ; c_2$ with the *concatenation* $P_1 @ P_2$ of the
730 RISC programs P_1 and P_2 that are respectively the outputs of successful consecutive
731 compilation of c_1 and c_2 by the wr -compiler. It is intended for cases where the **While**
732 (resp. RISC) program is currently in c_1 (resp. P_1).
733 It is an inductive case of \mathcal{R}_{wr} , in that:
734 ■ c_1 is required to be related by \mathcal{R}_{wr} to the present location in P_1 .
735 ■ For all local configurations that obey the **compiled-cmd-config-consistent** requirements,
736 c_2 is required to be related by \mathcal{R}_{wr} to the first instruction of P_2 . This quantification
737 ensures that \mathcal{R}_{wr} remains closed when execution progresses from the first program to
738 the second program.
739 It asserts that P_1 and P_2 are joinable (Section B.2), particularly relevant here to ensure
740 that P_1 can only jump to locations within or at the end of itself (i.e. the start of P_2).
- 741 ■ **join**: This case relates a **While** program c with an offset $pc > \text{length } P_1$ into a RISC
742 program $P_1 @ P_2$, assuming the inductive hypothesis that c is related by \mathcal{R}_{wr} with the
743 offset $pc - \text{length } P_1$ into the RISC program P_2 alone.
744 It is intended primarily for cases where the **While** (resp. RISC) program is currently in
745 the c_2 (resp. P_2) of some consecutively compiled $c_1 ; c_2$ (resp. P_1 concatenated with P_2)

746 but applies more broadly to allow any prepend of dead, unreachable instructions onto
747 the front of a RISC program without breaking \mathcal{R}_{wr} .

748 It also asserts that P_1 and P_2 are joinable, which is important here to ensure that P_2
749 cannot jump back into P_1 .

750 ■ **if_expr**: This case relates the expression evaluation part (for the expression e) of the
751 **While** program **if** e **then** c_1 **else** c_2 with the corresponding part (including the conditional
752 jump **Jz** at the end of expression evaluation) of the RISC program obtained by compiling
753 it with the *wr*-compiler.

754 It relies on both c_1 and c_2 being related by \mathcal{R}_{wr} to its compiled RISC counterparts when
755 started with initialisation states judged valid by *compiled-cmd-config-consistent*.

756 ■ **if_c1**: This case relates some **While** program c'_1 reachable from c_1 with the corresponding
757 part within the c_1 part of the RISC program obtained by compiling **if** e **then** c_1 **else** c_2
758 with the *wr*-compiler.

759 It relies on c_1 being related by \mathcal{R}_{wr} to its compiled RISC counterpart at the appropriate
760 program counter offset.

761 ■ **if_c2**: As for **if_c1**, but for c_2 .

762 ■ **epilogue_step**: This case relates a terminated **While** program to the silent control flow
763 steps navigating to the end of a RISC program from the end of the “then” and “else”
764 branches of a compiled if-conditional.

765 It works only for epilogue step forms (see Section 5.5.2).

766 It is inductive in that it asserts closedness of \mathcal{R}_{wr} over pairwise reachability from the pair
767 currently under consideration – the only case to do so directly.

768 ■ **while_expr**: This case relates the **While** program (**while** e **do** c)’s initial intermediate
769 step to **if** e **then** (c ; **while** e **do** c) **else stop**, and its expression evaluation part, with
770 the expression evaluation and conditional jump of the RISC program that **while** e **do** c
771 was compiled to by *compile-cmd*.

772 It relies on c being related by \mathcal{R}_{wr} to its compiled RISC counterpart when started with
773 initialisation states judged valid by *compiled-cmd-config-consistent*.

774 ■ **while_inner**: This case relates some program c_I ; **while** e **do** c reachable from c ; **while** e **do** c
775 to the loop body part of the RISC program compiled from **while** e **do** c .

776 It relies on c_I being related by \mathcal{R}_{wr} to its compiled RISC counterpart at the appropriate
777 program counter offset.

778 It also carries around the same reliance on c being related by \mathcal{R}_{wr} to its compiled RISC
779 counterpart for all initialisation states judged valid by *compiled-cmd-config-consistent*.

780 ■ **while_loop**: This case handles epilogue steps for the inner loop body program, and the
781 final jump back to the beginning of the **While**-loop.

782 It requires \mathcal{R}_{wr} to relate the terminated **While** program to the end of the compiled
783 loop body, and furthermore also carries around the same reliance on c being related
784 by \mathcal{R}_{wr} to its compiled RISC counterpart for all initialisation states judged valid by
785 *compiled-cmd-config-consistent*.

