Verified Secure Compilation for Mixed-Sensitivity Concurrent Programs

ROBERT SISON * [†], TOBY MURRAY *

 * School of Computing and Information Systems, University of Melbourne, Australia
 † CSIRO's Data61 and UNSW Sydney, Australia (e-mail: [firstname].[lastname]@unimelb.edu.au)

Abstract

Proving only over source code that programs do not leak sensitive data leaves a gap between reasoning and reality that can only be filled by accounting for the behaviour of the compiler. Furthermore, software does not always have the luxury of limiting itself to single-threaded computation with resources statically dedicated to each user to ensure the confidentiality of their data. This results in *mixed-sensitivity concurrent programs*, which might reuse memory shared between their threads to hold data of different sensitivity levels at different times; for such programs, a compiler must preserve the *value-dependent* coordination of such *mixed-sensitivity reuse* despite the impact of *concurrency*.

Here we demonstrate, using Isabelle/HOL, that it is feasible to verify that a compiler preserves *noninterference*, the strictest kind of confidentiality property, for mixed-sensitivity concurrent programs. First, we present notions of refinement that preserve a *concurrent value-dependent* notion of noninterference that we have designed to support such programs. As proving noninterference-preserving refinement can be considerably more complex than the standard refinements typically used to verify semantics-preserving compilation, our notions include a decomposition principle that separates the semantics-preservation from security-preservation concerns. Second, we demonstrate that these refinement notions are applicable to verified secure compilation, by exercising them on a single-pass compiler for mixed-sensitivity concurrent programs that synchronise using mutex locks, from a generic imperative language to a generic RISC-style assembly language. Finally, we execute our compiler on a nontrivial mixed-sensitivity concurrent program modelling a real-world use case, thus preserving its source-level noninterference properties down to an assembly-level model automatically. All results are formalised and proved in the Isabelle/HOL interactive proof assistant.

Our work paves the way for more fully featured compilers to offer verified secure compilation support to developers of multithreaded software that must handle data of multiple sensitivity levels.

1 Introduction

Here we show how to extend secure compilation support to programs that are designed to address two fundamental problems of scale: (1) the need to divide work in computer systems that handle information, and (2) the need to share scarce resources to be able to service every customer for whom that work is done. There will always be a program for which that sharing is not abstracted; that program's responsibility is to implement that sharing in such a way that it never allows the information of one customer to flow to

2 Verified Secure Compilation for Mixed-Sensitivity Concurrent Programs

another. In this paper, we prove formally that a compiler does not break that program's responsibility.

It is well known that program translations of the kind carried out by compilers can easily break security properties like confidentiality (Kaufmann et al., 2016; Barthe et al., 2018). This is especially the case for *mixed-sensitivity concurrent programs*, which feature both:

- Concurrency of access to memory locations shared between different threads of execution. A compiler must preserve both (1) the synchronisation that coordinates threads' access to shared memory, and (2) the absence of any internal timing leaks, to prevent them from manifesting as storage leaks (Volpano & Smith, 1998).
- Mixed-sensitivity reuse of shared memory to hold information of different sensitivity levels at different times. A compiler must preserve the program functionality that coordinates this reuse; this implies support for value-dependent classification policies, which allow the classification of a memory location to change dynamically depending on values held in other memory locations (Murray, 2015). Furthermore, it must do so accounting for the potential impact of concurrent access by other threads.

Although existing verified compilers for dialects of mainstream programming languages, like CompCert (Leroy, 2009) and CakeML (Kumar et al., 2014), have been proved to preserve program functionality (semantics) and some timing-sensitive forms of noninterference (Barthe et al., 2020), none are yet verified to preserve proofs of noninterference for mixed-sensitivity concurrent programs. Ideally such a compiler, applied to the threads of a proved-secure mixed-sensitivity concurrent program, would yield assembly code that, run concurrently, also composes into a secure mixed-sensitivity concurrent program.

To this end, here we present notions of concurrent value-dependent noninterferencepreserving refinement, which are compositional across the threads of mixed-sensitivity concurrent programs. In these notions, the usual square-shaped commuting diagram commonly used to depict (semantics-preserving) refinement (Figure 4a) has been replaced by a cube (Figure 3). This reflects that it preserves a 2-safety hyperproperty (Terauchi & Aiken, 2005; Clarkson & Schneider, 2010), which compares two executions rather than examining a single one. Our earlier work (Murray et al., 2016b) was the first to make this observation 76 and to propose a general cube-shaped refinement property; however other work on verified secure compilation targeted towards noninterference preservation (Barthe et al., 2018, 2020) since made the same observation. As these cube-shaped properties are significantly more complicated to prove than standard notions of semantics-preserving refinement typical in verified compilation (Leroy, 2009; Kumar et al., 2014), we present a principle of decomposing the cube (Figure 3) into three separate obligations (Figure 4): the first of 82 these is akin to semantics-preserving refinement, while the rest prevent the introduction of any termination- and timing-leaks. A simple comparison of proof effort for a refinement example (Figure 2) shows this approach can almost halve its complexity, and that it is applicable to proofs of refinement for programs with secret-dependent control flow—the example pads an **if** h **then** ... **else** ... **fi** conditional with **skip**s, so as not to introduce a timing leak of h.

We then go on to demonstrate that the decomposition principle we provide makes our notion of refinement a tractable target for verified secure compilation. Our compiler is

90 91

47

48

49

50

51

52 53

54

55

56

57

58

59

60

61 62

63

64

65

66

67

68

69

70

71

72

73

74

75

77

78

79

80

81

83

84

85

86

87

88

89

an executable function in Isabelle/HOL that translates mixed-sensitivity concurrent pro-93 grams that synchronise using mutex locks, from a generic imperative While language to 94 a generic RISC-style assembly language. In particular, it supports the class of programs 95 that avoid all *implicit flows*, where a secret determines the choice between two control 96 flow paths with different observable effects, by disallowing any secret-dependent control 97 flow—for example, disallowing if h then ... else ... fi conditionals to prevent any tim-98 ing leaks from h. This is a common approach against implicit flows, as it avoids any precise 99 source-level reasoning about time. To preserve confidentiality for programs that take that 100 approach, we instantiate the decomposition principle so that it enforces that our compiler 101 does not introduce any new secret-dependent control flow. Furthermore, as part of satisfy-102 ing the demands of our refinement notion, our compiler demonstrates a way of formalising 103 and proving when it is safe for a compiler to perform optimisations in the presence of 104 concurrency. To ensure that the contents of shared memory locations are preserved under 105 compilation despite potential interference from other threads, our compiler tracks which 106 shared memory locations are free from data races. It then makes use of this tracking to 107 avoid redundant loads from "stable" (i.e. race-free) shared variables safely, that would 108 otherwise be considered unsafe to omit. 109

Finally, to show that the compiler preserves noninterference for actual mixed-110 sensitivity concurrent programs, we execute it on a real-world use case: a model of the 111 software-componentised input-handling regime for the Cross Domain Desktop Compositor 112 (Beaumont et al., 2016), a device that enforces information-flow control over input clas-113 sified dynamically by a trusted user. We leave treatment on the design and application 114 of per-thread proof techniques establishing CVDNI for the successive versions of this 115 model to other works (Murray et al., 2018; Sison, 2020), and here focus on its CVDNI-116 preserving compilation—expanding on Sison & Murray (2019), the conference version 117 of this paper. This yields the first proofs of noninterference for an assembly-level model 118 of a nontrivial mixed-sensitivity concurrent program, demonstrating the power of verified 119 secure compilation to preserve security properties of compiled code. 120

The structure of our paper is as follows. First, we present language-independent notions of noninterference and its refinement, designed for mixed-sensitivity concurrent programs (Section 2). Our attention then turns to preliminaries for our compiler: the main properties of interest of the source While language it compiles (Section 3), and of the target RISC language it produces (Section 4). Then, after presenting the details of our compiler and its verification (Section 5), and the case study to which we apply it (Section 6), we discuss the most closely related work in the area (Section 7), before concluding (Section 8).

We expand on the conference version of this paper (Sison & Murray, 2019) as follows:

- Here we have adapted the noninterference properties to support assumptions on initial memory and extra security requirements; these will allow us to clarify exactly what our compiler is verified to preserve, and for which kinds of programs.
- We also in Section 2 present further preliminaries that will allow us to explain in greater detail the different ways to establish and use proofs about a verified compiler to obtain whole-system noninterference at the target-language level. These include:
 - 1. The side conditions and theorem of compositionality for the noninterference properties. In Section 3, we then for the first time present the proof,
- 137 138

121

122

123

124

125

126

128

129

130

131 132

133 134

135

Verified Secure Compilation for Mixed-Sensitivity Concurrent Programs

whose details were elided from Sison & Murray (2019), for a noncompositional 139 "global" part of this side condition, which is necessary to obtain whole-system 140 noninterference from per-thread noninterference both at source and target level. 141 2. A whole-system refinement theorem, adapted to support assumptions on ini-142 tial memory. This theorem was alluded to in Murray et al. (2016b) but, until 143 now, has never been formally presented outside of the Isabelle/HOL theories. 144 It gives us a means to prove preservation of whole-system noninterference by 145 the compiler, without having to re-prove the noncompositional side condition 146 at the target-language level. 147 • In Section 5, we then compare alternative methods of obtaining whole-system secu-148 rity at RISC level, that a developer would choose depending on whether all, or only 149 some threads are compiled with our compiler. In contrast, Sison & Murray (2019) 150 stopped after presenting the application of refinement decomposition principle. 151 • In Section 6, the case study to which we apply the compiler is significantly expanded, 152 being a 3-component version of the CDDC input-handling program-closer to a 153 version presented in Murray et al. (2018)-as opposed to the 2-component version 154 of Sison & Murray (2019). 155 • Furthermore, we present substantially more details of our case study in Section 6, 156 which were mostly elided from Sison & Murray (2019). These include formal state-157 ments of both the source-level properties preserved and the target-level properties 158 obtained, alongside explanations of all alternative methods for obtaining the latter 159 from the former. 160 · Finally, every lemma and theorem we prove is presented with a proof sketch or 161 explanation, which were largely absent from Sison & Murray (2019). 162 163 164 165 166 167 2 Noninterference and its refinement for mixed-sensitivity concurrent programs 168 To support mixed-sensitivity concurrent programs, we verify our compiler to preserve the 169 concurrent value-dependent noninterference (CVDNI) notions of Murray et al. (2016b). 170 In this section we present the definitions of CVDNI and its refinement, as we have adapted 171

In this section we present the definitions of CVDNI and its refinement, as we have adapted them from that work's Isabelle formalisation (Murray *et al.*, 2016*c*,a). In particular, the version of the theory we present here supports extra customisation of requirements beyond the prior work; we will need this to parameterise the theory with initial conditions needed for a compositionality property of our source language (Section 3), and our compiler's preservation of a ban on secret-dependent control flow (Section 5). Furthermore, it is simplified to the case where the shared memory is the same for both the original *abstract* and the refined *concrete* program—refinement adds no new shared variables. Later, we will instantiate this CVDNI theory to have our compiler's source and target languages (Sections 3, 4) respectively play the roles of the abstract and concrete programs' languages in the theory.

We begin by introducing with an illustrative example the challenges of verifying valuedependent noninterference in the presence of shared-variable concurrency (Section 2.1). Then we present the per-thread and whole-system noninterference properties themselves

183 184

172

173

174

175

176

177

178

179

180

181

182

	while TRUE do	
185	<pre>lock(workspace_lock);</pre>	
186	while !suspended do	
187	<pre>lock(source_lock);</pre>	
188	<i>workspace</i> := <i>source</i> ;	
189	/* operations on <i>workspace</i> */	
190	if <i>domain</i> = LOW then	
191	$low_sink := workspace$	
192	else	
193	<i>high_sink</i> := workspace;	
194	<i>workspace</i> := 0	
195	fi;	
196	unlock (<i>source_lock</i>)	
197	od;	
198	<pre>unlock(workspace_lock);</pre>	
199	while suspended do skip od	
200	od	



(b) The phone providing the High personality: $domain \neq LOW$, and *source* is classified High to reflect that the user might type in secrets.



(c) The phone displaying visual indicators that it is providing the Low personality: *domain* = LOW, and *source* is classified Low to reflect that we trust the user not to type in secrets.

(a) Input processing worker thread program

Fig. 1: Example: Touchscreen input processing for a dual-personality smartphone. Reproduced from Sison & Murray (2019).

(Section 2.2), followed by the notion of *per-thread refinement* that preserves the perthread property between the two languages (Section 2.3). As the cube-shapedness of noninterference-preserving refinement diagrams in general makes them difficult to apply directly to compiler verification, we present a decomposition principle (Section 2.4) that we will use to prove CVDNI-preserving refinement for our compiler. We then present requirements and a theorem for *whole-system refinement* by which we have that CVDNIpreserving refinement is compositional across the threads of the program being compiled, such that it yields the whole-system property at the target language level (Section 2.5).

2.1 Illustrative example of a mixed-sensitivity concurrent program

Consider the task of verifying a multithreaded system that manages the user interface (UI) for a *dual-personality smartphone*, a phone that provides clearly distinguished user contexts (*personalities*), typically for work versus leisure. Specifically, our task is to verify that it does not leak *sensitive* information intended only for one of those personalities, which we classify High (Figure 1b), to locations belonging to the other, which we classify Low (Figure 1c).

Here and generally throughout this paper, our *attacker model* is an entity that can read from the system's *untrusted sinks*: some subset of permanently Low-classified locations not subject to synchronisation. In our example, the untrusted sinks may include WLAN device registers in a hostile environment.

The smartphone's UI system consists of a number of threads running concurrently with a shared address space; we aim to verify that, as a whole, this system of threads satisfies the security requirement. However, to avoid a state space explosion that is exponential in the

229 230

201

202

203

204 205

206

207

208

209

210

211

212

213 214 215

216

217

218

219

220

221

222

223

224

225

226

227

number of threads, we must do this *compositionally*: one thread at a time, then combining the results of these analyses.

232 233

6

231

234

235

236

237

238

239

240

241

242

243

244

245

246

247

248

249

250

251

252

253

254

255

256

257

258

259 260

261

262

263

264

We focus on a particular worker thread (Figure 1a), the one responsible for sending touchscreen input from the *source* variable to its intended destination.

The first challenge is that the destination depends on which personality the phone is currently providing, which is indicated by the value of *domain*. This is reflected by the classification of *source* being dependent on the value of *domain: source* is classified Low exactly when *domain* = LOW (where LOW is a designated constant), and is classified High otherwise. Due to this dependency, *domain* is known as a *control variable* of *source*.

The second challenge is the worker thread runs in a shared address space that might be accessed or modified by other threads, for various purposes. One of these threads may be responsible for maintaining that *domain* = LOW exactly when the phone indicates it is providing the Low personality (Figure 1c), so the user knows not to type in anything sensitive. Another thread may be responsible for assigning *suspended* := TRUE when the user turns the phone's screen off, to make the worker stop processing touchscreen input. We may then wish for *workspace* to be usable by some other thread—for example, processing input from a fingerprint scanner—in such a way that it can assume that workspace no longer contains any sensitive values.

When we analyse one thread like this worker in terms of our compositional security property (Section 2.2), all the other threads in the system are trusted to do two things:

- 1. They follow a synchronisation discipline; in particular for this example, this is a mutual exclusion (mutex) locking discipline: If read- or write-access to a certain variable is governed by a lock, each thread may only access the variable in that manner if they hold that lock. Mutual exclusion then follows from the semantics of the locking primitives ensuring only one thread may hold a given lock at a time.
- 2. They themselves do not leak values from High-classified locations (we refer to such values themselves as High) to Low-classified locations that are read-accessible to other threads. Note that, here, it is our objective to prove that the thread we are analysing can be trusted in the same way.

Even under these assumptions, the concurrency gives rise to some tricky considerations.

First, it is important that no thread in the system (including the one under analysis) modifies any control variables carelessly. For example, writing *domain* := LOW immediately after the worker reads a High value from *source*, will cause it to leak to *low_sink*. To prevent this, the worker uses *source_lock*, granting it *exclusive write-access* to *source* and *domain*.

265 Furthermore, as noted above, we may want to ensure that a non-attacker-observable 266 location is nevertheless cleared of any sensitive values before being used by another thread. 267 In our example, we classify *workspace* Low for the analysis to enforce this when the worker 268 is suspended, but as the worker sometimes uses it to process High values, it is important 269 to know *workspace* is accessible only to the worker during that time. To ensure this, the 270 worker uses workspace_lock, granting it exclusive read- and write-access to workspace. It 271 is then responsible for clearing it of any High values by the time it releases that access.

- 272
- 273
- 274
- 275
- 276

2.2 Concurrent value-dependent notions of noninterference

Having illustrated the challenges with an example, we now present the definitions of per-278 thread and whole-system noninterference, the theorem by which the former composes into 279 the latter, and the compositionality side conditions demanded by that theorem. 280

As proved for each thread, CVDNI is defined by Murray et al. (2016b) in terms of: 281

- 1. A binary strong low-bisimulation (modulo modes) relation \mathcal{B} between program configurations, which serves as witness to CVDNI. In the style of other lowbisimulation-based noninterference definitions (Focardi et al., 1995; Sabelfeld & Sands, 2000; Mantel et al., 2011) it requires the program configurations it relates to agree on their "low"-observable portions, and demands that lock-step execution preserves that correspondence. Furthermore, it is rely-guarantee-style concurrency aware, following Mantel et al. (2011), but modified to allow value-dependent classifications (Murray, 2015) for mixed-sensitivity reuse (see next point). 290
 - 2. A *classification* function \mathscr{L} that determines the "low"-observable portion of a program configuration, thus affecting \mathscr{B} 's requirements. The innovation of \mathscr{L} , as parameterised first by Murray (2015) and then by Murray *et al.* (2016b) as reproduced here, is that \mathcal{L} can depend on values in the program configuration itself, thus expressing dynamic and not just static classifications.

The theory is parameterised over the type of values Val, a finite set of shared variables Var, and a deterministic evaluation step semantics \rightarrow between local configurations of a thread in a concurrent program. Each local configuration is a triple (*tps, mds, mem*):

- *tps* :: *ThreadPrivate* is the *thread-private state*, which the theory will consider to be permanently inaccessible to the attacker and not shared with the other threads. Note that, due to this inaccessibility, we allow the user of the theory to parameterise the type ThreadPrivate, and we do not impose any particular structure on it.
- $mds:: Mode \Rightarrow Var set$ is the (assume-guarantee) mode state, which is ghost state associating each of $Mode \triangleq \{AsmNoW, AsmNoRW, GuarNoW, GuarNoRW\}$ with a set of shared variables. Intuitively, it identifies the set of variables for which the thread currently Assumes it possesses (or Guarantees it respects) exclusive permission to Write (or Read and Write), granted (or obligated) for those variables typically by some synchronisation scheme. This facilitates compositional, rely-guarantee-style reasoning about such access (Jones, 1981; Mantel et al., 2011). For example, when our worker thread (of Figure 1a) holds source_lock, it assumes that no other threads write to source or its control variable domain (i.e. {source, domain} \subseteq mds AsmNoW), otherwise it guarantees it does not write to them (**GuarNoW**). Similarly, when it holds *workspace_lock* it assumes that no other threads *read or write* to *workspace* (i.e. *workspace* \in *mds* **AsmNoRW**), and at all other times it makes the corresponding guarantee (GuarNoRW).
- mem :: Mem is shared memory considered potentially accessible to the attacker and 317 other threads. To make what is accessible amenable to analysis, we impose the 318 structure $Mem \triangleq Var \Rightarrow Val$, a total map from shared variable names to values. 319
- 320
- 321

277

282

283

284

285

286

287

288

289

201

292

293

294

295

296

297

298 299

300

301

302

303

304

305

306

307

308

309

310

311

312

313

314

315

316

Verified Secure Compilation for Mixed-Sensitivity Concurrent Programs

The theory is then further parameterised by the value-dependent classification function

\mathscr{L} :: $Mem \Rightarrow Var \Rightarrow \{\text{High}, \text{Low}\}, \text{ inducing a function } \mathscr{C} \text{vars} :: Var \Rightarrow Var set that returns all the control variables of a given variable. In our worker thread example, \mathscr{L} mem x gives:$
 High when x is <i>high_sink</i>, meaning <i>high_sink</i> is classified High at all times. when x is <i>source</i>: Low if <i>mem domain</i> = LOW, and High otherwise. Low for all other variables x, meaning they are classified Low at all times.
The set $\mathscr{C} = \{y \mid \exists x. y \in \mathscr{C} \text{vars } x\}$ is then defined to contain all control variables in the system. Thus in our worker thread example, $\mathscr{C} \text{vars } source = \{domain\}$ and $\mathscr{C} = \{domain\}$. With these parameters having been set, we can now define notions of <i>observational equivalence</i> —underpinning noninterference properties—that are value dependent. The notion of observational equivalence of memories, used by the <i>whole-system</i> non-interference property to quantify over initial state pairs, is as follows: Variables that are value-dependently classified Low <i>according to both memories</i> are required to have the same value <i>in both memories</i> . Formally, as defined originally by Murray (2015):
Definition 2.1 (Low-equivalent memories).
$mem_1 = {}^{Low} mem_2 \triangleq \forall x. \mathscr{L} mem_1 \ x = Low \longrightarrow mem_1 \ x = mem_2 \ x$
Note that the asymmetry of Definition 2.1 (also Definition 2.3 to follow) referring only to mem_1 is resolved by requiring the classification function \mathscr{L} to classify all control variables as Low <i>statically</i> —that is, Low <i>always</i> , regardless of the memory state (cf. our restriction Proposition 3.4 on the classification of state used to implement locks, later in Section 3.2). To support compositionality for concurrent programs, however, the equivalence notion for the <i>per-thread</i> noninterference property is relaxed to be <i>modulo modes</i> in the style of Mantel <i>et al.</i> (2011): Here, Low-classified non-control variables $x \notin \mathscr{C}$ are only required to have the same value if they are assumed to be <i>readable</i> by other threads according to
the mode state. (Control variables $x \in \mathcal{C}$ are excluded from that relaxation, and are <i>always</i>

required to be equal.) Defined more formally, again as originally by Murray (2015):¹

Definition 2.2 (Readability of variable *x*, according to mode state *mds*).

readable $mds x \triangleq x \notin mds$ AsmNoRW

Definition 2.3 (Low-equivalence of memories, modulo the mode state *mds*).

```
\begin{array}{l} \textit{mem}_1 =_{\textit{mds}}^{\mathsf{Low}}\textit{mem}_2 \ \triangleq \\ \forall x. \ x \in \mathscr{C} \ \lor \ \mathscr{L} \textit{ mem}_1 \ x = \mathsf{Low} \ \land \ \mathsf{readable} \ \textit{mds} \ x \ \longrightarrow \ \textit{mem}_1 \ x = \textit{mem}_2 \ x \end{array}
```

Moreover, we will use notation $lc_1 = _{mds}^{Low} lc_2$ from Sison & Murray (2019) to lift Definition 2.3 to local program configurations, asserting also that the local configurations lc_1 and lc_2 have the same assume–guarantee mode state. Additionally, we will use notation $lc_1 = _{mds} lc_2$ to denote (only) that lc_1 and lc_2 have the same assume–guarantee mode state.

¹ Logical operator precedence here is just as in Isabelle/HOL—from most tightly to least: $\land, \lor, \longrightarrow$.

367 368 8

323

324

330

331

332 333

334

335

343

344

345

346

347

348

349

350

351 352 353

354

355 356

357 358

359

360 361

362

363

364

Thus, intuitively, the user of the theory should model the permanent untrusted output 369 sinks, of their whole concurrent program, as variables for which \mathcal{L} always returns Low, 370 ungoverned by any synchronisation scheme that the attacker cannot be trusted to follow. In our worker example program (of Figure 1a), *low_sink* is untrusted permanently in this 372 way, but workspace is untrusted only when unlocked. 373

371

374

375

376

377

378

379

380

381

382

383

384

385

386

387

388

389

390

391

392 393

394 395

396

397

407

We now have almost enough definitions to state the per-thread compositional security property. This property will assert the existence of a witness *bisimulation* relation *B* for every possible observationally equivalent pair of starting configurations. Specifically, this witness relation must be a strong low-bisimulation (modulo modes) (denoted by strong-low-bisim-mm \mathcal{B}), meaning that it must satisfy the following three conditions:

- 1. It must maintain observational indistinguishability by requiring that all configuration pairs it relates (i.e. $(lc_1, lc_2) \in \mathscr{B}$) that have the same mode state $(lc_1 =_{\mathsf{mds}} lc_2)$, are low-equivalent modulo modes $(lc_1 = \frac{Low}{mds} lc_2)$.
- 2. Furthermore, it must be a *bisimulation* by being symmetric (denoted by sym \mathscr{B}) and progressing to itself: Any step taken by one of the configurations $(lc_1 \rightsquigarrow lc'_1)$ must be matched by some step taken by the configuration related to it $(lc_2 \rightsquigarrow lc'_2)$, so the destinations remain related (i.e. $(lc'_1, lc'_2) \in \mathscr{B}$) and modes-equal $(lc'_1 =_{mds} lc'_2)$.
- 3. Finally, it must be *closed under globally consistent changes* made to memory by other threads (denoted by cg-consistent *B*)-that is, changes that preserve lowequivalence and are permitted by the current mode state *mds*. Specifically, other threads are permitted to change either of variable x's value or its classification only when x is considered *writable* by the current mode state (denoted by writable *mds x*, Definition 2.5). This is the most crucial element of the per-thread CVDNI property itself that ensures its compositionality for concurrent programs.

These requirements are formalised by Definition 2.4, using Definitions 2.5 and 2.6:

Definition 2.4 (Strong low bisimulation, modulo modes).

. . .

strong-low-bisim-mm
$$\mathscr{B} \triangleq$$
 cg-consistent $\mathscr{B} \land$ sym $\mathscr{B} \land$
 $(\forall lc_1 \ lc_2. \ (lc_1, lc_2) \in \mathscr{B} \land \ lc_1 =_{\mathsf{mds}} lc_2 \longrightarrow$
 $lc_1 =_{\mathsf{mds}}^{\mathsf{Low}} lc_2 \land$
 $(\forall lc'_1. \ lc_1 \rightsquigarrow lc'_1 \longrightarrow (\exists lc'_2. \ lc_2 \rightsquigarrow lc'_2 \land \ lc'_1 =_{\mathsf{mds}} lc'_2 \land (lc'_1, lc'_2) \in \mathscr{B})))$

Definition 2.5 (Writability of variable *x*, according to mode state *mds*).

writable
$$mds x \triangleq x \notin mds$$
 AsmNoW $\land x \notin mds$ AsmNoRW

Definition 2.6 (Closedness under globally consistent changes).

cg-consistent $\mathscr{B} \triangleq \forall tps_1 mem_1 tps_2 mem_2 mds$. 408 $(\langle tps_1, mds, mem_1 \rangle, \langle tps_2, mds, mem_2 \rangle) \in \mathscr{B} \longrightarrow$ 409 $(\forall mem'_1 mem'_2)$. $(\forall x. (mem_1 x \neq mem'_1 x \lor mem_2 x \neq mem'_2 x \lor$ 410 $\mathscr{L} mem_1 x \neq \mathscr{L} mem'_1 x) \longrightarrow \text{writable } mds x) \land mem'_1 = \overset{\text{Low}}{\underset{mds}{\text{mem}'_2}} mem'_2 \longrightarrow$ 411 412 $(\langle tps_1, mds, mem'_1 \rangle, \langle tps_2, mds, mem'_2 \rangle) \in \mathscr{B})$ 413 414

10 Verified Secure Compilation for Mixed-Sensitivity Concurrent Programs

Note that, to prevent unnecessary proof effort, strong-low-bisim-mm assumes instead of 415 asserting the initial modes-equality $(lc_1 = mds lc_2)$, as the security property that will use 416 strong-low-bisim-mm will take responsibility for asserting it (to follow, in Definition 2.7). 417 We now present definitions of the CVDNI security properties that differ from those 418 published in Murray et al. (2016b) and our conference paper Sison & Murray (2019), in 419 that they allow two additional forms of customisation as parameters to the theory, necessary 420 for a fuller written presentation of the formal verification of our compiler: 421 1. Initialisation requirements for the system, in the form of a predicate over shared 422 memory called INIT. 423 The per-thread and whole-system security properties are *relaxed* such that they only 424 quantify over initial shared memories that obey this predicate. 425 2. Extra requirements to be imposed on top of strong low-bisimulation modulo modes, 426 in the form of a predicate over bisimulation relations called EXTRA. 427 The per-thread security property is strengthened to impose these additional require-428 ments on any candidate security witness. 429 430 When dropped from each of the names of the properties "com-secure" and "sys-secure" 431 soon to be introduced, *INIT* and *EXTRA* default to (λ_{-} . True); in that case, the definitions 432 of those properties will then simplify to their original versions as presented in Murray et al. 433 (2016b); Sison & Murray (2019). 434 The per-thread security property is then as follows: 435 436 437 Definition 2.7 (Per-thread compositional security, with INIT, EXTRA requirements). 438 $\operatorname{com-secure}_{INIT}^{EXTRA}(tps, mds) \triangleq \forall mem_1 \ mem_2.$ 439 $mem_1 = \frac{Low}{mds} mem_2 \land INIT mem_1 \land INIT mem_2 \longrightarrow$ 440 441 $(\exists \mathscr{B}. \mathsf{strong-low-bisim-mm} \ \mathscr{B} \land \mathsf{EXTRA} \ \mathscr{B} \land$ 442 $(\langle tps, mds, mem_1 \rangle, \langle tps, mds, mem_2 \rangle) \in \mathscr{B})$ 443 444 445 We have proved in Isabelle/HOL that the compositionality theorem of Murray et al. 446 (2016b) holds regardless of the INIT, EXTRA chosen—in short, the INIT condition relaxes 447 the goal sufficiently to relax each of its assumptions the same way, and the EXTRA require-448 ment only strengthens its assumptions. Subject to some "sound mode use" side conditions 449 (to be discussed soon), it gives us that the parallel composition cms:: (*ThreadPrivate* \times 450 $(Mode \Rightarrow Var set))$ list of com-secure program threads will itself be a concurrent program 451 that enforces "sys-secure", a system-wide value-dependent noninterference property. Here, 452 the set operator returns the set of all the elements in a given list: 453 454 455 **Theorem 2.8** (Compositionality of com-secure *EXTRA*). 456 $\forall (tps, mds) \in set \ cms. \ com-secure_{INIT}^{EXTRA} \ (tps, mds)$ 457 $\forall mem. INIT mem \longrightarrow$ sound-mode-use (cms, mem) 458 459 sys-secure_{INIT} cms 460

We first introduce the elements of this whole-system property "sys-secure", before defining it formally (to follow, in Definition 2.9).

From all low-equivalent pairs of initial memories that both satisfy the *INIT* conditions, this whole-system property "sys-secure" asserts a form of low-equality between all global configuration pairs that are reachable via evaluation $-\rightarrow_{sched}$ to the same fixed schedule sched, for all such finite lists sched giving an order of steps of execution from each thread:

$$gc \dashrightarrow_{\square} gc' \triangleq (gc = gc')$$

$$gc \dashrightarrow_{n.ns} gc' \triangleq (\exists gc''. gc \rightsquigarrow_n gc'' \land gc'' \dashrightarrow_{ns} gc')$$

Here [] is an empty list, ... is the cons operator, and \sim_n means the *n*th thread in the global configuration takes one step.

In always comparing pairs of runs executing against the same schedule, the property models the class of schedulers whose decisions never depend on any secrets. Consequently, this excludes schedulers that are specialised, in the manner of Barthe *et al.* (2007a), to actively monitor the sensitivity level of each thread's control flow, so as to intervene and avoid interleaving it with others when it has become dependent on secrets. Instead, the CVDNI theory puts the onus on the developer of the program to prove that any branching on secret conditionals does not lead to timing-sensitive flows of the secret as discernible via low-classified sinks accessible to other threads in the system. Note that, as CVDNI prohibits mode state from ever becoming secret dependent, it will implicitly prohibit any leaks into parts of memory with which the mode state is directly associated—in Section 3, we will need to prohibit leaks into the memory we use to implement mutex locks, for this reason.

The special form of low-equality applied by the whole-system property is one that is modified from Definition 2.1, so that it only requires each Low-classified non-control vari-able $x \notin \mathscr{C}$ to be of equal value in both global configurations if the mode states of *all* threads consider x to be readable (Definition 2.2). Furthermore, the property ensures that paired global configurations continue to agree on the number of threads in the system, and on the mode states for all threads, written $cms'_1 = all-mds cms'_2 \triangleq (map mds cms'_1 = all-mds cms'_1$ map mds cms'_2), where the syntax "map mds cms" denotes the mapped projection that extracts a list of mode states from a list *cms* of *ThreadPrivate* \times (*Mode* \Rightarrow *Var set*) pairs. Finally, we will use syntax cms[i] to denote the *i*th element in list cms.

This whole-system noninterference property, written formally, is then as follows:

	Definition 2.9 (Whole-system value-dependent security, with <i>INIT</i> requirements).
507	
508	sys-secure _{INIT} $cms \triangleq \forall mem_1 \ mem_2$.
509	$INIT mem_1 \land INIT mem_2 \land mem_1 = {}^{Low} mem_2 \longrightarrow$
510	$(\forall sched \ cms'_1 \ mem'_1. \ (cms, mem_1) \dashrightarrow sched \ (cms'_1, mem'_1) \longrightarrow$
511	$(\exists cms'_2 mem'_2. (cms, mem_2) \dashrightarrow_{sched} (cms'_2, mem'_2)) \land$
513	$(\forall cms'_2 mem'_2. (cms, mem_2) \dashrightarrow _{sched} (cms'_2, mem'_2) \longrightarrow$
514	$length\ \mathit{cms}_1' = length\ \mathit{cms}_2'\ \land\ \mathit{cms}_1' = all-mds\ \mathit{cms}_2'\ \land$
515	$(\forall x. x \in \mathscr{C} \lor \mathscr{L} mem'_1 x = Low \land$
516	$(\forall i < \text{length } cms'_1. \text{ readable } cms'_1[i] x) \longrightarrow mem'_1 x = mem'_2 x)))$
517	
518	Finally, we must note that the use of assume-guarantee reasoning to obtain the com-
519	positionality of the per-thread property in the style of Mantel et al. (2011) gives rise to
520	requirements justifying the soundness of that reasoning; requirements that we will prove
521	our compiler to preserve. For CVDNI, these are summed up by the "sound-mode-use" side
523	condition of Theorem 2.8, which consists of a "local" and a "global" part:
524	Definition 2 10 (Cound mode use side condition)
525	Demitton 2.10 (Sound mode use side-condition).
526	sound-mode-use $(cms, mem) \triangleq$
527	$(\forall cm \in set\ cms.\ local-mode-compliance\ (cm, mem)) \land$
528	global-modes-compatibility (cms, mem)
529	
530	First, all threads must each obey a local mode compliance requirement. This says that
531	for all reachable local configurations of the program, at no point will the thread violate
532	any of its own guarantees not to access a particular location in the shared state, which
534	implies also not accessing any of its control variables. We leave precise definitions for "reachable les" and "desent read (or modified" to this paper's leabelle/HOL supplement
535	reachable-ics and doesni-read-(or-modify) to this paper's isabelie/HOL supplement material but mention here that it is the doesni-read-(or-modify) assertions that enforce
536	that any guarantees not to access some variable x will effectively apply also to all of x's
537	control variables:
538	
539	Definition 2.11 (Local mode compliance).
540	local-mode-compliance $lc \triangleq$
541	$\forall c m ds m am / c m ds m am) \in reachable les lc \rightarrow respects own guarantees (c m ds)$
542	where $\langle c, mas, mem \rangle \in \text{reachable-ics ic} \longrightarrow \text{respects-own-guarantees} (c, mas)$
543	where
545	respects-own-guarantees $(c, mds) \triangleq$
546	$(\forall x. (x \in mds \ \mathbf{GuarNoRW} \longrightarrow \text{doesnt-read-or-modify} \ c \ x) \land$
547	$(x \in mds \text{ GuarNoW} \longrightarrow \text{doesnt-modify } c x))$
548	
549	Then, all threads must together obey a <i>global modes compatibility</i> requirement. This
550	requirement says that the threads' mode states in all reachable global configurations of the
551	concurrent program (the reachable-mus-lists) are compatible—that is, it any one thread
552	

assumes a particular location will not be accessed for writing or reading, then all other threads must be guaranteeing not to access that location for the same purpose:

⁵⁵⁵ **Definition 2.12** (Global modes compatibility).

global-modes-compatibility $gc \triangleq \forall mdss \in reachable-mds-lists gc.$ compatible-modes mdsswhere

Note that this global modes compatibility requirement is *not compositional*; consequently, instead of obliging the program developer to prove it for the source programs to be fed to our compiler, we will prove it as an invariant maintained by the execution semantics of our source language—particularly, by its synchronisation primitives (see Section 3).

For more details and precise presentations of all the definitions we have adapted from Murray *et al.* (2016*b*,c) to enable the compiler verification work described in this paper, please refer to the Isabelle/HOL formalisation in our supplement material.

2.3 Cube-shaped refinement for preserving noninterference

Proof of *CVDNI-preserving refinement* (also *security-preserving* or *secure refinement*), for a single-threaded program that will be run as a thread of a concurrent program, requires the user of the theory to nominate two binary relations (both illustrated by Figure 2):

- 1. A *refinement relation* \mathscr{R} relating local configurations of the abstract program to local configurations of the concrete program: Abstract must simulate concrete, in a sense typical of much other work on program refinement, including compiler verification.
- 2. A concrete coupling invariant \mathscr{I} that allows us to use \mathscr{B} and \mathscr{R} to build a new strong low-bisimulation (modulo modes) for the concrete program, by discarding pairs of local configurations *after the refinement* that should not be reached in the same number of evaluation steps. It thereby witnesses that any changes a refinement (or compiler) might make to the execution time do not introduce any timing channels.

The essence of the proof technique is to require that a number of conditions—analogous to those for strong-low-bisim-mm (Definition 2.4)—be imposed on the nominated \mathscr{R} and \mathscr{I} , in relation to a given witness relation \mathscr{B} establishing com-secure (Definition 2.7) for the abstract program. The definitions to follow are adapted from Murray *et al.* (2016*b*) Section V, as we presented in Sison & Murray (2019)—for better readability, a simplified