786 **B More details of some definitions and proofs**

787 **B.1 Proofs about the decomposition principle**

► **Theorem 18** (Soundness of secure-refinement-decomp).

788 secure-refinement-decomp $\mathcal{B} \mathcal{R} \mathcal{I} \text{ abs-steps} \implies \text{secure-refinement } \mathcal{B} \mathcal{R} \mathcal{I}$
789 19

790 **Proof.** The only obligation for `secure-refinement` (Definition 6) not obtained immediately
 791 from `secure-refinement-decomp` (Definition 7) is the cube-shaped `coupling-inv-pres` (Figure 1).

792 The front face of the cube is just ordinary square-shaped refinement preservation (depicted
 793 in Figure 2a), given to us by `secure-refinement-decomp`. This gives us that a single concrete
 794 step from lc_{1C} is simulated by n abstract steps lc_{1A} , where n is given by *abs-steps*.

795 We are then obliged to prove a simulation in the other direction (the back face of the
 796 cube), that n abstract steps from all configurations lc_{2A} related by \mathcal{B} to lc_{1A} are simulated
 797 by some concrete step from lc_{2C} related by \mathcal{R} to lc_{2A} and by \mathcal{I} to lc_{1C} .

798 Here, we lean on the determinism of the abstract program’s evaluation semantics (required
 799 by the theory) to flip the direction of simulation, knowing that n abstract steps from lc_{2A} ,
 800 simulating a single concrete step from lc_{2C} , could only be the very same n abstract steps from
 801 lc_{2A} that we were required to consider. This allows us to use once again the square-shaped
 802 refinement preservation (Figure 2a) given to us by `secure-refinement-decomp`.

803 Consistency of refinement pacing and stopping behaviour (depicted in Figure 2b) given
 804 by `decomp-refinement-safe` (Definition 8) then respectively ensure that n (via *abs-steps*) is
 805 the correct number of abstract steps to consider, and that there will indeed be a concrete
 806 step from lc_{2C} to drive the matching simulation step.

807 Finally, the remainder of `decomp-refinement-safe` (depicted in Figure 2c) discharges the
 808 requirement of closedness and modes-equality maintenance of \mathcal{I} under lockstep execution,
 809 demanded by the bottom face of the cube. ◀

810 B.2 Label allocation and sequential composability

811 For allocating natural numbers to use as labels for RISC instructions the `wr-compiler` ensures
 812 freshness merely by using the highest number reached so far on a “next label” counter (nl
 813 in the invocation example (1)), incrementing the counter before passing it along to subsequent
 814 calls, and outputting the next available unused label on return (as nl' in the example).

815 We define two RISC programs P_1, P_2 to be joinable if they are both:

- 816 ■ **joinable-forward:** P_1 only ever jumps to labels that are either
 - 817 ■ labelling an instruction in P_1 itself, or
 - 818 ■ the label of the very first instruction in P_2 .
- 819 ■ **joinable-backward:** P_2 does not jump to any of the labels of instructions in P_1 .

820 We prove a lemma that says that two RISC programs that were compiled by the `wr-compiler`
 821 *consecutively*—in the sense that the relevant outputs from the first call are fed directly into
 822 the second call—are joinable.

823 B.3 More detail on compile-cmd-input-reqs and the wr-compiler proofs

824 The first two requirements to the predicate `compile-cmd-input-reqs` $C\ l\ nl\ c$ were given in
 825 Section 5.4. Its other two requirements reflect that the terminated `while` program `stop` has
 826 no valid compilation, and that the initial label (if provided) must be valid (see Section B.2
 827 for more information on label allocation).

► **Definition 19** (Requirements on inputs to `compile-cmd`).

$$\begin{aligned}
 & \text{compile-cmd-input-reqs } C\ l\ nl\ c \equiv c \neq \text{stop} \wedge \\
 & (\forall x. l = \text{Some } x \longrightarrow x < nl) \wedge \text{no-unstable-exprs } c\ C \wedge \text{regrec-stable } C
 \end{aligned}$$

831 These input conditions give us enough information to prove that every instruction of a
 832 *CompRec*-annotated RISC program output by a successful run of `compile-cmd` is annotated
 833 by a stable register record, and that the output *CompRec*'s register record is also stable:

► **Lemma 20** (Successful compilations output only stable register records).