500	if $h \neq 0$ then	reg3 = h	
599 600	x := y	if $reg3 \neq 0$ then	else
000	else	skip;	
602	x := y + z	skip;	reg2:=z;
602	fi	reg0:=y;	reg0 := reg1 + reg2;
603	(a) Abstract if -conditional.	x := reg0	x:=reg0
004	Relation \mathcal{R} pairs configurations of		fi
605	this program with configurations	(b) Concrete if-condition	nal. Relation <i>I</i> pairs configurations
607	are of the same-shaded region.	of this program as shown	n by the dashed lines.
608	Eise 2: Essentia from a CVDNU au		
609	Fig. 2: Excerpts from a C V DNI-pr	d - contain zoro, and w	ample with secret-dependent con-
610	trol flow: h contains a secret, y and z contain zero, and x is an untrusted sink. Reproduced		
611	from Sison & Murray (2019)—the	e example is originally	from Murray et al. (2016b).
612	version in which no new shared y	variables are added by	the refinement. Consequently, we
613	use the notation $-^{mem}$ to denote the	that two local configur	ations have equal mode state and
614	memory regardless of whether rel	lating configurations of	f the same or differing languages
615	Regarding the maintenance of t	nodes equivalence and	observational equivalence across
616	the relation the restrictions on refinement are tighter than those that were applied to		
617	strong-low-bisim-mm in that \mathscr{R} i	s required to preserve t	than those that were appred to be shared memory in its entirety:
618	strong for bisin hin, in that se	s required to preserve t	
619	Definition 2.13 (Preservation of m	nodes and memory).	
620	preserves modes mem	$\mathbb{R} \triangleq \forall l_{\mathcal{C}}, l_{\mathcal{C}} \in (l_{\mathcal{C}}, l_{\mathcal{C}})$	$c \in \mathbb{R} \longrightarrow lc = -m^{\text{mem}} lc =$
621	preserves-modes-mem 3	$e = vic_A ic_C. (ic_A, ic_C)$	$f) \in \mathcal{M} \longrightarrow \mathcal{M}_A{mds} \mathcal{M}_C$
622	Regarding the closedness under	r changes by other thre	ads that ensures compositionality
622 623	Regarding the closedness under for concurrency on <i>I</i> we again	r changes by other thre	ads that ensures compositionality
622 623 624	Regarding the closedness under for concurrency, on \mathscr{I} we again the base of \mathscr{R} with	r changes by other thre impose cg-consistent (eads that ensures compositionality (Definition 2.6) from Section 2.2.
622 623 624 625	Regarding the closedness under for concurrency, on \mathscr{I} we again However, in the case of \mathscr{R} , we cg-consistent that considers only	r changes by other thre impose cg-consistent (e instead impose "clo environmental actions	ads that ensures compositionality (Definition 2.6) from Section 2.2. osed-others", a simplification of that affect the memories on both
622 623 624 625 626	Regarding the closedness under for concurrency, on \mathscr{I} we again However, in the case of \mathscr{R} , we cg-consistent that considers only sides of the relation identically. F	r changes by other thre impose cg-consistent (e instead impose "clo environmental actions urthermore, closed-oth	ads that ensures compositionality (Definition 2.6) from Section 2.2. osed-others'', a simplification of that affect the memories on both
622 623 624 625 626 627	Regarding the closedness under for concurrency, on \mathscr{I} we again However, in the case of \mathscr{R} , we cg-consistent that considers only sides of the relation identically. F variables, not just those judged of	r changes by other thre impose cg-consistent (e instead impose "clo environmental actions urthermore, closed-oth pservable. Defined form	eads that ensures compositionality (Definition 2.6) from Section 2.2. osed-others", a simplification of that affect the memories on both hers ensures equality of <i>all</i> shared nally:
622 623 624 625 626 627 628	Regarding the closedness under for concurrency, on \mathscr{I} we again However, in the case of \mathscr{R} , we cg-consistent that considers only sides of the relation identically. F variables, not just those judged of	r changes by other thre impose cg-consistent (e instead impose "clo environmental actions urthermore, closed-oth oservable. Defined form	ads that ensures compositionality (Definition 2.6) from Section 2.2. osed-others", a simplification of that affect the memories on both hers ensures equality of <i>all</i> shared hally:
622 623 624 625 626 627 628 629 630	Regarding the closedness under for concurrency, on \mathscr{I} we again However, in the case of \mathscr{R} , we cg-consistent that considers only sides of the relation identically. F variables, not just those judged of Definition 2.14 (Closedness of re	r changes by other thre impose cg-consistent (e instead impose "clo environmental actions urthermore, closed-oth oservable. Defined form finements under chang	eads that ensures compositionality (Definition 2.6) from Section 2.2. osed-others", a simplification of that affect the memories on both hers ensures equality of <i>all</i> shared hally: es by others).
622 623 624 625 626 627 628 629 630 631	Regarding the closedness under for concurrency, on \mathscr{I} we again the However, in the case of \mathscr{R} , we cg-consistent that considers only sides of the relation identically. For variables, not just those judged of Definition 2.14 (Closedness of re- closed-others $\mathscr{R} \triangleq \forall tps_A tps_A$	r changes by other three impose cg-consistent (e instead impose "clo environmental actions urthermore, closed-oth oservable. Defined form finements under chang <i>c</i> mds mem mem'.	ads that ensures compositionality (Definition 2.6) from Section 2.2. osed-others", a simplification of that affect the memories on both hers ensures equality of <i>all</i> shared nally: es by others).
622 623 624 625 626 627 628 629 630 631 632	Regarding the closedness under for concurrency, on \mathscr{I} we again the However, in the case of \mathscr{R} , we cg-consistent that considers only sides of the relation identically. F variables, not just those judged of Definition 2.14 (Closedness of re- closed-others $\mathscr{R} \triangleq \forall tps_A tps_b$ $(\langle tps_A, mds, mem \rangle_A, \langle tps_b \rangle$	r changes by other three impose cg-consistent (e instead impose "clo environmental actions urthermore, closed-oth pservable. Defined form finements under chang $_{C}$ mds mem mem'.	eads that ensures compositionality (Definition 2.6) from Section 2.2. osed-others", a simplification of that affect the memories on both hers ensures equality of <i>all</i> shared hally: es by others).
622 623 624 625 626 627 628 629 630 631 632 633	Regarding the closedness under for concurrency, on \mathscr{I} we again a However, in the case of \mathscr{R} , we cg-consistent that considers only sides of the relation identically. F variables, not just those judged of Definition 2.14 (Closedness of re- closed-others $\mathscr{R} \triangleq \forall tps_A tps_A$ $(\langle tps_A, mds, mem \rangle_A, \langle tps_A tps_A$	r changes by other three impose cg-consistent (e instead impose "clo environmental actions urthermore, closed-oth pservable. Defined form finements under chang $_{C}$ mds mem mem'. $_{C}$, mds, mem \rangle_{C}) $\in \mathscr{R}$) /	eads that ensures compositionality (Definition 2.6) from Section 2.2. osed-others", a simplification of that affect the memories on both hers ensures equality of <i>all</i> shared hally: es by others).
622 623 624 625 626 627 628 629 630 631 632 633 634	Regarding the closedness under for concurrency, on \mathscr{I} we again However, in the case of \mathscr{R} , we cg-consistent that considers only sides of the relation identically. F variables, not just those judged of Definition 2.14 (Closedness of re- closed-others $\mathscr{R} \triangleq \forall tps_A tps_A$ $(\langle tps_A, mds, mem \rangle_A, \langle tps_A, (\forall x. (mem x \neq mem' x \lor$	r changes by other three impose cg-consistent (e instead impose "clo environmental actions urthermore, closed-oth pservable. Defined form finements under chang $_{C}$ mds mem mem'. $_{C}$, mds, mem \rangle_{C}) $\in \mathscr{R}$) (\mathscr{L} mem $x \neq \mathscr{L}$ mem'	ads that ensures compositionality (Definition 2.6) from Section 2.2. osed-others", a simplification of that affect the memories on both hers ensures equality of <i>all</i> shared nally: es by others). (A) (x) \longrightarrow writable <i>mds</i> x) \longrightarrow
622 623 624 625 626 627 628 629 630 631 632 633 633 634 635	Regarding the closedness under for concurrency, on \mathscr{I} we again the However, in the case of \mathscr{R} , we cg-consistent that considers only sides of the relation identically. For variables, not just those judged of Definition 2.14 (Closedness of re- closed-others $\mathscr{R} \triangleq \forall tps_A tps_A$ $(\langle tps_A, mds, mem \rangle_A, \langle tps_A tps_A, \langle tps_A, mds, mem' x \lor \langle \langle tps_A, mds, mem' \rangle_A, \langle tps_A, mds, mem' \rangle_A$	r changes by other three impose cg-consistent (e instead impose "clo environmental actions urthermore, closed-oth pservable. Defined form finements under chang $_{C}$ mds mem mem'. $_{C}$, mds, mem \rangle_{C}) $\in \mathscr{R}$) / \mathscr{L} mem $x \neq \mathscr{L}$ mem' S_{C} , mds, mem' \rangle_{C}) $\in \mathscr{R}$)	and that ensures compositionality (Definition 2.6) from Section 2.2. based-others'', a simplification of that affect the memories on both hers ensures equality of <i>all</i> shared hally: es by others). ($(x) \rightarrow (x) \rightarrow (x$
622 623 624 625 626 627 628 629 630 631 632 633 634 635 636	Regarding the closedness under for concurrency, on \mathscr{I} we again However, in the case of \mathscr{R} , we cg-consistent that considers only sides of the relation identically. F variables, not just those judged of Definition 2.14 (Closedness of re- closed-others $\mathscr{R} \triangleq \forall tps_A tps_A$ $(\langle tps_A, mds, mem \rangle_A, \langle tps_A, \langle tps_A, mds, mem' x \lor \langle \langle tps_A, mds, mem' \rangle_A, \langle tps_A, mds, mem' \rangle_A$	r changes by other three impose cg-consistent (e instead impose "clo environmental actions urthermore, closed-oth pservable. Defined form finements under chang $_{C}$ mds mem mem'. $_{C}$, mds, mem \rangle_{C}) $\in \mathscr{R}$) \mathscr{L} mem $x \neq \mathscr{L}$ mem' \mathscr{L}_{C} , mds, mem' \rangle_{C}) $\in \mathscr{R}$)	eads that ensures compositionality (Definition 2.6) from Section 2.2. osed-others", a simplification of that affect the memories on both hers ensures equality of <i>all</i> shared hally: es by others). $(x) \longrightarrow$ writable <i>mds</i> $x) \longrightarrow$
622 623 624 625 626 627 628 629 630 631 632 633 634 635 636 636 637	Regarding the closedness under for concurrency, on \mathscr{I} we again a However, in the case of \mathscr{R} , we cg-consistent that considers only sides of the relation identically. F variables, not just those judged of Definition 2.14 (Closedness of re- closed-others $\mathscr{R} \triangleq \forall tps_A tps_A$ $(\langle tps_A, mds, mem \rangle_A, \langle tps_A, \langle tps_A, mds, mem' \rangle_A, \langle tps_A, mds, mem' \rangle_A$	r changes by other three impose cg-consistent (e instead impose "clo environmental actions urthermore, closed-oth pservable. Defined form finements under chang $_{C}$ mds mem mem'. $_{C}$, mds, mem \rangle_{C}) $\in \mathscr{R}$) $f \ \mathscr{L}$ mem $x \neq \mathscr{L}$ mem' S_{C} , mds, mem' \rangle_{C}) $\in \mathscr{R}$) -requirement for confid-	ads that ensures compositionality (Definition 2.6) from Section 2.2. psed-others", a simplification of that affect the memories on both hers ensures equality of <i>all</i> shared hally: es by others). $(x) \rightarrow$ writable <i>mds</i> $x) \rightarrow$ dentiality preservation is to prove

programs, using a cube-shaped "refinement and coupling invariant preservation" diagram (coupling-inv-pres, depicted in Figure 3), whose edges are configuration pairs in \mathcal{B} , \mathcal{R} , and \mathcal{I} . (Reducing its difficulty is the focus of the decomposition principle in Section 2.4.)

All that then remains is for the nominated concrete coupling invariant \mathscr{I} to be symmetric (sym \mathscr{I}), and the predicate secure-refinement puts together all the requirements:



Fig. 3: Definition and graphical depiction of refinement preservation obligation for secure-refinement (Definition 2.15). Reproduced from Sison & Murray (2019)—the definition is a simplified restatement of its original formalisation in Murray *et al.* (2016*b*).

Definition 2.15 (Requirements for confidentiality-preserving secure refinement).

 $\begin{array}{l} \mathsf{secure-refinement} \ \mathscr{B} \ \mathscr{R} \ \mathscr{I} \ \triangleq \ \mathsf{preserves-modes-mem} \ \mathscr{R} \ \land \ \mathsf{closed-others} \ \mathscr{R} \ \land \\ \mathsf{cg-consistent} \ \mathscr{I} \ \land \ \mathsf{sym} \ \mathscr{I} \ \land \ \mathsf{coupling-inv-pres} \ \mathscr{B} \ \mathscr{R} \ \mathscr{I} \end{array}$

The soundness theorem for confidentiality-preserving refinement by Murray *et al.* (2016*b*) then gives us that, under these conditions, the concrete relation " \mathscr{B}_{C} of $\mathscr{B} \mathscr{R} \mathscr{I}$ ", derived from a witness strong-low-bisim-mm relation \mathscr{B} , refinement relation \mathscr{R} , and coupling invariant \mathscr{I} , is itself a witness strong-low-bisim-mm for the concrete program. For readability, from here onwards we will use $\mathbf{a}_1, \mathbf{c}_1, \ldots$ instead of lc_{1A}, lc_{1C}, \ldots for local configuration variables when comparing abstract and concrete executions simultaneously:

Definition 2.16 (Concrete bisimulation relation derived from \mathscr{B}, \mathscr{R} and \mathscr{I}).

$$\begin{split} \mathscr{B}_{\mathsf{C}} \mathsf{of} \ \mathscr{B} \ \mathscr{R} \ \mathscr{I} \ \triangleq \ \{(\mathbf{c}_1, \mathbf{c}_2) \mid \exists \mathbf{a}_1 \ \mathbf{a}_2. \ (\mathbf{a}_1, \mathbf{c}_1) \in \mathscr{R} \ \land \ (\mathbf{a}_2, \mathbf{c}_2) \in \mathscr{R} \land \\ (\mathbf{a}_1, \mathbf{a}_2) \in \mathscr{B} \ \land \ \mathbf{c}_1 = _{\mathsf{mds}}^{\mathsf{Low}} \mathbf{c}_2 \ \land \ (\mathbf{c}_1, \mathbf{c}_2) \in \mathscr{I} \} \end{split}$$

Theorem 2.17 (Preservation of strong-low-bisim-mm by secure-refinement).

 $\frac{\text{strong-low-bisim-mm }\mathcal{B} \quad \text{secure-refinement } \mathcal{B} \, \mathcal{R} \, \mathscr{I}}{\text{strong-low-bisim-mm } (\mathcal{B}_{\mathsf{C}} \text{of } \, \mathcal{B} \, \mathcal{R} \, \mathscr{I})}$

2.4 Decomposition principle and its impact on refinement proofs

We now present, as we first did in Sison & Murray (2019), an alternative way to prove secure-refinement (Definition 2.15) that obviates the need to use the cube-shaped, two-sided refinement obligation (depicted by Figure 3), by decomposing its concerns into:

1. Proving \mathscr{R} closed using a square-shaped simulation diagram (depicted by Figure 4a) akin to the *backward simulations* commonly used to prove semantics-preserving refinement by compilers (e.g. for CompCert (Leroy, 2009)), and



- stops $c_1 = stops c_2$ ensures that the refinement has not introduced any termination leaks, by asserting consistent stopping behaviour for I-related concrete program configurations, which we know to be observationally indistinguishable.
- 739 • abs-steps $\mathbf{a}_1 \mathbf{c}_1 = abs$ -steps $\mathbf{a}_2 \mathbf{c}_2$ ensures that the refinement has not introduced 740 any timing leaks, by asserting consistency of the pace of the refinement for 741 \mathcal{R} -related program configurations, which we again know to be observationally 742 indistinguishable. 743
 - The final \forall -quantified clause asserts \mathscr{I} 's suitability as a coupling invariant, in that it must remain closed under lockstep evaluation of the concrete program configurations it relates. Furthermore it must maintain mode state equality with each lockstep evaluation, which ensures that the refinement has not introduced any inconsistencies in the memory access assumptions and guarantees needed for the concurrent compositionality of the property.

Note that the \mathscr{B} - and \mathscr{R} -edges in Figure 4c may capture useful facts about a particular 750 program verification technique and compiler (respectively), so their availability as assump-751 tions is intended to reduce greatly the effort needed to specify a coupling invariant \mathscr{I} and 752 prove it satisfies the condition. 753

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

Theorem 2.20 (Soundness of the decomposition principle).

secure-refinement-decomp $\mathscr{B}\mathscr{R}\mathscr{I}$ abs-steps \implies secure-refinement $\mathscr{B}\mathscr{R}\mathscr{I}$

Proof The only obligation for secure-refinement (Definition 2.15) not obtained immediately from secure-refinement-decomp (Definition 2.18) is the cube-shaped coupling-inv-pres (Figure 3). We discharge this as follows:

The front face of the cube is just ordinary square-shaped refinement preservation (depicted in Figure 4a), given to us by secure-refinement-decomp: that a single concrete step from \mathbf{c}_1 is simulated by *n* abstract steps from \mathbf{a}_1 , where *n* is given by *abs-steps*.

We are then obliged to prove a simulation in the other direction (the back face of the cube), that n abstract steps from all configurations \mathbf{a}_2 related by \mathscr{B} to \mathbf{a}_1 are simulated by some concrete step from \mathbf{c}_2 related by \mathscr{R} to \mathbf{a}_2 and by \mathscr{I} to \mathbf{c}_1 .

Here, we lean on the determinism of the abstract program's evaluation semantics 772 (required by the theory) to flip the direction of simulation, knowing that *n* abstract steps 773 from \mathbf{a}_2 , simulating a single concrete step from \mathbf{c}_2 , could only be the very same *n* abstract 774 steps from \mathbf{a}_2 that we were required to consider. This allows us to obtain that simulation by using, once again, the square-shaped refinement preservation (Figure 4a) given to us by 776 secure-refinement-decomp. 777

Consistency of refinement pacing and stopping behaviour (depicted in Figure 4b) 778 given by decomp-refinement-safe (Definition 2.19) then respectively ensure that n (via 779 *abs-steps*) is the correct number of abstract steps to consider, and that there will indeed be 780 a concrete step from c_2 to drive the matching simulation step.

781 782

737

738

744

745

746

747

748

749

754

755

756 757 758

759

760 761 762

763

764

765

766

767

768

769

770

771

18 Verified Secure Compilation for Mixed-Sensitivity Concurrent Programs

Finally, the remainder of decomp-refinement-safe (depicted in Figure 4c) discharges the requirement of closedness and modes-equality maintenance of \mathscr{I} under lockstep execution, demanded by the bottom face of the cube.

To demonstrate how the decomposition principle reduces proof complexity and effort, we returned to the example program refinement discussed in Section V-E of Murray *et al.* (2016*b*) and proved in its Isabelle formalisation (Murray *et al.*, 2016*a*), an excerpt of which is shown in Figure 2. The abstract program (9 imperative commands) branches on a sensitive value, and executes a single atomic expression assignment in each branch. Its refinement (to 16 commands) models expansion of the expressions into multiple steps, resolving a timing disparity between the two branches by padding with **skip**.

We use proof size as a proxy for proof effort, since the former is known to be strongly 793 linearly correlated with the latter (Staples et al., 2014). Formalised in Isabelle/HOL as 794 EgHighBranchRevC.thy (Murray et al., 2016a), the proof line count for that theory 795 stood at about 4.6K lines of definitions and proof, of which approx. 3.6K line were proofs. 796 Adapting the proof instead to use the decomposition principle (secure-refinement-decomp, 797 Definition 2.18), the proof line count drops from 3.6K to approx. 2K, a 44% reduction. 798 Regarding definition changes, the new proof makes less than 10 lines of adaptations to a 799 coupling invariant and pacing function used by the old proof, and adds about 30 lines worth 800 of new helper definitions, for use with the decomposition principle. The rest of the theory 801 and its external dependencies remain in common between the two versions. 802

As would be expected, the bulk of the deletions are from the full cube-shaped refinement diagram proof (Figure 3) of secure-refinement (Definition 2.15) for the refinement relation. The surviving parts of that proof just become the square-shaped refinement diagram proof (Figure 4a) of the decomposition principle (Definition 2.18), without much modification. The deletions are replaced by newly added proofs of the decomposition principle's more security-focused side conditions (Definition 2.19, depicted by Figures 4b, 4c).

2.5 Compositional whole-system secure refinement

We now present the whole-system refinement theorem from Murray *et al.* (2016*b*,a), which we adapt here to support the specification of *INIT* requirements (as in com-secure EXTRA, Definition 2.7), and simplify to the case of refinements that add no shared variables.

The main usefulness of this theorem is that, beyond demanding secure-refinement (Definition 2.15), which dealt with the preservation of per-thread security as witnessed by a strong-low-bisim-mm (Definition 2.4), it deals additionally with the preservation of the sound-mode-use side conditions (Definition 2.10) that will be demanded by the compositionality theorem for CVDNI (Theorem 2.8) at the target language level.

Notably, although it imposes the requirement for the refinement to preserve the "local" part of sound-mode-use (Definition 2.11), it *automatically* preserves the noncompositional "global" part of this side condition (Definition 2.12) as a consequence of the requirements imposed by the per-thread secure refinements. Thus, our source-level proof of the global condition (see Section 3) will be sufficient, and there will be no need for us to prove anything extra about our compiler for it to preserve that to the target-language level. We now present the requirements and theorem for whole-system refinement formally.

827

809 810

811

812

813

814

815

816

817

818

819

783

784

785

First, in addition to the per-thread refinement notion secure-refinement (Definition 2.15) that we addressed in Sections 2.3 and 2.4, our whole-system refinement theorem will require that the refinement relation \mathcal{R} established by the compiler additionally preserves the compositional local mode compliance property for each thread. Here, "respects-own-guarantees" is from Definition 2.11:

835	Definition 2.21 (Refinement \mathscr{R} preserves local mode compliance).
836	preserves-local-compliance $\mathscr{R} \triangleq \forall tns, mds, mem, tns, mds, mem to the second secon$
837	preserves rocal compliance $\mathcal{D}^{e} = \mathcal{A}_{psA} \operatorname{mas}_{A} \operatorname{mas}_{A} \operatorname{ps}_{C} \operatorname{mas}_{A} \operatorname{mas}_{A}$
838	respects-own-guarantees $(tps_A, mds_A) \land$
839	$(\langle tps_A, mds_A, mem_A \rangle_A, \langle tps_C, mds_C, mem_C \rangle_C) \in \mathscr{R} \longrightarrow$
840	respects-own-guarantees (tps_C, mds_C)
841	

We define a new "compositional refinement" predicate to capture all per-thread requirements that will be demanded by our compositional whole-system refinement theorem. This bundles together preserves-local-compliance and secure-refinement so as to preserve the strong-low-bisim-mm relations (Definition 2.4) that witness noninterference for each thread of the abstract program. Alongside all these requirements just described, it also requires the concrete coupling invariant I to cover all possible initial memory pairs that are low-equal modulo modes (Definition 2.3) and satisfy the $INIT_C$ conditions that will parameterise the target language-level CVDNI property:

Definition 2.22 (Requirements for compositional whole-system refinement).

853	
054	compositional-refinement $\mathscr{B} \mathscr{R} \mathscr{I} \triangleq$
854	
855	secure-refinement $\mathscr{B} \mathscr{K} \mathscr{I} \land$ strong-low-bisim-mm $\mathscr{B} \land$
856	preserves-local-compliance \mathscr{R} \wedge
857	$(\forall tps_C \ mds \ mem_1 \ mem_2. \ mem_1 = _{mds}^{\sf Low} \ mem_2 \ \land \ INIT_C \ mem_1 \ \land \ INIT_C \ mem_2 \ \longrightarrow$
858	$(\langle tps_C, mds, mem_1 \rangle_{C}, \langle tps_C, mds, mem_2 \rangle_{C}) \in \mathscr{I})$

With these requirements, we prove using Isabelle/HOL that a whole-system refinement 861 theorem, proved originally by Murray et al. (2016b,a), can be adapted to support the spec-862 ification of INIT requirements on initial memory at both abstract- and concrete-level. (As 863 with Theorem 2.8, the relaxation of the goal by $INIT_C$ is enough to permit the relaxations of 864 its assumptions by $INIT_C$, $INIT_A$.) First we will state the theorem, then we will explain it, 865 line-by-line. This theorem proves that abstract-level sound-mode-use (including its global 866 part) by a system of secure mixed-sensitivity concurrent program threads (i.e. list cms_A, as 867 witnessed by bisimulations \mathscr{B} s for each thread) is sufficient for a set of per-thread secure 868 refinements (in terms of the lists $\mathscr{B}s$, $\mathscr{R}s$, $\mathscr{I}s$ of bisimulation, refinement, and concrete 869 coupling invariant relations for each thread, respectively) to yield a concrete-level secure 870 concurrent program (i.e. list *cms*_C that satisfies sys-secure, Definition 2.9): 871

872 873

829

830

831

832

833 834

842

843

844

845

846

847

848

849

850 851 852

8

859 860

875	Theorem 2.23 (Whole-system compositionality of per-thread secure refinement).
876	$(\forall mem. INIT_C mem \longrightarrow INIT_A mem) \land$
877	$(\forall mem. INIT_A mem \longrightarrow \text{sound-mode-use} (cms_A, mem)) \land$
878	length cms_A = length \mathscr{B} s = length \mathscr{R} s = length \mathscr{I} s = length $cms_C \land$
879	$(\forall i < \text{length } cms_C)$.
880	compositional-refinement $\mathscr{B}s[i] \mathscr{R}s[i] \mathscr{I}s[i] \wedge$
881	$(\forall mem, INIT_{C} mem \longrightarrow ((cms_{A}[i], mem), (cms_{C}[i], mem)) \in \mathscr{R}s[i]) \land$
882	$(\forall mem_1 mem_2 INIT_4 mem_1 \land INIT_4 mem_2 \land mem_1 = Low_{(a)} mem_2 \longrightarrow$
884	$((\operatorname{sum}_1 \operatorname{mem}_2), \operatorname{mem}_1) (\operatorname{sum}_1 \operatorname{mem}_1) \subset \mathscr{B}[i])$
885	$((\textit{cms}_{A}[\textit{l}],\textit{mem}_{1}),(\textit{cms}_{A}[\textit{l}],\textit{mem}_{2})) \in \mathscr{B}S[\textit{l}]))$
886 887	sys-secure _{INIT_C} cms _C
888 889	The premises of this theorem can be understood as follows:
890	• $(\forall mem. INIT_C mem \longrightarrow INIT_A mem)$:
891 892	The concrete-level " $INIT_C$ " initial condition must be no weaker than the abstract-level " $INIT_A$ " one.
893	• $(\forall mem. INIT_A mem \longrightarrow \text{sound-mode-use} (cms_A, mem)):$
894	For the abstract program, sound-mode-use (Definition 2.10) must hold for all
895	possible initial memories.
896	• The firsts of finitial infread-private and mode states at abstract and concrete level (resp. cms., cms.) and lists of bisimulation, refinement, and concrete coupling
897	$(1csp. cmsA, cmsC)$, and fists of distinuiation, remember, and concrete coupling invariant relations (resp. \mathcal{B}_S , \mathcal{A}_S , \mathcal{A}_S) must all be for the same number of threads
899	 Then, for all threads <i>i</i> in the system:
900	
901	- The relations $\mathcal{B}, \mathcal{H}, \mathcal{Y}$ for thread <i>i</i> must meet the requirements for "compositional whole-system refinement" (Definition 2.22)
902	- The refinement relation \mathscr{R} for thread <i>i</i> must hold initially, i.e. cover its initial
903	thread-private and mode states at concrete and abstract level (resp. cms_C , cms_A),
904	for all initial memories that satisfy the concrete $INIT_C$ requirement.
905	– The abstract bisimulation relation \mathcal{B} for thread <i>i</i> must hold initially, i.e. must
907	relate its initial thread-private and mode state to itself, for all pairs of memories
908	that are low-equal modulo that mode state, and that both satisfy the abstract $INIT_A$
909	requirement.
910	Given all these assumptions, Theorem 2.23 yields a whole-system noninterference
911	property sys-secure _{INIT_C} for the resulting concurrent program (with the list of initial
912	thread-private and mode states cms_C) that assumes that the initial memory satisfies Ini_C .
913	
915	3 Source language: While with mutex locks
916	
917	In this section, we give a focused presentation of our compiler's source language, centered on its properties that apple the composition of par thread proofs of CVDNL preserving
918	refinement to the compiler's target RISC language. Our Isabelle/HOL supplement provides
919	remement to the compiler's target it too language. Our isabene/from supprement provides
920	

While with mutex locks (hereafter While) is a generic imperative language with support for conditional looping, consisting of the commands *cmd* over arithmetic expressions *exp*:

924 925

921

922

923

- 926
- 927
- 928

 $exp ::= n | v | exp \oplus exp$ cmd ::= skip | cmd ; cmd | if exp then cmd else cmd fi |while exp do cmd od | v := exp | stop |
lock(k) | unlock(k)

929 The language is parameterised over shared program-variable identifiers v:: Var, shared 930 lock-variable identifiers k :: Lock, constant values n :: Val, and binary arithmetic operators 931 \oplus :: Val \Rightarrow Val \Rightarrow Val that each have a big-step evaluation semantics; these induce a big-932 step evaluation semantics for *exp* as a whole. The commands *cmd* then have a small-step 933 operational semantics, wherein **skip** and variable assignment v := exp execute in one step 934 to **stop** (which itself does not step to anything); conditional branch **if** steps to the appropri-935 ate *cmd* depending on whether its expression evaluates to zero; and conditional loop **while** 936 steps to an **if**-conditional between either (1) the loop body sequenced with a repetition of 937 the while command, or (2) stop. Finally the sequential command c_1 ; c_2 executes to c_2 938 when c_1 executes to **stop**, and to c'_1 ; c_2 (c'_1 being c_1 's destination) otherwise. Of these 939 aforementioned commands, only variable assignments can modify the shared memory 940 (program-variables only), and none can directly modify the mode state or lock-variables.

941 We will give special focus to the addition to the While language of the mutex 942 synchronisation primitives lock(k) and **unlock**(k), which are the sole means of modi-943 fying lock variables and mode state. These replace both the ad-hoc mode annotations 944 and the **await**(v) synchronisation primitive that were previously offered for While by 945 Murray et al. (2016b). After briefly noting here how While instantiates the underly-946 ing theory from Section 2, we will present these new primitives' operational semantics, 947 which depends on the program developer supplying details of the locking discipline as 948 a parameter (Section 3.1) subject to some restrictions (Section 3.2). We will then prove 949 that global-modes-compatibility (Definition 2.12) is invariant for systems of While pro-950 grams running concurrently (Section 3.3), subject to some initial conditions (Section 3.4). 951 Discharging this once-off noncompositional proof obligation is crucial in enabling both 952 composition of per-thread noninterference properties (using Theorem 2.8), and composi-953 tional whole-system secure refinement of noninterference down to RISC by our compiler 954 (using Theorem 2.23).

⁹⁵⁵ While instantiates the concurrent value-dependent noninterference theory described in ⁹⁵⁶ Section 2.2. This instantiation assumes that the underlying concurrent execution model ⁹⁵⁷ (e.g. operating system, scheduler) for the While language prevents threads from seeing ⁹⁵⁸ each others' current program location. Thus the While program command c :: cmd being ⁹⁵⁹ executed (understood as the current program location) is modelled as the thread-private ⁹⁶⁰ state of the local configuration triple: $\langle c, mds, mem \rangle_w$. (The subscript w distinguishes ⁹⁶¹ While program triples from RISC ones, which are subscripted r.)

To ease formalisation of lock(k) and unlock(k), we instantiate the shared *mem* :: *Mem* type as a total mapping from a sum type to values *Val*. This sum type, with constructors

964 965

962

963

Lock, Var, distinguishes lock-variable identifiers *k* :: *Lock* (which can only be read or written by the lock primitives) from program-variable identifiers *v* :: *Var* (which can be read or written by the rest of the commands). In Isabelle/HOL's datatype notation, this is:

 $Mem \triangleq (Lock \ Lock \ | \ Var \ Var) \Rightarrow Val$

For readability, we will elide this distinction between *Lock* and *Var*—or applications of their constructors Lock and Var—from the presentation whenever clear from the context.

3.1 Locking discipline and its semantics

The program developer provides the details of the program's locking discipline in the form of a *lock interpretation* parameter *lock-interp* :: *Lock* \Rightarrow (*Var set* \times *Var set*), which gives for each lock the two non-overlapping sets of program-variables over which acquiring the lock grants exclusive permission to write, (resp.) read and write. For readability, this presentation will elide *lock-interp* from the arguments of definitions, and use the notation *varsNoW*, *varsNoRW* :: *Lock* \Rightarrow *Var set* to refer to its fst and snd projection.

Alongside encoding the mutex primitives' usual effect on control flow—most crucially, lock(k) should refuse to proceed meaningfully if the lock k is already held—we will now specify for them an evaluation semantics that furthermore encodes the permissions implied by the locking discipline, as assumptions and guarantees expressed in the mode state. This semantics assumes that, initially, no locks are held, and all threads are making guarantees not to access the variables they govern (conditions we will define formally in Section 3.4).

The following two helpers specify how acquiring (resp. releasing) a lock affects the mode state under a given lock interpretation *lock-interp*. When a thread acquires a lock it gains more assumptions, and makes fewer guarantees about the region of memory concerned:

Definition 3.1 (Impact on mode state *mds* of acquiring lock *k*).

 $mds \oplus k \triangleq \lambda m.$ case m of **GuarNoW** \Rightarrow mds **GuarNoW** - varsNoW k | **AsmNoW** \Rightarrow mds **AsmNoW** \cup varsNoW k | **GuarNoRW** \Rightarrow mds **GuarNoRW** - varsNoRW k| **AsmNoRW** \Rightarrow mds **AsmNoRW** \cup varsNoRW k

The converse occurs when releasing a lock: the thread drops the assumptions it was making about that region of memory, and once again makes guarantees not to access it.

Definition 3.2 (Impact on mode state *mds* of releasing lock *k*).

 $mds \ominus k \triangleq \lambda m.$ case m of **GuarNoW** \Rightarrow mds **GuarNoW** \cup varsNoW k | **AsmNoW** \Rightarrow mds **AsmNoW** - varsNoW k | **GuarNoRW** \Rightarrow mds **GuarNoRW** \cup varsNoRW k| **AsmNoRW** \Rightarrow mds **AsmNoRW** - varsNoRW k

The operational semantics for **lock**(*k*) is then given by two rules: LOCKACQ when lock *k* is available, and LOCKSPIN when it is already held. For these, we use predicate ev_{Lock} :: *Val* \Rightarrow *bool* with designated constants $True_{Lock}$, $False_{Lock}$:: *Val* to indicate that the lock is, resp. is not held—i.e. ev_{Lock} ($True_{Lock}$) = True, and ev_{Lock} ($False_{Lock}$) = False.²

Apart from impacting the mode state as already specified (by Definition 3.1), attempting to acquire an available lock will succeed in the usual manner, setting the lock-variable to the designated constant ($True_{Lock}$) to prevent subsequent lock acquisition attempts:

$$\neg ev_{Lock} (mem (Lock k)) \qquad mem' = mem[Lock k \mapsto True_{Lock}]$$
$$\frac{mds' = mds \oplus k}{\langle lock(k), mds, mem \rangle_{w} \rightsquigarrow_{w} \langle stop, mds', mem' \rangle_{w}} LOCKACQ$$

Attempting to acquire an already-held lock results in a stuttering evaluation step:

$$\frac{\operatorname{ev}_{Lock} (mem (\operatorname{Lock} k))}{\langle \operatorname{lock}(k), mds, mem \rangle_{\mathsf{w}} \rightsquigarrow_{\mathsf{w}} \langle \operatorname{lock}(k), mds, mem \rangle_{\mathsf{w}}} \operatorname{LOCKSPIN}$$

Then, the operational semantics for **unlock**(k) is given by two rules, of which only one, LOCKREL, will ever be used by programs that follow locking discipline. This rule requires that the mode state *mds* is consistent with the present thread having previously acquired the lock k: In short, it should have all the assumptions, but none of the guarantees, associated with the variables governed by the lock. To specify this, we define the following helper:

Definition 3.3 (Mode state is consistent with holding a lock *k*).

lock-held-mds-correct *mds* $k \triangleq$ $\forall x. (x \in varsNoW \ k \longrightarrow x \notin mds \ GuarNoW \land x \in mds \ AsmNoW) \land$ $(x \in varsNoRW \ k \longrightarrow x \notin mds \ GuarNoRW \land x \in mds \ AsmNoRW)$

With that condition satisfied, the LOCKREL rule specifies that an unlock(k) will proceed successfully, to enact lock release on the memory and mode state as expected:

 $\frac{\mathsf{lock-held-mds-correct} \ mds \ k \ mem' = mem[\mathsf{Lock} \ k \mapsto \mathsf{False}_{Lock}]}{mds' = mds \ominus k} \ LockRel}$ $\frac{\mathsf{unlock}(k), mds, mem_{\mathsf{W}} \rightsquigarrow_{\mathsf{W}} \langle \mathsf{stop}, mds', mem' \rangle_{\mathsf{W}}}{\mathsf{Volume}}$

To ensure that the While evaluation semantics is defined for all possible configurations, the LOCKINVALID rule defines a stuttering evaluation step for attempts to **unlock**(k) that violate the locking discipline due to not having previously acquired the lock k:

$$\frac{\neg \text{ lock-held-mds-correct } mds \ k}{\langle \mathbf{unlock}(k), mds, mem \rangle_{\mathsf{w}} \rightsquigarrow_{\mathsf{w}} \langle \mathbf{unlock}(k), mds, mem \rangle_{\mathsf{w}}} \text{ LockInvalid}$$

As mode state is nominally a form of ghost state, having the operational semantics appear to depend on it in this manner is rather unusual. To remove the semantics' reliance on ghost state, the program developer must use a check for local-mode-compliance (Definition 2.11) that only ever admits programs that satisfy the lock-held-mds-correct

² All three of ev_{Lock} , $True_{Lock}$, $False_{Lock}$ are parameters that are set by the user of the theory, with the proviso that their choice of parameters satisfy that $ev_{Lock}(True_{Lock})$ and $\neg ev_{Lock}(False_{Lock})$ hold as required.

check whenever attempting to **unlock**(k). For such programs, the operational semantics is equivalent to one that (1) omits the lock-held-mds-correct check from the LOCKREL rule, and (2) omits the LOCKINVALID rule from the While-language semantics entirely. An example of such a check is included in our Isabelle/HOL supplement.

3.2 Restrictions on locking disciplines

Here we lay out some cleanliness conditions on locking disciplines, giving particular focus to those relevant to our locking semantics (Section 3.1), and to our verification efforts for While's global compositionality property (Section 3.3) and our compiler (Section 5).

Of these, only one is a hard consequence of the underlying CVDNI theory we presented in Section 2: The per-thread CVDNI property com-secure (Definition 2.7) effectively compels us to enforce that secrets are never allowed to leak into the locking state. Otherwise, mode state would become tainted upon any attempt to acquire a lock whose status is secret, which would violate com-secure's requirement that modes-equality must be maintained at all times (note the $=_{mds}$ enforced by strong-low-bisim-mm, Definition 2.4). To ensure that com-secure will always treat the locking state as an untrusted sink, we impose the following requirement on the \mathscr{L} parameter supplied by the program developer:

Proposition 3.4 (\mathscr{L} must permanently assign Low classification to all lock-variables *k*).

 $\forall k mem. \mathscr{L} mem (Lock k) = Low$

The remaining restrictions are consequences of various simplifications of convenience.

First, note that the type signature of the *lock-interp* parameter (given in Section 3.1) only allows locks to govern program variables, not other locks. We justify this simplification with the fact that if some lock k governed lock k', then k would already have to be held whenever acquiring k'—otherwise, the change to k' would violate a no-write assumption implied by the locking discipline. This, however, would make k' entirely redundant with k.

Second, the lock acquisition and release semantics we gave in Section 3.1 is rather simplified, in that releasing a lock will drop the assumptions of all its variables from the mode state, even if another lock for that variable is still held! Thus, we signal that it only works for disciplines wherein no more than one lock governs each program variable, by asserting:

Proposition 3.5 (No variable can be managed by more than one lock).

$$\forall v \ k. \ v \in varsNoW \ k \ \cup varsNoRW \ k \longrightarrow$$
$$(\forall k'. \ v \in varsNoW \ k' \ \cup varsNoRW \ k' \longrightarrow k' = k)$$

We believe that it would be feasible to relax Proposition 3.5, by generalising While's locking semantics to allow disciplines wherein multiple locks must be held to access a given variable. To satisfy CVDNI-preserving refinement (particularly Definition 2.13), a compiler would need to preserve the lock memory operations that implement the more sophisticated bookkeeping needed, as ours does for the current, much simpler locking semantics.

Next, we assume that the program developer has not specified any "vacuous" locks

(i.e. ones that govern no variables), and that all locks grant at most one of **AsmNoW** or **AsmNoRW** (i.e. not both) on any given variable. These two assumptions allow us to exclude various pathological cases from our reasoning in Section 3.3 and Section 5, respectively:

Proposition 3.6 (Every lock governs access to some variable).

 $\forall k. varsNoW k \cup varsNoRW k \neq \emptyset$

Proposition 3.7 (The lock interpretation sets for any given lock k do not overlap).

 $\forall k. varsNoW k \cap varsNoRW k = \emptyset$

The final two restrictions simplify the possible interactions between locks and control variables: We disallow locks from being control variables, and require variables to be governed by the same lock as their control variables. In particular, they will help us establish (in Section 5.4) that the compiler produces programs that satisfy local-mode-compliance.

First, recall we mentioned that, as part of local-mode-compliance (Definition 2.11), the doesnt-read-(or-modify) assertions entail that any guarantees not to access some variable v will effectively apply also to all of v's control variables. Disallowing lock-variables from being control variables thus ensures that **lock**(k) and **unlock**(k), because they only access lock-variable k, cannot violate doesnt-read-(or-modify) for any program-variables:

Proposition 3.8 (Lock-variables k cannot be control variables).

 $\forall k. (Lock k) \notin \mathscr{C}$

Finally, requiring variables to be governed by the same lock as their control variables effectively ensures they are always locked simultaneously. Apart from making it easier for programs to satisfy local-mode-compliance, this also naturally prevents leaks caused by other threads changing a variable's classification to Low when it still contains High data:

Proposition 3.9 (Variables are always governed by the same lock as their control variables).

$$\forall c \ v \ k. \ \mathsf{Var} \ c \in \mathscr{C} \mathsf{vars} \ (\mathsf{Var} \ v) \longrightarrow (c \in varsNoW \ k = v \in varsNoW \ k) \land$$
$$(c \in varsNoRW \ k = v \in varsNoRW \ k)$$

3.3 Proof of global modes compatibility as an invariant

This section will present proof that global-modes-compatibility (Definition 2.12) holds as an invariant for concurrent While programs (Section 3.3) when initialised to have no locks held (Section 3.4). Consequently, it is sufficient for a developer to use a local compliance check (Sison, 2020) to obtain the sound-mode-use condition (Definition 2.10) needed for per-thread security proofs to be compositional via Theorem 2.8.