$$834 \quad (PCs, l', nl', C', failed) = \text{compile-cmd } C \ l \ nl \ c \ \wedge \ \text{compile-cmd-input-reqs } C \ l \ nl \ c \ \wedge \\ 835 \quad failed = \text{False} \implies (\forall pc < \text{length } PCs. \text{regrec-stable } (\text{map snd } PCs \ ! \ pc)) \ \wedge \ \text{regrec-stable } C'$$

837 **Proof.** By induction on the structure of the `While` language program c , making reference to
 838 the implementation of `compile-cmd`. For cases that must compile expressions, we furthermore
 839 prove and make use of a lemma by induction on the structure of expressions, making reference
 840 to the implementation of the expression compiler function `compile-expr` called by `compile-cmd`.
 841 In essence, we prove that expressions that appear in register records must be stable because
 842 they are always only ever parts of expressions on variables that were stable in the input
 843 program when read into registers, and furthermore, when compiling an `unlock()` the `wr-`
 844 `compiler` will always flush all register records making reference to any variables that the
 845 `unlock()` makes unstable. ◀

846 Before proceeding, we name the parts of `compiled-cmd-config-consistent` more explicitly:

► **Definition 21** (Configuration consistency requirements for compiled commands).

$$847 \quad \text{compiled-cmd-config-consistent } C \ \text{regs} \ \text{mds} \ \text{mem} \equiv \\ 848 \quad \text{regrec-mem-consistent } (\text{regrec } C) \ \text{regs} \ \text{mem} \ \wedge \ \text{asmrec-mds-consistent } (\text{asmrec } C) \ \text{mds}$$

► **Definition 22** (Consistency between a register record, register bank, and shared memory).

$$850 \quad \text{regrec-mem-consistent } \Phi \ \text{regs} \ \text{mem} \equiv \forall r \ e. \ \Phi \ r = \text{Some } e \implies \text{regs } r = \text{ev}_{\text{exp}} \ \text{mem } e$$

► **Definition 23** (Consistency between an assumption record and a mode state).

$$851 \quad \text{asmrec-mds-consistent } \mathcal{S} \ \text{mds} \equiv \mathcal{S} = (\text{mds } \mathbf{AsmNoW}, \ \text{mds } \mathbf{AsmNoRW})$$

852 ► **Lemma 24** (\mathcal{R}_{wr} preserves modes and memory). `preserves-modes-mem` \mathcal{R}_{wr}

853 **Proof.** By induction on the structure of \mathcal{R}_{wr} . For all cases of $(lc_w, lc_r) \in \mathcal{R}_{\text{wr}}$, $lc_w =_{\text{mds}}^{\text{mem}} lc_r$
 854 is either asserted directly by the guards or obtainable from the inductive hypothesis. ◀

855 ► **Lemma 25** (\mathcal{R}_{wr} is closed under changes by others). `closed-others` \mathcal{R}_{wr}

856 **Proof.** By induction on the structure of \mathcal{R}_{wr} . Changes by others (Definition 4) only modify
 857 writable variables the same way for both configurations, so preservation of $=_{\text{mds}}^{\text{mem}}$ is immedi-
 858 ate. Also, `regrec-mem-consistent` is unaffected because `compile-cmd` only creates `regrec-stable`
 859 records (referring to no writable variables). No other \mathcal{R}_{wr} guards mention shared memory. ◀

► **Lemma 26** (Successfully compiled programs maintain config consistency requirements).

$$860 \quad (PCs, l', nl', C', failed) = \text{compile-cmd } C \ l \ nl \ c \ \wedge \ \text{compile-cmd-input-reqs } C \ l \ nl \ c \ \wedge \\ 861 \quad failed = \text{False} \ \wedge \ pc < \text{length } PCs \ \wedge \ P = \text{map fst } PCs \ \wedge \ Cs = \text{map snd } PCs \ \wedge \\ 862 \quad \text{compiled-cmd-config-consistent } (Cs \ ! \ pc) \ \text{regs} \ \text{mds} \ \text{mem} \ \wedge \\ 863 \quad \langle \langle (pc, P), \text{regs} \rangle, \text{mds}, \text{mem} \rangle_r \rightsquigarrow_r \langle \langle (pc', P), \text{regs}' \rangle, \text{mds}', \text{mem}' \rangle_r \rangle \implies \\ 864 \quad \text{compiled-cmd-config-consistent } (\text{if } pc' < \text{length } P \ \text{then } (Cs \ ! \ pc') \ \text{else } C') \ \text{regs}' \ \text{mds}' \ \text{mem}'$$

866 **Proof.** We in fact prove it separately for `regrec-mem-consistent` and `asmrec-mds-consistent`,
 867 in both cases by induction on the structure of the `While` program c . In each case, we use
 868 the simplifiers for the `compile-cmd` implementation to yield the corresponding RISC program
 869 fragment in question, and then prove the lemma for each of the possible locations of pc in
 870 the compiled program. For both proofs, there is some trickiness in accounting for (and ruling
 871 out) which destination pc' must be considered for each of these cases of pc , particularly for
 872 those `While` programs that compile to RISC programs that may have jumps in them.