1150

1110

1111

1112 1113 1114

1115

1116

1118

1119

1120

1121

1127

1129

1130 1131

1132

1133

1134

1135 1136

1137

Recall from Section 2.2 that this compatibility requirement formalises that for all reach-1151 able global configurations of a concurrent program, any assumptions made by any of the 1152 threads must be met by corresponding guarantees made by all of the other threads: 1153 1154 Definition 2.12 (Global modes compatibility). 1155 1156 global-modes-compatibility $gc \triangleq \forall mdss \in reachable-mds-lists gc.$ compatible-modes mdss1157 where 1158 reachable-mds-lists $gc \triangleq$ 1159 { $mdss \mid \exists cms' mem' sched. gc \rightarrow sched (cms', mem') \land map mds cms' = mdss$ } 1160 1161 compatible-modes $mdss \triangleq \forall i x. i < \text{length } mdss \longrightarrow$ 1162 $(x \in mdss[i] \text{ AsmNoRW} \longrightarrow$ 1163 $(\forall j < \text{length } mdss. \ j \neq i \longrightarrow x \in mdss[i] \ \text{GuarNoRW})) \land$ 1164 $(x \in mdss[i] \text{ AsmNoW} \longrightarrow$ 1165 1166 $(\forall i < \text{length } mdss. \ i \neq i \longrightarrow x \in mdss[i] \text{ GuarNoW}))$ 1167 1168 The approach to establish global-modes-compatibility here will be to define three mode 1169 management requirements that taken together imply compatible-modes, and to prove them 1170 invariant for concurrent While programs when initialised such that they hold to begin with. 1171 The first of these pertains to variables whose access is governed by some lock, accord-1172 ing to the locking discipline. To define it, we need, alongside lock-held-mds-correct 1173 (Definition 3.3) from Section 3.1, a predicate that specifies the correct mode state for not 1174 holding a lock k: It should make all of the guarantees, and have none of the assumptions 1175 associated with the variables governed by k^{3} Stated formally: 1176 1177 1178 **Definition 3.10** (Mode state is consistent with *not* holding a lock *k*). 1179 lock-not-held-mds-correct *mds* $k \triangleq$ 1180 $\forall x. (x \in varsNoW \ k \longrightarrow x \in mds \ GuarNoW \land x \notin mds \ AsmNoW) \land$ 1181 $(x \in varsNoRW \ k \longrightarrow x \in mds \ GuarNoRW \land x \notin mds \ AsmNoRW)$ 1182 1183 1184 Note that our simplifying exclusion of "vacuous" locks (Proposition 3.6) ensures 1185 we never have to deal with a case where lock-held-mds-correct mds k and 1186 lock-not-held-mds-correct mds k hold simultaneously. 1187 The requirement on global configurations regarding these lock-managed variables is 1188 then as follows: If and only if a given lock is held by anybody, then exactly one thread has 1189 a mode state consistent with holding it; furthermore, all other threads will have a mode 1190 state consistent with not holding it. Formally, with mdss $gc \triangleq map mds$ (cms gc): 1191

³ Note that this not merely the negation of lock-held-mds-correct mds k (Definition 3.3)!

1195 1196

1107	Definition 3.11 (Lock-managed variable modes are compatible with memory).
1197	lock-managed-modes-mem-compatible $gc \triangleq$
1199	$\forall k$, if $(e_{V,cck} ((mem gc) k))$ then
1200	$\exists i, i \leq \text{length} (\text{cms } gc) \land$
1201	$\Box : i < \text{length} (\text{chis} ge) / (\text{mass} ge)^{[i]} k \land$
1202	$()/(z + k_{1} + k_{2} + k_{3} + k_{3$
1203	$(\forall j < \text{length} (\text{cms } gc), i \neq j \longrightarrow$
1204	lock-not-held-mds-correct (mdss gc)[j] k)
1205	else $\forall i < \text{length} (\text{cms } gc).$
1206	lock-not-held-mds-correct (mdss $gc)[i]$ k
1207	
1208	The second requirement pertains to variables whose access is entirely ungoverned by
1210	any locks in the locking discipline. For these we specify a more direct check that if any
1211	thread in the global configuration has an assumption about access to any of these variables,
1212	then all other threads must be providing the corresponding guarantee to that assumption:
1213	$\mathbf{D}_{\mathbf{r}} = \mathbf{C}_{\mathbf{r}} + $
1214	Definition 3.12 (Unmanaged variable modes are compatible).
1215	unmanaged-var-modes-compatible $gc riangleq orall i x. i < length (mdss gc) \longrightarrow$
1216	$(x \notin [] varsNoRW k \longrightarrow$
1217	k::Lock
1218	$(x \in (mdss\ gc)[i]$ AsmNoRW \longrightarrow
1219	$(\forall j < length (mdss \ gc). \ j \neq i \longrightarrow x \in (mdss \ gc)[j] \ \mathbf{GuarNoRW}))) \land$
1220	$(x \notin [])$ varsNoW k \longrightarrow
1222	k::Lock
1223	$(x \in (mdss\ gc)[i]\ \mathbf{AsmNoW} \longrightarrow$
1224	$(\forall j < \text{length } (\text{mdss } gc). \ j \neq i \longrightarrow x \in (\text{mdss } gc)[j] \ \mathbf{GuarNoW})))$
1225	
1226	Also proved invariant is a third, minor property that enforces globally that no assump-
1227	tions or guarantees are ever recorded regarding access to lock-variables:
1228	
1229	Definition 3.13 (No assumptions and guarantees on lock variables).
1230	no-lock-mds $mds \triangleq \forall l m$. Lock $l \notin mds m$
1232	no-lock-mds-gc $ac \triangleq \forall mds \in set (mdss ac)$ no-lock-mds mds
1233	10^{-10} cm ⁻¹⁰ cm
1234	This follows trivially from (1) our simplification (discussed in Section 3.2) only to allow
1235	locks to protect access to program variables and not other locks, and (2) the resulting fact
1236	that no While primitives ever touch any mode state pertaining to lock variables. Thus,
1237	further details on this third management requirement will be elided.
1238	We then have straightforwardly from their definitions that together, these three mode
1239	management requirements imply compatible modes for a given global configuration:
1240	
1241	
1242	

lock-mar	aged-modes-mem-compatible gc unmanaged-var-modes-compatible gc no-lock-mds-gc gc
	compatible-modes (mdss gc)
Proofs o of an arbitr For the f	f invariance then proceed by induction over the single-step evaluation semantic ary thread taking a step to progress the system to a new global configuration. irst management requirement (Definition 3.11):
Lemma 3.	15 (Single-step preservation of lock-managed-modes-mem-compatible).
	$\begin{array}{l} lock-managed-modes-mem-compatible\;(\mathit{cms},\mathit{mem})\\ \langle c_i,\mathit{mds}_i,\mathit{mem}\rangle_{w}\;\;\sim_{w}\;\langle c_i',\mathit{mds}_i',\mathit{mem}'\rangle_{w}\;\;\;i$
	lock-managed-modes-mem-compatible (cms', mem')
Proof By <i>i</i> that is tak	induction over the single-threaded evaluation semantics of the program at indexing a step.
it is not alr consistent Similarly lock k's m with not ho The othe state nor a	eady set – it would then become the single unique thread to set lock k s memory ready set – it would then become the single unique thread whose mode state i with holding k . Otherwise, the mode states and memory remain unchanged. y, unlock (k) preserves the property because its only possible change is to unserve emory, and return the unique thread holding lock k to a mode state consistent olding k . er While commands preserve the property because they do not touch the mode by lock-variables.
For the s	econd management requirement (Definition 3.12):
Lemma 3.	16 (Single-step preservation of unmanaged-var-modes-compatible).
	$\label{eq:comparison} \begin{array}{l} \text{unmanaged-var-modes-compatible } (cms, mem) \\ \langle c_i, mds_i, mem \rangle_{w} \sim \sim_{w} \langle c_i', mds_i', mem' \rangle_{w} i < \text{length } cms \\ cms' = cms[i := (c_i', mds_i')] cms[i] = (c_i, mds_i) \end{array}$
	${\tt unmanaged-var-modes-compatible} \ ({\it cms', mem'})$
Proof Aga at index <i>i</i> the We provide pertaining commands have any effective	in, by induction over the single-threaded evaluation semantics of the program hat is taking a step. The and use lemmas that $lock(k)$ and $unlock(k)$ do not touch any mode state to variables that are unmanaged by any locks, and that the remaining While do not touch the mode state at all. Therefore evaluation steps cannot possibly effect on the compatibility of modes on these variables.
These sit step evaluation with the fa- in a straight	ngle-step evaluation results lift easily to invariance results over the global multi tition semantics quantified over arbitrary schedules. These invariance results ct that the management requirements ensure compatibility (Lemma 3.14), yield tforward manner the desired global compatibility invariance theorem:

Theorem 3.17 (Mode management requirements ensure global compatibility). 1289 lock-managed-modes-mem-compatible gc unmanaged-var-modes-compatible gc 1290 no-lock-mds-gc gc 1291 global-modes-compatibility gc 1292 1293 1294 3.4 Initial conditions ensuring global modes compatibility 1295 1296 We now define conditions on memory and mode state consistent with no locks being held, 1297 and show that initialising a system under these conditions is enough to satisfy the global 1298 compatibility part (Definition 2.12) of the sound-mode-use side condition (Definition 2.10) 1299 of the compositionality theorem for our security property (Theorem 2.8). 1300 We define the following predicate for initial memory: 1301 Definition 3.18 (A requirement for initial memory that no locks are held). 1302 1303 no-locks-held *mem* $\triangleq \forall k. \neg ev_{Lock} (mem k)$ 1304 1305 We then define an initial mode state $mds_0 :: Mode \Rightarrow Var set$ that provides all guaran-1306 tees demanded by the lock interpretation parameters varsNoW, varsNoRW (described in 1307 Section 3.1) for all lock variables in the system, and makes no assumptions: 1308 1309 **Definition 3.19** (Initial mode state mds₀). 1310 $\mathsf{mds}_0 \triangleq \lambda \ m. \ \mathsf{case} \ m \ \mathsf{of} \ \mathbf{GuarNoW} \Rightarrow \bigcup_{k::Lock} varsNoW \ k$ $| \ \mathbf{GuarNoRW} \Rightarrow \bigcup varsNoRW \ k$ 1311 1312 1313 k::Lock 1314 | AsmNoW $\Rightarrow \emptyset$ 1315 AsmNoRW $\Rightarrow \emptyset$ 1316 1317 We are then able to show that these conditions are enough to satisfy the requirements 1318 we just showed (in Section 3.3) ensure global modes compatibility for While: 1319 1320 Lemma 3.20 (Initialising with no-locks-held, mds₀ ensures global modes compatibility). 1321 1322 no-locks-held mem $\forall (c, mds) \in set \ cms. \ mds = mds_0$ 1323 global-modes-compatibility (*cms*, *mem*) 1324 1325 Proof Theorem 3.17 obliges us to show that the mode management conditions 1326 (Definitions 3.11, 3.12, and 3.13) hold. This follows straightforwardly from all the relevant 1327 definitions. 1328 1329 1330 4 Target language: RISC with mutex locks 1331 Here we introduce RISC with mutex locks (hereafter RISC), the target of our compiler. 1332 This is a generic RISC-style assembly language based on the RISC architecture targeted 1333 1334

by the compilation scheme of Tedesco *et al.* (2016). A RISC program text is a list of RISC instructions *I*, each optionally associated with a label:

1336 1337

1335

- 1338
- 1339
- 1340
- 1341

```
I ::= [l:]B

B ::= \text{Load } r v \mid \text{Store } v r \mid \text{Jmp } l \mid \text{Jz } l r \mid \text{Nop}

MoveK r n \mid \text{MoveR } r r \mid \text{Op } \oplus r r

LockAcq k \mid LockRel k
```

Here we fix the types of the constant values n: Val, binary arithmetic operators \oplus :: 1342 $Val \Rightarrow Val \Rightarrow Val$, shared program variables v :: Var, and shared lock variables k :: Lock to 1343 be the same as those for the source While language being compiled. Thus, the only new 1344 types here compared to Section 3 are for the register identifiers r :: Reg, and labels l :: Lab. 1345 RISC has a small-step operational semantics that is largely unchanged from Tedesco 1346 et al. (2016), in that each step updates a distinguished program counter register, which 1347 1348 captures the current thread's program location as an index into its RISC program text. The 1349 instructions MoveK, MoveR, Load, and Store, for moving values to and between the 1350 registers and shared memory, and the "no-op" instruction Nop, all increment the program counter; the "jump if zero" instruction Jz l r updates it to the index of the instruction at l 1351 if r contains zero (else increments it); the unconditional **Jmp** l does so unconditionally. 1352

Modifying this instruction set from Tedesco *et al.* (2016), we then customise the **Op** instruction, and add **LockAcq** *k* and **LockRel** *k* instructions, to have semantics mirroring those of \oplus , **lock**(*k*), and **unlock**(*k*) from While respectively. Whereas the RISC equivalents for the LOCKACQ and LOCKREL evaluation rules (described by Section 3.1 for the While language) increment the program counter, those for LOCKSPIN and LOCKINVALID leave it unchanged. There is no RISC evaluation rule that changes the program text.

Although it has only direct-addressing **Load** and **Store** instructions, our RISC target language is adequate for implementing all features of While present in Section 3, with the big-step semantics of *exp* replaced by small-step operations on registers. We relegate RISC's full formal semantics to this paper's supplement Isabelle/HOL material.

Our defining LockAcq k and LockRel k to have the same operational semantics on shared memory and mode state as While's lock(k) and unlock(k) has two consequences:

• Our compiler will expect the program developer to supply the details of the locking discipline for the While program being compiled, so as to be able to ensure that the RISC program it produces follows the same discipline.

 We then have that global compatibility is invariant for RISC execution, by a nearidentical argument to the one we presented in Section 3.3, when initialised with the conditions we presented in Section 3.4. This presents one option for obtaining RISC-level composition of per-thread noninterference properties; however, invoking it directly will not be necessary when using the compositional whole-system secure refinement method of Section 2.5. (The alternative options and their application will be demonstrated further, respectively in Section 5.4, Section 6.3.)

As for While in Section 3, we instantiate here for RISC the CVDNI theory of Murray *et al.* (2016*b*) as recalled in Section 2.2, assuming that the underlying concurrency model (e.g. OS, scheduler etc.) prevents one thread from reading the program text of another.

1379 1380

1376

1377

1378

1363

1364 1365

1366

1367

For RISC, we furthermore assume that the context switching mechanism ensures effectively that no thread can read or interfere with the contents of the registers (including the program counter) when active for another thread. Based on these assumptions, we model all three of the program counter register's value pc :: nat, RISC program text P :: I list, and register bank $regs :: Reg \Rightarrow Val$, as thread-private state in the local configuration triple: $\langle ((pc, P), regs), mds, mem \rangle_r$. (We use the subscript r to distinguish RISC triples.)

1387 1388 1389

1390

1400

1401

1402

1403

1404 1405

1406

1407

1408

1409

1410

1411

1412

1413

1414

1415

1416

1417

1426

5 Verified secure compiler for mixed-sensitivity concurrent While programs

This section presents the COVERN wr-compiler: the first compiler proved to preserve 1391 proofs of noninterference for mixed-sensitivity concurrent programs. By using assume-1392 guarantee modes (Mantel et al., 2011) and the decomposition principle of Section 2.4 to 1393 prove it introduces (resp.) no race conditions or timing leaks, we demonstrate the applica-1394 bility to compiler verification of the CVDNI-preserving refinement notion of Section 2.3 1395 originally posed by Murray et al. (2016b). Here the decomposition principle (Figure 4) is 1396 crucial because, in separating the concern of preventing new timing leaks, it avoids directly 1397 having to prove the cube-shaped refinement diagram (Figure 3) arising from its need to 1398 preserve a 2-safety hyperproperty (Terauchi & Aiken, 2005; Clarkson & Schneider, 2010). 1399

To preserve security for mixed-sensitivity concurrent programs, CVDNI-preserving refinement demands small-step preservation of *the contents of all shared memory loca-tions* including those that control value-dependent classifications and implement locks. As it is unusual for verified compilers to make such promises, we show that a valid approach is to take advantage of CVDNI's assume–guarantee framework to:

- 1. *test and preserve any absence of race conditions* implied (via the framework) by mutex lock-based synchronisation of access to such locations, and then
- 2. *use this absence of race conditions* to establish the small-step preservation of their contents demanded for security-preserving refinement.

In doing so, we prove that some optimisations the wr-compiler performs with its knowledge of the locking discipline—it avoids unnecessary Loads and recalculation of common subexpressions over shared memory when locked—are safe to allow without violating CVDNI.

In preserving CVDNI, the wr-compiler preserves security proofs that are produced by the program verification techniques of Sison (2020) for While with mutex locks, which in turn were adapted from Murray *et al.* (2016b,c). We will present such an application of our compiler, to a case study program verified using these techniques, in Section 6.

Section 5.1 will focus on the wr-compiler's particular adaptations to CVDNI (beyond 1418 the fault-resilient noninterference targeted by the original compilation scheme of Tedesco 1419 et al. (2016)), in the form of static checks and invariants that (resp.) test for and maintain 1420 the absence of race conditions on lock-protected shared variables. Section 5.2 formalises 1421 a ban, preserved by the wr-compiler, on secret-dependent control flow. Section 5.3 then 1422 presents formal proof (structured by our decomposition principle of Section 2.4) that the 1423 wr-compiler implements CVDNI-preserving refinement. Section 5.4 ultimately presents 1424 proofs of overall security preservation results useful to users of the wr-compiler: Namely, 1425

it can be used either to preserve security down to RISC for an entire concurrent Whilelanguage program, or to preserve the per-thread security for threads that will be run alongside others written directly in the RISC-language.

5.1 Preserving race-free expression evaluation

Recall from Section 2.3 that CVDNI-preserving refinement (Murray *et al.*, 2016*b*) demands that all shared memory contents be preserved, between each target- and source-language configuration that it relates. This is security critical for mixed-sensitivity concurrent programs, as it ensures that any future influence of those contents on value-dependent classifications (via control variables) or readability by other threads (in the case of the While and RISC languages, via lock variables) is preserved.

The wr-compiler's approach to preserving the contents of shared memory is to ensure:

- 1. That values calculated by expressions are *preserved* by compilation—that is, they have the same value when written back to shared memory (or conditionally branched on) by the RISC program, as they did in the original While program; and
- 2. That expression evaluation is *race-free*—that is, free of any race conditions with other threads that would render the calculated expression inaccurate.
- To this end, the wr-compiler requires of the original While program that whenever each thread attempts to evaluate an expression, it must hold locks ensuring the stability of *all* variables referenced by the expression.
- Thus, its knowledge and enforcement of the locking discipline is crucial, not only to show that its optimisations preserve CVDNI, but that *any meaningful operation* over shared memory preserves it. It therefore tests for *and rejects* programs that exhibit potential race conditions due to their failure to follow locking discipline—these result in a failed compilation.

The wr-compiler tracks two kinds of information to achieve these outcomes: the contents of registers as expressions over shared variables, and assumptions on access to variables by other threads. The structures the wr-compiler uses to do this are, respectively:

- A register record Φ:: RegRec ≜ Reg → exp. This draws inspiration from that used by the compilation scheme of Tedesco et al. (2016) (originally of type Reg → Var) to avoid generating unnecessary Load instructions to registers that already contain a variable; in addition, here we extend it to track entire expressions on shared variables.
- An assumption record S :: AsmRec ≜ (Var set × Var set) that tracks which variables at a given point in the source While program are "stable" due to having, respectively, an AsmNoW or AsmNoRW assumption.

The wr-compiler's main function compile-cmd then outputs every register-assumption record pair (or *compilation record*) $C = (\Phi, \mathscr{S}) :: CompRec \triangleq RegRec \times AsmRec asso$ ciated with the program state*before*execution of each instruction in the output RISCprogram.⁴ A typical invocation to compile some <math>c :: cmd takes an *initial compilation*

- ⁴ For readability, we will use regrec, asmrec to denote a *CompRec*'s (resp.) fst, snd projections.
- 1471 1472

1427

1428

1429 1430

1431 1432

1433

1434

1435

1436

1437

1438

1439

1440

1441 1442

1443

1444

1445

1446

1447

1454

1455

1456

1457

1458 1459

1460 1461

1462

1463

1464 1465

1466

1467

1468

record C, and returns the *CompRec*-annotated RISC program $PCs :: (I \times CompRec)$ list (i.e. map fst *PCs* recovers an unannotated RISC text), and a *final compilation record* C':

Example 5.1 (Example invocation of the COVERN wr-compiler).

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

The remainder of this section will focus on formal properties of the compilation records output alongside each RISC text: Section 5.1.1 will elaborate on checks enforced on input programs with the help of *AsmRecs*, and Section 5.1.2 will present a resulting property that *RegRecs* track stable expressions, needed to prove security preservation (in Section 5.3).

Remaining details (e.g. l, l', nl, nl' for label allocation) will be relegated to appendices that we provide as supplement material. We note here only that (1) compile-cmd may return True for *failed* to reject the input program, such as when it detects a race condition (described further in Section 5.1.1), or if expression depth exceeds the limit assumed by the register allocation scheme model (elided to Appendix B); also, (2) relative to the label allocation scheme (elided to Appendix A) we proved that the control flow of each program fragment compiled by the wr-compiler remains self-contained even when composed sequentially with other such fragments.