873 Control flow trickiness aside, the intuition for `regrec-mem-consistent` is that it tests the
 874 correctness of the compilation of expressions, and so for this we must prove a sub-lemma for
 875 maintenance of `compiled-cmd-config-consistent` by induction on the structure of expressions
 876 e that are encountered in the `While` programs `if e then c_1 else c_2 , while e do c' , $v := e$` .
 877 Additionally, `unlock()` flushes register record entries mentioning variables that are to become
 878 unstable, and `while e do c'` conservatively flushes entries to force evaluation of the loop
 879 condition expression. This is safe trivially because flushing entries can never make a consistent
 880 register record inconsistent. The rest of the cases for c are straightforward because they do
 881 not touch the register record.

882 Then for `asmrec-mds-consistent`, the substantial part of the proof is as a test of the
 883 correctness of the compiler's bookkeeping of assumptions being consistent with the semantics
 884 of `lock()` and `unlock()`. The other cases for c do not touch the mode state. ◀

▶ **Lemma 27** (Correctness of the expression compiler).

885 $(PCs, r, C', failed) = \text{compile-expr } C \text{ A l } e \wedge failed = \text{False} \implies (\text{regrec } C') r = \text{Some } e$

887 **Proof.** By induction on the structure of expressions e , using the simplification rules for the
 888 implementation of `compile-expr`, and also relying on the assumptions of correctness of the
 889 register allocation scheme supplied by the instantiator of the theory (see Section 5.3.1). ◀

▶ **Lemma 28** (\mathcal{R}_{wr} is a refinement paced by `abs-stepswr`).

890 $\forall lc_w, lc_r. (lc_w, lc_r) \in \mathcal{R}_{wr} \longrightarrow (\forall lc'_r. lc_r \rightsquigarrow_r lc'_r \longrightarrow$
 891 $(\exists lc'_w. lc_w \rightsquigarrow_w^{(\text{abs-steps}_{wr} lc_w lc_r)} lc'_w \wedge (lc'_w, lc'_r) \in \mathcal{R}_{wr}))$

893 **Proof.** By induction on the structure of \mathcal{R}_{wr} .

894 The base case `stop` is immediate, because it pertains to a terminated `While` and RISC
 895 program. The base cases that proceed in one step to a terminating program configuration
 896 (`skip_nop`, `assign_store`, `lock_acq`, `lock_rel`) are fairly straightforward because after
 897 dealing with the single step, the resulting obligation can then be handled by the `stop` case.
 898 This leaves the last remaining base case `assign_expr`, which proceeds in one step either to
 899 itself, or to `assign_store`. In all of these cases, we use Lemma 26 to obtain the preservation
 900 of the guards demanded by the \mathcal{R}_{wr} introduction rule for the destination configuration of the
 901 step. Particularly, the `assign_store` case must make use of `regrec-mem-consistent` and the
 902 correctness of `compile-expr` (Lemma 27) in order to ensure that once the expression evaluation
 903 result is written back to shared memory, $lc'_w =_{\text{mds}}^{\text{mem}} lc'_r$ holds as demanded by the `stop` case.

904 The inductive cases that concern expression evaluation (`if_expr`, `while_expr`) are much
 905 like `assign_expr` in that they have the possibility of progressing in one step to themselves.
 906 Unlike `assign_expr` however, their other possibility is a conditional jump based on the result
 907 of that expression. Again we use Lemma 27 to obtain that the result is an accurate calculation
 908 of the expression, and this time we prove by the two different cases whether `if_expr` ends
 909 up in `if_c1` or `if_c2`, or if `while_expr` ends up in `while_inner` or at `stop` (having jumped

910 to the exit label). In these cases, the guards over which the inductive references to \mathcal{R}_{wr} have
 911 been quantified are versatile enough to discharge themselves (when $*_expr$ steps to itself),
 912 or to discharge any reachable initial starting state for the nested compiled RISC program,
 913 given that Lemma 26 ensures the invariance of these guards.