5.1.1 Requirements on inputs to the wr-compiler

We define a shared variable v to be recorded as assumed *stable* if it and all its control variables (i.e. $\mathscr{C}vars v$) cannot presently be written to by another thread—that is, if they are recorded as having either of **AsmNoW** or **AsmNoRW** active on them. Formally:

Definition 5.1 (Stability of variable *v* according to assumption record \mathscr{S}).

var-stable $\mathscr{S} v \triangleq v \in (\mathsf{fst} \ \mathscr{S} \cup \mathsf{snd} \ \mathscr{S}) \land (\forall v' \in \mathscr{C} \mathsf{vars} v. v' \in (\mathsf{fst} \ \mathscr{S} \cup \mathsf{snd} \ \mathscr{S}))$

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

Definition 5.2 (Stability of the register record in compilation record *C*).

regrec-stable $C \triangleq \forall e \in ran (regrec C). (\forall v \in exp-vars e. var-stable (asmrec C) v)$

We then implement a collection of stability-checks :: $cmd \times CompRec \Rightarrow bool$ (called no-unstable-exprs in Sison & Murray (2019)) as a recursive function on the structure of While programs, that compile-cmd will use to ensure the following requirements of the given *cmd* if started with a configuration consistent with the given *CompRec*:

• The first two requirements ensure that programs comply with the locking discipline:

34 Verified Secure Compilation for Mixed-Sensitivity Concurrent Programs

	- The requirement regarding reading from shared variables establishes the main
1519	outcome of freedom from race conditions we described at the beginning of
1520	Section 5.1: The program must not refer to expressions on any unstable variables.
1521	As a matter of convenience, instead of introducing a dedicated primitive to the
1522	While language for reading atomically from a single (otherwise-unstable) device
1523	memory location, in our case study model of Section 6 we model such interactions
1524	using a simple assignment $x := y$ protected by a "read-atomicity" lock on a shared
1525	memory location y that models the hardware interface. ⁵
1526	- If the program assigns to an unstable shared variable, then it must not be a lock-
1527	governed one according to the locking discipline. This prevents the violation of
1528	any guarantees not to write to the variable (due to not holding its lock).
1529	• The remaining two requirements follow some simplifying assertions, originally
1530	made by the security type system of Murray et al. (2016b), that ensure mode state
1531	remains consistent after conditional branching and looping:
1532	
1533	- The two sides of any H -conditional branches in the program must both end with,
1534	the mode state as contained by the accumulation accord
1535	The similar reasons we require any while loops in the program to restore the
1530	- For similar reasons, we require any while-loops in the program to restore the
1537	original set of locks held on loop entry (again, as captured by the assumption
1520	We believe these to be reasonable simplifications given that in our setting the set
1539	of variables governed by each lock does not shange at matime in such a way that
1540	of variables governed by each lock does not change at running in such a way that
1542	would require access to them to be lock protected (or not) in a conditional mannet.
1543	Together, regrec-stable C and stability-checks c C make up the main two require-
1544	ments of a predicate compile-cmd-input-reqs C l nl c imposed on the input arguments
1545	to compile-cmd. (Its other two requirements reflect that the terminated While program
1546	stop has no valid compilation, and that the initial label, if provided, must be valid-details
1540	are relegated to supplementary Appendix A.) If any of these requirements are violated,
1548	compile-cmd rejects the program with <i>failed</i> = True:
1549	
1550	Definition 5.3 (Requirements on inputs to compile cmd)
1550	bennuon 3.5 (Requirements on inputs to complication).
1552	compile-cmd-input-reqs $C \ l \ nl \ c \ riangle$ stability-checks $c \ C \ \wedge$ regrec-stable $C \land$
1553	$c \neq stop \land (\forall x. \ l = \text{Some } x \longrightarrow x < nl)$
1554	
1555	5.1.2 Proof that all tracked register contents are stable
1556	5.1.2 Trooj indi dii tracked register contents dre stable
1557	Imposing the predicate compile-cmd-input-reqs (Definition 5.3) gives us enough informa-
1558	tion to prove a lemma that compile-cmd only ever outputs stable register records, that attest
1559	to the fact that registers contain the results of evaluating expressions on stable variables.
1560	
1561	
1562	when such atomic hardware primitives exist on a given architecture, we expect it would be straightforward for source languages to expose them and oblige their architecture-specific compilers to compile them to that
1563	single atomic instruction in the target language's semantics, which would eliminate the need for such locks.
1564	

Stated more precisely, every RISC program returned by a successful invocation of compile-cmd is annotated by *CompRecs* all with stable register records, and furthermore that the final *CompRec*'s register record is also stable:

Lemma 5.4 (Successful compilations output only stable register records).

$$(PCs, l', nl', C', False) = compile-cmd C l nl c$$
 compile-cmd-input-reqs C l nl c

 $(\forall pc < \text{length } PCs. \text{ regrec-stable } (\text{map snd } PCs)[pc]) \land \text{ regrec-stable } C'$

Proof By induction on the structure of the While language program *c*, making reference to the implementation of compile-cmd.

For cases that must compile expressions, we furthermore prove and make use of a lemma by induction on the structure of expressions, making reference to the implementation of the expression compiler function compile-expr called by compile-cmd. In essence, we prove that (sub)expressions appearing in register records must be stable, for two reasons:

First, they are always only ever subexpressions over variables that must have been stable in the input program when their contents were first loaded into registers.

Second, when compiling an unlock(k), the wr-compiler will always flush all register records that make reference to any variables that the unlock(k) makes unstable.

5.2 Preserving a ban on secret-dependent control flow

The wr-compiler assumes that input While programs have no *conditional branches on* High-*sensitivity values* (High-*branching*), and therefore no secret-dependent control flow. This is a restriction commonly applied as a means to prevent all implicit flows, including timing leaks. This restriction will then be preserved by the wr-compiler for its output RISC programs, reflected primarily in the design of the concrete coupling invariant \mathscr{I}_{wr} (see Section 5.3.3).

Specifically, the wr-compiler assumes that the confidentiality of input While programs is witnessed by a strong low-bisimulation modulo modes with an extra requirement (supplied as a parameter, as in Section 2.2) that effectively disallows any present or past High-branching. Relying on the fact that a low-bisimulation already asserts Lowequivalence of memories, the extra requirement asserts that it furthermore pairs only configurations at the same program location, and that any **if**-conditional expressions must evaluate to the same value in both configurations' memories. Here, the helper function leftmost-cmd gives the leftmost in a sequence of ;-separated While-language commands:

Definition 5.5 (An extra requirement for low-bisimulations \mathscr{B} to ban High-branching).

¹⁶⁰³ no-high-branching $\mathscr{B} \triangleq$

 $\forall c \ c' \ mds \ mem \ mem'. \ (\langle c, mds, mem \rangle_{\sf w}, \langle c', mds, mem' \rangle_{\sf w}) \in \mathscr{B} \ \longrightarrow \ c = c' \ \land$

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

Then, in Section 5.4, we will prove that the wr-compiler produces confidential RISC programs with no secret-dependent control flow, as witnessed by a low-bisimulation that

asserts a similar extra requirement for RISC programs. In effect, this is the *pc-security* notion of Molnar *et al.* (2006), but also explicitly equating the program text:

Definition 5.6 (A pc-security–like requirement for RISC bisimulations *B*).

¹⁶¹⁵ pc-security $\mathscr{B} \triangleq \forall pc \ pc' \ P \ P' \ regs \ regs' \ mds \ mem \ mem'.$

 $(\langle ((pc, P), regs), mds, mem \rangle_{\mathsf{r}}, \langle ((pc', P'), regs'), mds, mem' \rangle_{\mathsf{r}}) \in \mathscr{B} \longrightarrow pc = pc' \land P = P'$

5.3 Use of the decomposition principle

Having covered the most relevant aspects of the wr-compiler's implementation, we now present the refinement relation \mathscr{R}_{wr} (in Section 5.3.1), pacing function abs-steps_{wr} (in Section 5.3.2), and concrete coupling invariant \mathscr{I}_{wr} (in Section 5.3.3), parameters we use to apply the decomposition principle we presented in Section 2.4 to prove (in Section 5.3.4) that successful compilations are legitimised by secure-refinement (Definition 2.15)—the desired confidentiality-preserving notion of refinement for mixed-sensitivity concurrent programs.

The strategy laid out by the decomposition principle will be to prove that these parameters satisfy decomp-refinement-safe (Definition 2.19) for a targeted class of input While-language programs—ones with no secret-dependent control flow, as we specified in Section 5.2—meaning (for such programs) we can use the parameters to enforce that Wrcompiler introduces no secret-dependent inconsistencies in termination, timing behaviour, or assume–guarantee modes.

In doing so we avoid a direct proof of the cube-shaped refinement diagram (Figure 3) of Murray *et al.* (2016*b*)—which would have involved reasoning about both \mathscr{R}_{wr} and \mathscr{I}_{wr} at once—and instead prove (with the assistance of abs-steps_{wr}) a square-shaped refinement diagram for \mathscr{R}_{wr} (Figure 4a) more typically found in compiler verification.

5.3.1 Refinement relation \mathscr{R}_{wr} and its invariants

In this section we introduce the refinement relation \mathscr{R}_{wr} that characterises compilation of programs from While to RISC using the wr-compiler, and prove it satisfies the two properties demanded of \mathscr{R}_{wr} (alone) by formal secure-refinement (Definition 2.15):

- 1. Preservation of modes and all contents of shared memory (preserves-modes-mem, Definition 2.13), and
- 2. Closedness under changes by other threads (closed-others, Definition 2.14).

An actual proof of refinement (using the square-shaped diagram of Figure 4a) for \mathscr{R}_{wr} will be deferred to Section 5.3.2, which introduces the abs-steps_{wr} function pacing it.

Just like the earlier example of a secure refinement relation (in Figure 2), the refinement relation \mathscr{R}_{wr} pairs abstract (here, While-language) with concrete (here, RISC-language) program configurations. For example, the if_expr case of \mathscr{R}_{wr} relates the expression-evaluation part of the While command if *e* then c_1 else c_2 fi, with the corresponding part of the RISC program obtained by running compile-cmd on it, including the conditional jump Jz after expression evaluation. (This case is depicted in Figure C.1, and a relevant excerpt of the compile-cmd implementation provided in Figure C.2 for comparison, both in

1655 1656

1611

1612 1613

1614

1616 1617 1618

1619

1634

1635

1636

1637 1638

1639

1640

1641

1642 1643

1644

1645

1646

1647

1648

1649

1650

1651

1652

1653

supplementary Appendix C. An informal description of all the cases of \mathscr{R}_{wr} , their purpose, and the invariants they maintain, can also be found in Appendix C.)

We define almost all the cases of \mathscr{R}_{wr} to assert at least one successful run of compile-cmd (where *failed* = False). We then define a guard that we impose to restrict the scope of \mathscr{R}_{wr} only to consider local program configurations consistent with the relevant compilation record produced by compile-cmd. In short, this ensures the actual values in the register bank *regs* equal any expression the register record says they should have, as evaluated under the current *mem*; and furthermore, that the assumption record is consistent with the **AsmNoW** and **AsmNoRW** modes in the actual *mds*. Formally:

Definition 5.7 (Configuration consistency requirements for compiled commands).

```
compiled-cmd-config-consistent C regs mds mem \triangleq
```

regrec-mem-consistent (regrec C) regs mem \land asmrec-mds-consistent (asmrec C) mds where

regrec-mem-consistent Φ regs mem $\triangleq \forall r \ e. \ \Phi \ r =$ Some $e \longrightarrow regs \ r = ev_{exp}$ mem e(Consistency between register record, register bank, and shared memory)

asmrec-mds-consistent \mathscr{S} mds $\triangleq \mathscr{S} = (mds \text{ AsmNoW}, mds \text{ AsmNoRW})$

(Consistency between an assumption record and a mode state)

Apart from using compiled-cmd-config-consistent to restrict the scope of \mathscr{R}_{wr} in this manner, we will also impose it in Section 5.3.4 as *initial configuration requirements* for compiled programs: Only configurations obeying them may be used to initialise a RISC program compiled by the wr-compiler with initial *CompRec C*.

The cases of \mathscr{R}_{wr} also tend to assert regrec-stable (Definition 5.2), which we already proved holds for all compilation records produced by the wr-compiler (Lemma 5.4).

Finally, whenever a case of \mathscr{R}_{wr} is inductive (e.g. the if_expr case, for its nested calls to compile-cmd for each of its "then" and "else" branches) it quantifies over all configurations that obey compiled-cmd-config-consistent (Definition 5.7) and regrec-stable (Definition 5.2) relative to the initial compilation record given to each nested call to compile-cmd.

With \mathscr{R}_{wr} thus specified, we can now prove the two requirements for secure-refinement that pertain to \mathscr{R}_{wr} alone: preserves-modes-mem (Definition 2.13), and closed-others (Definition 2.14). In short, preserves-modes-mem is largely enforced by the definition of \mathscr{R}_{wr} , but closed-others relies in part on \mathscr{R}_{wr} only ever talking about stable register records:

Lemma 5.8 (\mathscr{R}_{wr} preserves modes and memory).

preserves-modes-mem \mathscr{R}_{wr}

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

For all cases of $(lc_w, lc_r) \in \mathscr{R}_{wr}$, $lc_w =_{mds}^{mem} lc_r$ is either asserted directly by the guards or obtainable from the inductive hypothesis.

Lemma 5.9 (\mathscr{R}_{wr} is closed under changes by others).

closed-others	\mathscr{R}_{w}
---------------	-------------------

Proof By induction on the structure of \mathscr{R}_{wr} .

Changes by others (Definition 2.14) only modify writable variables the same way for both configurations, so preservation of $=_{mds}^{mem}$ is immediate. Also, regrec-mem-consistent is unaffected because by Lemma 5.4, compile-cmd only creates regrec-stable records i.e. referring to no writable variables. No other \Re_{wr} guards mention shared memory.

5.3.2 Refinement pacing function abs-stepswr

In this section we nominate a pacing function, $abs-steps_{wr}$, specifying the number of evaluation steps with which a While program should simulate each step of the RISC program to which the wr-compiler compiled it. Using the square-shaped "refinement preservation" diagram of Figure 4a (part of Definition 2.18), we then prove that the \Re_{wr} relation we introduced in Section 5.3.1 is a refinement when "paced" by abs-steps_{wr} in this manner.

Here we define abs-steps_{wr} to depend only on the current program location; consequently, as long as the wr-compiler introduces no secret-dependent control flow, it will also introduce no timing leaks—that is, no secret-dependent variations to the pacing of the program, as disallowed by Figure 4b (part of Definition 2.19)—which we will be obliged to prove in Section 5.3.4. To this end, abs-steps_{wr} primarily looks at the form of the RISC instruction (sometimes While command) about to be executed, dividing them into three categories:

- Instructions output by compile-expr: Load, Op, and MoveK. For these, abs-steps_{wr} returns 1 if the leftmost-cmd (the leftmost in a sequence of ;-separated commands) of the While program is "while e do c od", to allow it to step to "if e then (c; while e do c od) else stop fi" concurrently with the first RISC step of the compiled expression itself. Otherwise, abs-steps_{wr} returns 0, to indicate the While program standing still while the RISC program takes *new* steps to evaluate the expression.
 - "Epilogue" steps: **Jmp** and **Nop** when used for control flow at the end of a smaller compiled program in the context of a larger one. For these, abs-steps_{wr} returns 0.

• All other RISC instructions are assumed to proceed at a lockstep pace with the While command they were compiled from, and for these abs-steps_{wr} returns 1.

Having nominated abs-steps_{wr} and \mathcal{R}_{wr} , we now have the parameters over which we are obliged, by secure-refinement-decomp (Definition 2.18), to prove refinement preservation (Figure 4a). To this end, we prove firstly that every step of execution of a RISC program, produced by the wr-compiler from a While program, maintains the consistency demanded by compiled-cmd-config-consistent between configurations and compilation records:

Lemma 5.10 (Successfully compiled programs maintain config consistency requirements).

 $(PCs, l', nl', C', failed) = \text{compile-cmd} C l nl c \qquad \text{compile-cmd-input-reqs} C l nl c$ $failed = \text{False} \qquad pc < \text{length} PCs \qquad P = \text{map fst} PCs \qquad Cs = \text{map snd} PCs$ $\text{compiled-cmd-config-consistent} Cs[pc] \ regs \ mds \ mem$ $\langle ((pc, P), regs), mds, mem \rangle_{r} \rightsquigarrow_{r} \langle ((pc', P), regs'), mds', mem' \rangle_{r})$

compiled-cmd-config-consistent (if pc' < length P then Cs[pc'] else C') regs' mds' mem'

Proof Unfolding Definition 5.7, we in fact prove it separately for regrec-mem-consistent and asmrec-mds-consistent, both times by induction on the structure of While program *c*.

In each case, we use the simplifiers for the compile-cmd implementation to yield the corresponding RISC program fragment in question, and then prove the lemma for each of the possible locations of pc in the compiled program. For both proofs, there is some trickiness in accounting for (and ruling out) which destination pc' must be considered for each of these cases of pc, particularly for those While programs that compile to RISC programs that may have jumps in them.

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

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

Also, we must prove a correctness lemma for the expression compiler:

Lemma 5.11 (Correctness of the expression compiler).

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

Proof By induction on the structure of expressions e, using the simplification rules for the implementation of compile-expr, and also relying on assumptions of correctness of the register allocation scheme supplied by the instantiator of the theory.

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

Lemma 5.12 (\mathscr{R}_{wr} is a refinement paced by abs-steps_{wr}).

1790 1791 1792

1782

1787

1788 1789

 $\forall lc_w \ lc_r. \ (lc_w, lc_r) \in \mathscr{R}_{\mathsf{Wr}} \longrightarrow (\forall lc'_r. \ lc_r \rightsquigarrow_r lc'_r \longrightarrow (\exists lc'_w. \ lc_w \rightsquigarrow_w^{\mathsf{(abs-steps_{\mathsf{Wr}} \ lc_w \ lc_r)} \ lc'_w \ \land \ (lc'_w, lc'_r) \in \mathscr{R}_{\mathsf{Wr}}))$

Proof By induction on the structure of \mathscr{R}_{wr} . (Refer to supplementary Appendix C for an informal description of all cases of \mathcal{R}_{wr} .)

The base case stop is immediate, as it pertains to a terminated While and RISC pro-1797 gram. The base cases that proceed in one step to a terminating program configuration 1798 (skip_nop, assign_store, lock_acq, lock_rel) are fairly straightforward because after dealing with the single step, the resulting obligation can then be handled by the stop case. 1800 This leaves the last remaining base case assign_expr, which proceeds in one step either to itself, or to assign_store. In all these cases, we use Lemma 5.10 to obtain the preserva-1802 tion of the guards demanded by the \mathscr{R}_{wr} introduction rule for the destination configuration 1803 of the step. Particularly, the assign_store case must make use of regrec-mem-consistent 1804 and the correctness of compile-expr (Lemma 5.11) to ensure that once the evaluated expres-1805 sion is written back to shared memory, $lc'_w = \frac{mem}{mds} lc'_r$ holds as demanded by the stop 1806 case. 1807

The inductive cases that concern expression evaluation (if_expr, while_expr) are 1808 much like assign_expr in that they have the possibility of progressing in one step to 1809 themselves. Unlike assign_expr however, their other possibility is a conditional jump 1810 based on the result of that expression. Again we use Lemma 5.11 to obtain that the result is 1811 an accurate calculation of the expression, and this time we prove by the two different cases 1812 whether if_expr ends up in if_c1 or if_c2, or if while_expr ends up in while_inner 1813 or at stop (having jumped to the exit label). In these cases, the guards over which the 1814 inductive references to \mathcal{R}_{wr} have been quantified are versatile enough to discharge them-1815 selves (when *_expr steps to itself), or to discharge any reachable initial starting state for 1816 the nested compiled RISC program, given that Lemma 5.10 ensures the invariance of these 1817 guards. 1818

This just leaves the inductive cases that pertain to configurations inside a nested compiled RISC program (if_c1, if_c2, while_inner), or at the end of one (epilogue_step, while_loop). In these cases, the inductive hypotheses obtained from the inductive reference to \mathscr{R}_{wr} are always enough to satisfy the guards demanded by the possible destination cases. Like in the proof of Lemma 5.10, the trickiness mostly comes from accounting for all the possible cases of control flow (ruling out spurious destinations) that need to be considered.

5.3.3 Concrete coupling invariant \mathcal{I}_{wr}

The next element needed is the concrete coupling invariant \mathscr{I}_{wr} . Recall from Section 5.2 that the no-high-branching requirement (Definition 5.5) ensures that input While programs have no secret-dependent control flow; here we choose Iwr to ensure that the wr-compiler has not introduced any new secret-dependent control flow in the output RISC program.

We define \mathscr{I}_{wr} formally to assert that the witness strong low-bisimulation (modulo modes) to be derived for the output program only pairs local configurations that are at the same location pc = pc' of the same RISC program P = P':

Definition 5.13 (Concrete coupling invariant \mathscr{I}_{wr} for compiled programs).

 $\mathscr{I}_{\mathsf{wr}} \triangleq \{ (\langle ((pc, P), regs), mds, mem \rangle_{\mathsf{r}}, \langle ((pc', P'), regs'), mds', mem' \rangle_{\mathsf{r}}) \mid (pc, P) = (pc', P') \}$

1839 1840

1838

1795

1796

1799

1801

1819

1820

1821

1822

1823

1824

1825 1826

1827

1828

1829

1830

1831

1832

1833

From this definition, pc-security (Definition 5.6) is clearly immediate for any concrete bisimulation \mathscr{B}_{C} of $\mathscr{B} \mathscr{R} \mathscr{I}_{wr}$ (Definition 2.16) derived using \mathscr{I}_{wr} .

5.3.4 Proof of CVDNI-preserving refinement

With \mathscr{R}_{wr} , abs-steps_{wr}, and \mathscr{I}_{wr} nominated, we are ready to prove confidentialitypreserving refinement using the decomposition principle secure-refinement-decomp (Definition 2.18).

To this end, we now prove the suitability of these three parameters, for While programs that do not branch on High-sensitivity values (as we specified earlier, in Section 5.2):

Lemma 5.14 (\mathscr{R}_{wr} , abs-steps_{wr}, \mathscr{I}_{wr} are safe for secure-refinement decomposition).

 $\frac{\text{strong-low-bisim-mm }\mathcal{B} \quad \text{no-high-branching }\mathcal{B}}{\text{decomp-refinement-safe }\mathcal{B} \ \mathcal{R}_{\text{wr}} \ \mathcal{I}_{\text{wr}} \ \text{abs-steps}_{\text{wr}}}$

Proof Unfolding Definition 2.19 gives us the following obligations. (See also Figure 4.)

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

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

For closedness of \mathscr{I}_{wr} under lockstep execution, the only non-straightforward cases to consider are conditional branching, and the locking primitives. For conditional branching, we use no-high-branching for \mathscr{B} with memory preservation via \mathscr{R}_{wr} (Lemma 5.8) to ensure that the conditional branching outcome is the same on both sides.

Finally, as the only operations that touch mode state, the locking primitives are the only non-straightforward cases for modes-equality maintenance under lockstep execution. As all lock memory is classified Low (Proposition 3.4), we use strong-low-bisim-mm for \mathscr{B} with memory preservation via \mathscr{R}_{wr} to ensure the RISC configurations behave consistently.

Lemma 5.15 (\mathscr{R}_{wr} , abs-steps_{wr}, \mathscr{I}_{wr} meet decomposed secure-refinement requirements).

 $\frac{\text{strong-low-bisim-mm }\mathcal{B} \qquad \text{no-high-branching }\mathcal{B}}{\text{secure-refinement-decomp }\mathcal{B} \,\mathcal{R}_{\text{wr}} \,\mathcal{I}_{\text{wr}} \,\text{abs-steps}_{\text{wr}}}$

Proof Unfolding Definition 2.18, the obligations pertaining only to \mathscr{R}_{wr} and abs-steps_{wr} are discharged by Lemma 5.12, Lemma 5.9, and Lemma 5.8. Pertaining to \mathscr{I}_{wr} : Clearly \mathscr{I}_{wr} is symmetric, and furthermore it is cg-consistent (Definition 2.6) because the actions over which \mathscr{I}_{wr} must be closed modify only the shared memory, and \mathscr{I}_{wr} places only restrictions on the program text and current location. The final obligation (regarding decomp-refinement-safe) is discharged by Lemma 5.14.

From this it follows immediately via Theorem 2.20 that \mathscr{R}_{wr} with the help of \mathscr{I}_{wr} describes a confidentiality-preserving refinement for non-High-branching While programs:

Corollary 5.16 (\mathscr{R}_{wr} is a secure refinement for non-High-branching programs).

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

Finally we prove that successful compilation produces a RISC program related by \mathscr{R}_{wr} to its input While program, when started with corresponding (same *mds*, *mem*) and reasonable (according to compiled-cmd-config-consistent) initial configurations:

Theorem 5.17 (Successful compilations are refinements in \mathcal{R}_{wr}).

(PCs, l', nl', C', failed) =compile-cmd C l nl c compile-cmd-input-reqs C l nl cfailed =False compiled-cmd-config-consistent C regs mds mem P = map fst PCs

 $(\langle c, mds, mem \rangle_{w}, \langle ((0, P), regs), mds, mem \rangle_{r}) \in \mathscr{R}_{wr}$

Proof By induction on the structure of the While-language.

The compiler input and initial configuration conditions we impose allow us to have each of **skip**, *cmd*; *cmd*, **if** *exp* **then** *cmd* **else** *cmd* **fi**, **while** *exp* **do** *cmd* **od**, v := exp, **lock**(k), and **unlock**(k) and their compiled output meet the guards of the introduction rules for the cases skip, seq, if_expr, while_expr, assign_expr, lock_acq, and lock_rel of \Re_{wr} (described further in supplementary Appendix C) that we designed for them, respectively.

5.4 Proof of compositional noninterference preservation

Going beyond the level of detail of our presentation in Sison & Murray (2019), we now present the final few steps to obtain preservation of whole-system security for concurrent compositions of RISC threads when all are obtained via compilation by the wr-compiler (Section 5.4.1). In addition to this, we obtain preservation of per-thread compositional security for each program thread compiled, and other properties that may be useful for their composition with RISC threads proved secure directly at the RISC level (Section 5.4.2).

5.4.1 Whole-system security preservation

To use the whole-system refinement theorem (Theorem 2.23), we are obliged to show that, in addition to establishing a secure-refinement (Definition 2.15, which we just showed in Section 5.3), the wr-compiler also preserves local-mode-compliance as demanded by compositional-refinement (Definition 2.22). Then, as we noted in Section 2.5, there is no need for us to prove preservation of the non-compositional global-modes-compatibility condition—the whole-system refinement theorem takes care of that.

The local compliance preservation result follows from a property of the refinement relation, \mathscr{R}_{wr} . Here, "respects-own-guarantees" is from Definition 2.11:

Lemma 5.18 (Each step from a RISC configuration in \mathscr{R}_{wr} respects its own guarantees).

$$\langle c, mds, mem \rangle_{\mathsf{w}}, \langle ((pc, P), regs), mds, mem \rangle_{\mathsf{r}}) \in \mathscr{R}_{\mathsf{w}}$$

respects-own-guarantees (((pc, P), regs), mds)

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

(

Knowing that the While command does not access lock-governed variables without holding the relevant lock (via the stability-checks asserted as part of compile-cmd-input-reqs by every relevant case of \mathcal{R}_{wr}), we are obliged to show that the RISC instruction paired to it by \mathcal{R}_{wr} similarly respects the guarantee modes implied by the locking discipline (as specified in Section 3.1). We do so with a mixed Isar/"apply"-style proof that exercises the relevant cases of the RISC semantics, using lemmas about control flow under sequential composition (mentioned in Section 5.1; see also supplementary Appendix A). Propositions 3.8 and 3.9 also play a role in excluding certain cases from consideration.

Lemma 5.19 (Refinements in \mathcal{R}_{wr} ensure local mode compliance).

 $(\langle c, \textit{mds}, \textit{mem} \rangle_{\sf w}, \langle ((\textit{pc}, \textit{P}), \textit{regs}), \textit{mds}, \textit{mem} \rangle_{\sf r}) \in \mathscr{R}_{\sf wr}$

local-mode-compliance $\langle ((pc, P), regs), mds, mem \rangle_r$

Proof Unfolding Definition 2.11, we must show that what was proved by Lemma 5.18 holds for every RISC configuration reachable from $\langle ((pc, P), regs), mds, mem \rangle_r$.

First, we prove a lemma that establishes that every such reachable RISC configuration is also paired by \mathscr{R}_{wr} to some While configuration. Specifically, we prove that \mathscr{R}_{wr} is closed under a notion of "pairwise reachability under mode-permitted havoc", wherein:

- 1. Every one step by the RISC program is matched by either zero or one step by the While program, as specified by abs-steps_{wr} (Section 5.3.2).
- 2. Between each evaluation step, arbitrary changes are allowed to occur to the memory locations judged by the mode state to be writable (Definition 2.5).

Because all such RISC configurations reachable from the initial one are in \mathscr{R}_{wr} , it then follows from Lemma 5.18 that they respect their own guarantees, as required.

We then initialise the compiler with an empty C_0 :: *CompRec* that knows nothing about the register contents, and does not assume any variables to be stable:

Definition 5.20 (Empty compilation record C₀).

 $C_0 \triangleq ((\lambda_- . None), (\emptyset, \emptyset))$

With these definitions we have the desired consistency result:

Lemma 5.21 (Initial C₀, mds₀ are consistent with no-locks-held).

no-locks-held $mem \implies$ compiled-cmd-config-consistent C₀ regs mds₀ mem

Proof This is straightforward by unfolding Definitions 5.7, 3.18, 3.19, and 5.20, also relying on the cleanliness conditions Proposition 3.5 and Proposition 3.7 on locking disciplines specified in Section 3.2.

We now have enough information to derive a whole-system security result, for 1982 concurrent RISC programs obtained by running the wr-compiler on any list "cs" of 1983 secure While commands (one for each thread in the program). As we explained in 1984 Section 5.2, the COVERN wr-compiler's preservation of security is only for programs with 1985 no-high-branching (Definition 5.5); furthermore, so that we can derive global compatibil-1986 ity for multiple of these programs run concurrently as threads (as per Section 3.4), we 1987 will impose no-locks-held (Definition 3.18) as an initial condition. Therefore, the secu-1988 rity preservation theorem we choose to prove here demands that every thread of the input 1989 While program be com-secure^{no-high-branching} (Definition 2.7, with additional requirements 1990 as specified). It then promises that the output program is sys-secure_{no-locks-held}: 1991

Theorem 5.22 (Secure threads compiled by the wr-compiler form a secure system).

 $\underset{1995}{\overset{1994}{\text{ length } cms_r = \text{length } cs \land}}$

1979

1980

1981

1992 1993

1996

1997 1998

1999

2000

2001

2003 2004

2005 2006

2007

2008

2009

2010

2011 2012

2013

2014

2015

2016

2017

2018

2019

2020

2021

2022

 $\forall i < \text{length } cms_r. \exists l \ nl \ PCs \ l' \ nl' \ C' \ regs.$

 $\mathsf{com}\text{-}\mathsf{secure}^{\mathsf{no}\text{-}\mathsf{high}\text{-}\mathsf{branching}}_{\mathsf{no}\text{-}\mathsf{locks}\text{-}\mathsf{held}} \ (\mathit{cs}[i], \mathsf{mds}_0) \ \land \\$

 $(\forall \textit{mem}. \text{ no-locks-held } \textit{mem} \longrightarrow \textsf{local-mode-compliance } \langle c, \textsf{mds}_0, \textit{mem} \rangle_{\sf w}) \land$

 $(PCs, l', nl', C', False) = \text{compile-cmd } C_0 \ l \ nl \ cs[i] \land \text{compile-cmd-input-reqs } C_0 \ l \ nl \ cs[i] \land cms_r[i] = (((0, \text{map fst } PCs), regs), \text{mds}_0)$

sys-secure_{no-locks-held} cms_r

Proof We invoke Theorem 2.23, supplying:

• no-locks-held for the *INIT* parameter at both While and RISC level.

• $\mathscr{B}_{all}, \mathscr{R}_{wr}, \mathscr{I}_{wr}$ to be respectively the witness bisimulation, refinement relation, and coupling invariant for all compiled threads, where we define \mathscr{B}_{all} to be the arbitrary union of all strong low-bisimulations modulo modes that disallow high-branching:

 $\mathscr{B}_{\mathsf{all}} \triangleq \bigcup \{ \mathscr{B} \mid \mathsf{strong-low-bisim-mm} \ \mathscr{B} \land \mathsf{no-high-branching} \ \mathscr{B} \}$

• mds₀ to be the initial mode state for all While threads in *cs*.

The first thing we must prove is that the original program satisfies sound-mode-use (Definition 2.10) when initialised with mds_0 and no-locks-held; we have the local part from this theorem's local-mode-compliance assumption, and the global part from Lemma 3.20.

We then discharge the demands of compositional-refinement $\mathscr{B}_{all} \mathscr{R}_{wr} \mathscr{I}_{wr}$ (Definition 2.22) using Corollary 5.16, Lemma 5.19, and by unfolding Definition 5.13.

It only remains for us to show that the initial RISC–While and While–While configuration pairs of interest are captured respectively by \mathscr{R}_{wr} and \mathscr{R}_{all} . We obtain the former using this theorem's assumptions and Lemma 5.21 to discharge the guards of Theorem 5.17. Finally, we use the assumption that the original program is com-secure^{no-high-branching} and unfold Definition 2.7 to obtain that there exists some strong-low-bisim-mm \mathscr{B} that enforces

no-high-branching for every configuration pair with low-equal memories (modulo mds_0) and no-locks-held initially; therefore, these state pairs must all be captured by \mathscr{B}_{all} .

5.4.2 Per-thread compositional security preservation

For system developers who may want to run programs compiled from While to RISC concurrently with other programs written directly in RISC, per-thread security preservation results may be useful. To compose the security proofs for those threads, direct RISClevel lemmas for the "sound-mode-use" side conditions of the compositionality theorem (Theorem 2.8) will also be needed. We therefore present these as an alternative method to obtain compositional security results for RISC programs, applicable when only partially produced by compilation from While by the wr-compiler.

Given the facts we established in Section 5.3, we have straightforwardly that such programs' executions are captured by the bisimulation derived from $\mathcal{B}, \mathcal{R}_{wr}, \mathcal{I}_{wr}$, when started with reasonable initial configurations corresponding to those paired by \mathcal{B} :

Lemma 5.23 (Programs witnessed by \mathscr{B} are captured by \mathscr{B}_{C} of $\mathscr{B} \, \mathscr{R}_{wr} \, \mathscr{I}_{wr}$ once compiled).

strong-low-bisim-mm \mathscr{B} $(\langle c, mds, mem_1 \rangle_w, \langle c, mds, mem_2 \rangle_w) \in \mathscr{B}$

(PCs, l', nl', C', failed) =compile-cmd C l nl c compile-cmd-input-reqs C l nl cfailed = False compiled-cmd-config-consistent $C regs mds mem_1$ P =map fst PCs compiled-cmd-config-consistent $C regs mds mem_2$

 $(\langle ((0, P), regs), mds, mem_1 \rangle_r, \langle ((0, P), regs), mds, mem_2 \rangle_r) \in \mathscr{B}_{\mathsf{C}} \mathsf{of} \ \mathscr{B} \ \mathscr{R}_{\mathsf{wr}} \ \mathscr{I}_{\mathsf{wr}}$

Proof Straightforward from the definition of \mathscr{B}_{C} of (Definition 2.16), using Theorem 5.17 to show membership of \mathscr{R}_{wr} , and the definition of strong-low-bisim-mm (Definition 2.4) to show that the memories are low-equal modulo modes, as required by \mathscr{B}_{C} of. Finally, membership of \mathscr{I}_{wr} (Definition 5.13) follows from the fact that the paired configurations are at the same location (program counter 0) of the same program *P*.

We are ready to state the per-thread security preservation result formally. Given an input While command that satisfies com-secure $\frac{no-high-branching}{no-locks-held}$ with mds₀ initially, it promises that the RISC program output by the wr-compiler is com-secure $\frac{pc-security}{no-locks-held}$ with mds₀:

Theorem 5.24 (Preservation of per-thread confidentiality by the wr-compiler).

 $\begin{array}{c} \mathsf{com-secure}_{\mathsf{no-locks-held}}^{\mathsf{no-high-branching}} \ (c, \mathsf{mds}_0) \\ (\mathit{PCs}, \mathit{l}', \mathit{nl}', \mathit{C}', \mathsf{False}) = \mathsf{compile-cmd} \ \mathsf{C}_0 \ \mathit{l} \ \mathit{nl} \ c \\ \mathsf{compile-cmd-input-reqs} \ \mathsf{C}_0 \ \mathit{l} \ \mathit{nl} \ c \end{array}$

 $\mathsf{com}\operatorname{-secure}_{\mathsf{no-locks-held}}^{\mathsf{pc}\operatorname{-security}}(((0, \mathsf{map} \mathsf{fst} PCs), regs), \mathsf{mds}_0)$

Proof We are given by com-secure^{no-high-branching} (Definition 2.7) that for low-equal starting configurations (modulo modes) of c with no locks held, there exists some witness \mathscr{B} satisfying both strong-low-bisim-mm and no-high-branching.

From this and Lemma 5.23 we have that the output program's corresponding execution is captured by a RISC semantics-level relation \mathscr{B}_{C} of $\mathscr{B}_{Wr} \mathscr{I}_{Wr}$ derived from this \mathscr{B} , with Lemma 5.21 discharging the compiled-cmd-config-consistent requirements.

Corollary 5.16 then gives us that secure-refinement \mathscr{B} \mathscr{R}_{wr} \mathscr{I}_{wr} holds, and from this and strong-low-bisim-mm \mathscr{B} using Theorem 2.17 we have strong-low-bisim-mm (\mathscr{B}_{C} of \mathscr{B} \mathscr{R}_{wr} \mathscr{I}_{wr}). This is enough to show com-secure^{pc-security} no-locks-held for the RISC program, by Definition 2.7; as Section 5.3.3 noted, pc-security (Definition 5.6) is immediate from the definition of \mathscr{I}_{wr} .

To prove a whole-system security result at the RISC level for the compiled program, we must also prove sound-mode-use (Definition 2.10). To that end, we prove a local and global result for RISC programs output by the wr-compiler when given a secure While program. The former follows from the local compliance result in the preceding section:

Lemma 5.25 (Threads compiled by the wr-compiler obey local compliance).

 $(PCs, l', nl', C', False) = compile-cmd C_0 l nl c$ compile-cmd-input-reqs C_0 l nl c no-locks-held mem

local-mode-compliance $\langle ((0, map \text{ fst } PCs), regs), mds_0, mem \rangle_r$

Proof We use Theorem 5.17 and Lemma 5.21 to obtain membership in \mathscr{R}_{wr} , which then allows us to use Lemma 5.19.

Then we prove invariance of global modes compatibility (as in Section 3.3) for compiled RISC programs, due to RISC's identical semantics to While regarding locking and modes:

Lemma 5.26 (Initialising RISC with no-locks-held, mds₀ ensures global compatibility).

 $\frac{\text{no-locks-held } mem \qquad \forall (((pc, P), regs), mds) \in \text{set } cms_r. mds = \text{mds}_0}{\text{global-modes-compatibility } (cms_r, mem)}$

Proof We firstly prove versions of Lemma 3.15, Lemma 3.16, and Theorem 3.17 for RISC, following exactly the same reasoning as we did in Section 3.3 for While. This is because the RISC instructions LockAcq k and LockRel k are (like lock(k) and unlock(k) in While) the only ones in their language that modify mode state, and their semantics regarding mode state and lock memory are identical to those of the lock(k) and unlock(k) commands. The present result then follows for the same reason that Lemma 3.20 did for While.

With this result, it is now possible to invoke Theorem 2.8 to compose RISC-level perthread security and mode compliance, whether they were obtained via the wr-compiler (using Theorem 5.24 and Lemma 5.25, respectively), or proved directly at RISC level.

We remark that, for programs wholly compiled by the wr-compiler, Theorem 5.22 can be subsumed by a whole-system preservation result that no longer demands local-mode-compliance for each thread, due to our ability to obtain it directly at RISC level:

Theorem 5.27 (Secure threads compiled by the wr-compiler form a secure system).

 $\forall i < \text{length } cms_r. \exists c \ l \ nl \ PCs \ l' \ nl' \ C' \ regs.$ $\text{com-secure}_{no-locks-held}^{no-locks-held} \ (c, \ mds_0) \land$ $(PCs, l', nl', C', \ False) = \text{compile-cmd } C_0 \ l \ nl \ c \land \ \text{compile-cmd-input-reqs } C_0 \ l \ nl \ c \land$ $cms_r[i] = (((0, \ map \ fst \ PCs), \ regs), \ mds_0)$

sys-secure_{no-locks-held} cms_r

Proof By Theorem 2.8 and unfolding Definition 2.10, we are required to prove security and local mode compliance for every thread of the compiled RISC program, and global modes compatibility between them all as a whole, assuming no-locks-held and using mds_0 initially. These requirements are immediate using Theorem 5.24, Lemma 5.25, and Lemma 5.26.

6 Case study: Cross Domain Desktop Compositor input handler

This section presents—as the main case study for the COVERN wr-compiler—a mixedsensitivity concurrent program whose source-level noninterference properties are preserved by verified secure compilation down to an assembly-level model.

The Cross Domain Desktop Compositor (CDDC) of Beaumont *et al.* (2016) is a desktop device that gives trusted users the option of replacing multiple monitors, keyboards, and mice with a single multi-level secure user interface (via a single monitor, keyboard, and mouse, as depicted in Figure 5a) when using several desktop computers simultaneously.