914 This just leaves the inductive cases that pertain to configurations inside a nested com-
 915 piled RISC program (`if_c1`, `if_c2`, `while_inner`), or at the end of one (`epilogue_step`,
 916 `while_loop`). In these cases, the inductive hypotheses obtained from the inductive reference
 917 to \mathcal{R}_{wr} are always enough to satisfy the guards demanded by the possible destination cases.
 918 Like in the proof of Lemma 26, the trickiness mostly comes from accounting for all the possible
 919 cases of control flow (ruling out spurious destinations) that need to be considered. ◀

920 ▶ **Lemma 29.** $\text{strong-low-bisim-mm } \mathcal{B} \wedge \text{no-high-branching } \mathcal{B} \implies$
 921 $\text{decomp-refinement-safe } \mathcal{B} \mathcal{R}_{wr} \mathcal{I}_{wr} \text{abs-steps}_{wr}$

922 **Proof.** Definition 8 gives us the following obligations.

923 For consistent stopping behaviour, we prove a lemma that RISC programs stop if and
 924 only if their pc is outside the program text P , i.e. $pc > \text{length } P$. Because \mathcal{I}_{wr} equates pc
 925 and P for the two configurations, then clearly both have identical stopping behaviour.

926 For consistency of change in timing behaviour, abs-steps_{wr} depends only on `While` and
 927 RISC program locations, and `no-high-branching` and \mathcal{I}_{wr} forces them (resp.) to be equal for
 928 the local configurations under consideration.

929 For closedness of \mathcal{I}_{wr} under lockstep execution, the only non-straightforward cases to
 930 consider are conditional branching, and the locking primitives. For conditional branching, we
 931 use `no-high-branching` for \mathcal{B} with memory preservation via \mathcal{R}_{wr} (Lemma 11) to ensure that
 932 the conditional branching outcome is the same on both sides.

933 Finally, as the only operations that touch mode state, the locking primitives are the only
 934 non-straightforward cases for mode state equality maintenance under lockstep execution. As
 935 all lock memory is classified `Low` (see Section 5.3.2), we use `strong-low-bisim-mm` for \mathcal{B} with
 936 memory preservation via \mathcal{R}_{wr} to ensure the RISC configurations behave consistently. ◀

937 ▶ **Lemma 30.** $\text{strong-low-bisim-mm } \mathcal{B} \wedge \text{no-high-branching } \mathcal{B} \implies$
 938 $\text{secure-refinement-decomp } \mathcal{B} \mathcal{R}_{wr} \mathcal{I}_{wr} \text{abs-steps}_{wr}$

939 **Proof.** Referring to Definition 7, the obligations pertaining only to \mathcal{R}_{wr} and abs-steps_{wr} are
 940 discharged by Lemma 14, Lemma 12, and Lemma 11. Pertaining to \mathcal{I}_{wr} : clearly \mathcal{I}_{wr} is
 941 symmetric, and furthermore it is `cg-consistent` (Definition 3) because the actions over which
 942 \mathcal{I}_{wr} must be closed modify only the shared memory, and \mathcal{I}_{wr} places only restrictions on the
 943 program text and current location. The final obligation is discharged by Lemma 29. ◀

▶ **Theorem 31** (Successful compilations are refinements in \mathcal{R}_{wr}).

$$\begin{aligned}
 & (PCs, l', nl', C', failed) = \text{compile-cmd } C \ l \ nl \ c \wedge \text{compile-cmd-input-regs } C \ l \ nl \ c \wedge \\
 & failed = \text{False} \wedge \text{compiled-cmd-config-consistent } C \ \text{regs} \ \text{mds} \ \text{mem} \wedge P = \text{map fst } PCs \\
 & \implies \langle \langle c, \text{mds}, \text{mem} \rangle_w, \langle \langle (0, P), \text{regs} \rangle, \text{mds}, \text{mem} \rangle_r \rangle \in \mathcal{R}_{wr}
 \end{aligned}$$

948 **Proof.** By induction on the structure of `While`. The compiler input and initial configuration
 949 conditions we impose allow us to have each of `skip`, `cmd ; cmd`, `if exp then cmd else cmd`,
 950 `while exp do cmd`, `v := exp`, `lock(k)`, and `unlock(k)` and their compiled output meet the
 951 guards of the introduction rules for the cases `skip`, `seq`, `if_expr`, `while_expr`, `assign_expr`,
 952 `lock_acq`, and `lock_rel` of \mathcal{R}_{wr} that were designed for them respectively. ◀