Here we present as case study a program (replacing customised hardware) that handles the incoming mouse and keyboard inputs to the CDDC. This program has served as a particularly good case study, because it features both of the characteristics for which proving information-flow security is this work's main focus:

- Concurrency—here, between *software components* whose execution is interleaved (by the seL4 operating-system microkernel (Klein *et al.*, 2014)), and that interact via shared memory.
- Mixed-sensitivity reuse—here, of system resources (notably the input devices) and memory locations, for input whose sensitivity level can be different at different times.

By exercising the COVERN Wr-COMPiler on a While model of this case study, we show this compiler verification-based approach to be feasible for obtaining the preservation of noninterference properties proved at While level, straightforwardly and for little extra effort, down to a RISC model of the program.

The section will proceed as follows. Following an overview in Section 6.1 of the main characteristics of the case study, Section 6.2 presents the formal security properties proved about its While model—as our focus is its compilation, further details on this model and the proof techniques used to prove these properties at While level are left to Sison (2020).





(a) CDDC hardware use-case setup.

The bar painted at the top of the screen indicates the computer set to receive all keyboard events. Mouse events are delivered to the owner of the topmost window underneath the mouse cursor.

(b) CDDC hardware architecture.

The HID switch—implemented in software on top of seL4—runs on an ARM Cortex A9 core, and operates a compositor device implemented (as in Beaumont *et al.* (2016)) using an FPGA.

Fig. 5: Functional schematics for Cross Domain Desktop Compositor hardware. Reproduced from Murray *et al.* (2018).

Section 6.3 then presents the formal preservation of security properties down to a RISC model, obtained from running the verified wr-compiler of Section 5 on the While model.

6.1 Overview of the case study

The case study is a software implementation of the *human interface device (HID) switch* in the CDDC (see Figure 5b). In short, this part of the CDDC is responsible for determining the destination of all *HID input (keyboard* and *mouse* device) events, and ensuring that the user remains informed of that destination (by operating a *video compositor* device, which renders display elements for that purpose on a shared monitor, as depicted in Figure 5a).

6.1.1 Information-flow security

The HID switch's responsibilities are security critical, as the CDDC is intended to provide an interface to multiple desktop computers belonging to different security domains; hence, the user of the CDDC is expected to choose the sensitivity of the data they input, based on the computer to which they expect it to be delivered. Furthermore, part of the CDDC's functionality is to allow users to choose which computer they are interacting with, by clicking on (accordingly responsive) display elements using the mouse. Thus, the desired information-flow security property for the HID switch is that, in providing this functionality, it never delivers inputs to a destination contrary to the user's expectations.

We simplify analysis to the classic High $\not\rightarrow$ Low security policy over the basic two-point {High, Low} security lattice, and model the HID switch to service only two potential destination computers.⁶ One computer is designated as belonging to the High security domain,

⁶ Aside from presenting a more minimal case study, any verification for an arbitrary security lattice can be reduced to multiple applications of verification to the basic High → Low policy, with the locations reclassified appropriately. Furthermore, the design of the CDDC's HID switch program is symmetrical for each user.



Fig. 6: Functional schematic of seL4 component architecture for CDDC HID switch. Reproduced from Murray *et al.* (2018).

and is the only legitimate destination for High-sensitivity input events; the other is designated as belonging to the Low security domain. The hardware and connections that the SWITCH component uses to forward events to these computers are modelled as shared variables classified statically: one High, the other Low (as depicted in Figure 7b). The attacker is then considered to be an entity that can read at any time from the Low-classified one.

6.1.2 Shared-variable concurrency

```
lock(hid_read_atomicity_lock);
temp := hid_keyboard_available
                                                     if (current_event_type = KEYBOARD) then
unlock(hid_read_atomicity_lock);
                                                       if (active_domain = DOM_LOW) then
if (temp != 0) then
                                                         output_event_buffer0 := current_event_data
 lock(input_event_lock);
                                                       else
  input_event_data := 0;
                                                         output_event_buffer1 := current_event_data
  input_event_type := KEYBOARD;
                                                       fi
 input_event_data := hid_keyboard_source;
                                                     else
  unlock(input_event_lock)
                                                       skip
else
                                                     fi
  skip
fi
```

(a) Receipt from input device by INPUT driver.
The hid_keyboard_source variable is valuedependently classified by the value of its sole control variable, indicated_domain (modelling trusted user input to the keyboard). (b) Delivery to output device by SWITCH. The output-event buffers 0 and 1 are statically classified Low and High respectively (modelling an attacker-controlled computer that receives all data written to buffer 0).

Fig. 7: Examples of external device interactions by the CDDC HID switch, as modelled in While—here, for the keyboard events. The full model is in the Isabelle/HOL supplement.

2252

2225

2226

2227 2228

2229

2230

2231

2232

2233

2235 2236 2237

2238

2239

2240

2241

2242

2243

- 2253
- 2254

```
compositor_cursor_position :=
                                                             if (switch_state_mouse_down = 0 &&
2255
                                                                 current_event_data = MOUSE_DOWN &&

    current event data:

                                                                 active_domain != cursor_domain) then
2256
         lock (compositor read atomicity lock):
                                                               active domain := cursor domain:
          cursor_domain :=
                                                               lock(input event lock):
2257

→ compositor_domain_under_cursor;

                                                               input_event_data := 0;
                                                               input_event_type := NONE;
          unlock (compositor_read_atomicity_lock);
2258
                                                               hid keyboard source := 0:
          if (cursor_domain = DOM_INVALID) then
                                                               indicated domain := active domain:
2259
           cursor_domain := active_domain
                                                               unlock(input_event_lock)
2260
          else
                                                             else
           skin
                                                               skip
2261
          fi
                                                             fi
2262
         (a) Querying the compositor to determine the
                                                             (b) Instructing the compositor to indicate a
2263
         topmost domain under the mouse cursor.
                                                             change to the active domain.
2264
             Fig. 8: Excerpts of the SWITCH component interfacing with the compositor device.
2265
2266
```

```
2267
                                                                 lock(input_event_lock);
          /* Permanently grab this lock */
                                                                 if (indicated_domain = active_domain)
2268
          lock(switch_private_lock);
          current_event_data := 0;
                                                                 then
2269
                                                                   current_event_type := input_event_type;
          current_event_type := NONE;
                                                                   current_event_data := input_event_data
2270
                                                                 else
          lock(input event lock);
                                                                  skin
          input_event_data := 0;
2271
                                                                 fi·
          hid_keyboard_source := 0;
                                                                 unlock(input_event_lock)
          indicated_domain := active_domain;
2272
          unlock(input_event_lock)
2273
```

 (a) Initialising private variables, input-event buffer, and compositor-indicated domain, to an arbitrary initial value for active_domain.
 Zeroing the data fields prevents leaking any High-sensitivity data they might initially contain. (b) Copying from the input-event buffer to private variables. The security analysis shows that repeating the previous event is a safe course of action when the environment misbehaves by violating indicated_domain = active_domain.

Fig. 9: Excerpts of the SWITCH component interacting with the input-event buffer.

The software implementation (replacing the original FPGA-based implementation (Beaumont *et al.*, 2016)) of the CDDC's HID switch is a system of software components written in C, that all run in user mode on top of the seL4 microkernel (Klein *et al.*, 2014). Here, we have abstracted from the seL4-based C implementation's details, to model in the While language the basic functionality of its three main software components (as depicted in Figure 6) as a shared-variable concurrent program of three threads:

- The INPUT driver is responsible for taking events from input-device interfaces and placing them on an input-event buffer for consumption by the SWITCH (Figure 7a).
- The SWITCH is responsible for inspecting all input events on the buffer from the INPUT driver, querying the compositor device (as modelled in Figure 8a) and OVERLAY driver to determine if any constitute a user-directed change to the destination of subsequent events, and if so, updating the compositor device to display that change (as modelled in Figure 8b). Finally, it is responsible for delivering all events to their destination computer via the appropriate *output-device* interface (Figure 7b).
- The OVERLAY driver is responsible for servicing remote procedure calls (RPCs, made by the SWITCH) that query a subset of the compositor-device interface, regarding the position of certain mouse-clickable elements the compositor is rendering as part of a visual overlay on the trusted user's video monitor. (As no mixed-sensitivity

2300

2274

2275

2276

2277

2278

2279 2280 2281

2282 2283

2284 2285

2286 2287

2288

2289

2290

2291

2292

2293

reuse occurs in this part of the model, we leave its details to the Isabelle/HOL supplement.)

The device interfaces, shared buffers (for input events and RPC mechanisms), and local variables used by each component are all modelled as program variables in shared memory. Consequently in the While model, mutex locks are used to model all synchronisation and restriction of concurrent access by the components to those variables.

So that we do not need to add separate While semantics for interacting with private as opposed to shared memory, we model thread-private memory as shared program variables protected by a permanently held lock acquired at initialisation time (e.g. as in Figure 9a). We consider this to be a stand-in for the memory isolation properties established by the underlying operating system between the program threads that it hosts.

6.1.3 Mixed-sensitivity reuse

Inherently to the CDDC's role as a multi-level secure user interface, its HID switch receives
 data of differing sensitivity levels (at different times) from a single set of input device
 memory locations (e.g. as modelled in Figure 7a, for the keyboard events), rather than
 from those of distinct device sets for each sensitivity level.

Furthermore, the HID switch propagates all input event data (regardless of sensitivity) through a single set of memory locations (the input-event buffer and SWITCH-internal copies of its contents, as modelled in Figure 9), rather than duplicating those memory locations for each security domain. Consequently in the While model, all of these memory locations that are subject to mixed-sensitivity reuse are assigned value-dependent classifications, reflecting the trusted user's expectation of the sensitivity level of the data they contain:

• To model a user that we trust to type sensitive information into the keyboard only when the compositor device indicates the High domain computer is *active* (i.e. set to receive all keyboard events), we have the INPUT driver draw keyboard events from a shared variable named hid_keyboard_source (as depicted in Figure 7a) that has classification dependent on a control variable indicated_domain modelling the relevant state of the compositor (here, DOM_HIGH is a designated constant):

 $\begin{cases} \mathsf{High}, & \text{if indicated_domain} = \mathsf{DOM_HIGH} \\ \mathsf{Low}, & \text{otherwise.} \end{cases}$

- In contrast, as clicking on composited user interface elements has the potential ability to change the future indicated_domain (which, as a control variable, is never allowed to receive any High-sensitivity data), the model trusts the user not to encode sensitive information into the mouse input in any way. Thus, the INPUT driver always draws mouse events from a statically Low-classified shared variable.
- Consequently, as the data portion input_event_data of the input-event buffer⁷ between the INPUT driver and SWITCH may carry either keyboard data of valuedependent sensitivity or Low-sensitivity mouse data, we assign it a classification
 - ⁷ We model in While only a single-place buffer, which could easily be extended to a buffer of arbitrary size by duplicating the same basic pattern of access, classification, and lock-protection, for multiple places.
- 2345 2346

2343

2344

2301

2302

2312

2313

2325

2326

2327

2328

2329

2330

2331

2332 2333 2334

2335

2336

2337

2338

52 Verified Secure Compilation for Mixed-Sensitivity Concurrent Programs

dependent on the values of both its control portion and the indicated_domain:

- $\begin{cases} \mathsf{High}, & \text{if input_event_type} = \texttt{KEYBOARD} \land & \texttt{indicated_domain} = \texttt{DOM_HIGH} \\ \mathsf{Low}, & \text{otherwise.} \end{cases}$
- Finally, we model the seL4-based SWITCH component's copying of the event from the buffer into its own local variables, giving its data portion a classification dependent on its own private view of the currently active domain (modelled as active_domain):
 - $\begin{cases} \mathsf{High}, & \text{if current_event_type} = \texttt{KEYBOARD} \land & \texttt{active_domain} = \texttt{DOM_HIGH} \\ \mathsf{Low}, & \text{otherwise.} \end{cases}$

To ensure active_domain remains authoritative with what is composited by the CDDC into the display, in the While model the SWITCH initialises indicated_domain to match the initial value of active_domain (as depicted in Figure 9a), updates it whenever active_domain changes (as depicted in Figure 8b), and checks at runtime that active_domain = indicated_domain when copying data from the buffer to its own private variables (as depicted in Figure 9b).

The CVDNI properties' (1) value dependence on control variables, (2) quantification over all initial values for the control variables and (3) assumptions of environmental havoc on write-unprotected shared variables between evaluation steps (Definition 2.6) then ensure noninterference between High inputs and Low-classified sinks, regardless of the initial and dynamically changing sensitivity of all such locations subject to mixed-sensitivity reuse.

6.2 CVDNI properties of the While model to be preserved

This section will now give a brief formal exposition of the security properties of the CDDC HID switch's While-language model that our compiler will preserve down to RISC.

As the per-thread proof techniques for While that we used for the case study are outside the scope of this paper, we note only that they consist of an adaptation to mutex locks by Sison (2020) of a security type system and local mode compliance check developed by Murray *et al.* (2016*b*,c). Nevertheless, we have provided their full formalisation in our Isabelle/HOL supplement, and we refer the reader to these prior works on their design, and particularly to Sison (2020) for further discussion on their application to this case study.

In short, from applying local type checks on the While-language commands for each of the three software components (INPUT, SWITCH, and OVERLAY) to obtain per-thread security (com-secure, Definition 2.7) and modes compliance (local-mode-compliance, Definition 2.11), we have from Theorem 2.8 that the concurrent program of all three components satisfies the whole-system security property (sys-secure, Definition 2.9) as instantiated to specify that no locks are held initially.

So that we can use the approach we gave in Section 3 to obtain the global modes compatibility part of the sound-mode-use side-condition (Definition 2.10), we specify no-locks-held (Definition 3.18) as the *INIT* requirement on memory, and use the initial mode state mds_0 (Definition 3.19) for all of the components in the system.

This no-locks-held predicate and mds₀ are both defined relative to a lock interpretation parameter that we supply (as required by Section 3.1) for the CDDC model. The locks in the CDDC model fall under the following categories:

- The locks coordinating inter-component interactions grant exclusive read-write access to the shared variables they govern.
 - There are also locks granting the SWITCH and INPUT components exclusive readwrite access to a set of "private" variables each, for internal use. The components acquire these prior to entering their main loop, and never release them.
- 2401 • Finally the model uses read-atomicity locks—a practice introduced in Section 5.1.1. These grant exclusive write access to shared variables used to model hardware 2402 interfaces, to make explicit an assumption (normally implicit in the atomicity of 2403 expression evaluation in the While language) that these variables will not have their 2404 value changed by the environment during a simple assignment from those variables. 2405 Note that these read-atomicity locks are not needed to prove confidentiality for 2406 the While model, but rather we add them to satisfy the requirements demanded 2407 2408 by the wr-compiler so that it can preserve confidentiality (via small-step semantic preservation) down to the RISC model. 2409

The While-language proof techniques we apply to each thread of the program yield com-secure^{no-high-branching}, a stronger version of the per-thread CVDNI property that enforces no-high-branching (Definition 5.5). Furthermore, we have trivially from the definition of com-secure (Definition 2.7) that if a program is secure without imposing any initial conditions, then it remains secure if we impose any *INIT* parameter arbitrarily. Therefore, for each thread we have com-secure^{no-high-branching}_{no-locks-held} (Definition 2.7, with *INIT* \triangleq no-locks-held and *EXTRA* \triangleq no-high-branching):

Lemmas 6.1 (Per-thread confidentiality results for CDDC While model).

com-secure ^{no-high-branching}	$(OVERLAY,mds_0)$
com-secure ^{no-high-branching}	(Input, mds_0)
com-secure ^{no-high-branching}	$(Switch,mds_0)$

From this and Theorem 2.8, using local compliance checks and Lemma 3.20 to discharge the sound-mode-use (Definition 2.10) side condition, we have a whole-system confidentiality theorem for the system of all three components running concurrently:

Theorem 6.2 (Whole-system confidentiality result for the CDDC While model).

sys-secure_{no-locks-held} [(OVERLAY, mds₀), (INPUT, mds₀), (SWITCH, mds₀)]

6.3 Confidentiality-preserving compilation to RISC model

We now turn to applying the COVERN wr-compiler of Section 5 to our While-language model of the CDDC's HID switch; we then have automatically that it preserves the security properties presented in Section 6.2 down to the compiler's RISC-language output.

2437 2438

2396 2397

2398 2399

2400

2418

2425

2426

2427 2428 2429

2430

2431 2432 2433

2434

2435

The wr-compiler is *executable* in the Isabelle proof assistant. Using Isabelle's eval tactic, we execute the wr-compiler's main function, compile-cmd (whose implementation was described in Section 5.1) on the While-language models for all three of the CDDC's INPUT driver, SWITCH, and OVERLAY driver components, to obtain their RISC-language compilations. (Recall from Section 5.1 that we obtain the RISC text trivially as the map fst of the *CompRec*-annotated RISC program, which is the fst output of compile-cmd.)

Definition 6.3 (RISC-language program texts of CDDC model's components).

Our approach to obtain per-thread confidentiality for each of these RISC texts will be to use the theorem of its preservation by the wr-compiler (Theorem 5.24). Recall, this was:

Theorem 5.24 (Preservation of per-thread confidentiality by the wr-compiler).

com-secure _{no-locks-held}	$^{g}\left(c,mds_{0} ight)$
$(PCs, l', nl', C', False) = compile-cmd C_0 l nl c$	compile-cmd-input-reqs C ₀ <i>l nl c</i>
com-secure $_{no-locks-held}^{pc-security}$ (((0, map fs	$t PCs), regs), mds_0)$

Then, for compile-cmd to execute successfully (i.e. to return *failed* = False), the model must pass the stability-checks discussed in Section 5.1. All three of OVERLAY, INPUT, and SWITCH pass the checks (1) because they use locks to protect the atomicity of reads from (otherwise unstable) variables used to model hardware interfaces, and (2) as a consequence of having passed the local security and mode compliance checks mentioned in Section 6.2.

We are now in a position to prove a whole-system confidentiality result for the compiled RISC model—here, with each thread's register bank initialised to zero: $\operatorname{regs}_0 \triangleq (\lambda_-, 0)$.

Theorem 6.4 (Whole-system confidentiality result for the CDDC RISC model).

$$\begin{split} \text{sys-secure}_{\text{no-locks-held}} & [(((0, \text{OVERLAY}_{\text{RISC}}), \text{regs}_0), \text{mds}_0), \\ & (((0, \text{INPUT}_{\text{RISC}}), \text{regs}_0), \text{mds}_0), \\ & (((0, \text{SWITCH}_{\text{RISC}}), \text{regs}_0), \text{mds}_0)] \end{split}$$

Proof A few approaches are available; we obtained formal proofs of this theorem in Isabelle/HOL using all three of the following alternatives (unfolding Definition 6.3):

Option 1. Use either of Theorem 5.22 or Theorem 5.27, both of which established whole-system security for RISC outputs of the wr-compiler when executed on com-secure^{no-high-branching} While programs (which we have here from Lemmas 6.1). This is the easiest option to take for programs that are already verified in the While language, and then compiled successfully to RISC by the wr-compiler. It is possible to take here because all of OVERLAY_{RISC}, INPUT_{RISC}, and SWITCH_{RISC} were obtained in this manner.

Option 2. Use Theorem 5.24 and Lemma 5.25 to obtain com-secure^{pc-security}_{no-locks-held} and local-mode-compliance (resp.) for each of OVERLAY_{RISC}, INPUT_{RISC}, and SWITCH_{RISC}, then use Theorem 2.8 directly to obtain sys-secure_{no-locks-held}. This option can be used for systems where some of the threads are written directly in RISC; for such threads, com-secure^{pc-security}_{no-locks-held} and local-mode-compliance would need to be proved directly at RISC level. However, Lemma 5.26 still discharges the global-modes-compatibility requirement for RISC, provided all threads are initialised with mds₀ and no-locks-held.

Option 3. Use Theorem 2.23 directly. This option can be used for systems where all the RISC threads are secure refinements (according to Definition 2.15) of the threads of some While program that satisfied sound-mode-use with no-locks-held initially, but some were obtained by other means than the wr-compiler (i.e. not all via the refinement \mathcal{R}_{wr}).

7 Related work

First in Section 7.1, we describe other recent and related works that address concerns of noninterference proof compositionality in a concurrent setting (of the kind we tackled in Section 3). The remaining sections focus on related works on verified compilation: The works in Section 7.2 and Section 7.3, like ours, focus on compilation preserving a form of noninterference. In Section 7.4 we describe our work's relationship with varieties of robust property preservation, and other compilation verification efforts in Section 7.5.

7.1 Compositionality of concurrent noninterference proofs

Alternative approaches exist to establishing the non-compositional global modes compatibility condition we proved as invariant to concurrent While executions in Section 3. For the precursor (non-value-dependent) notion of concurrent noninterference to CVDNI, Mantel *et al.* (2011) originally proposed that such a condition be met by a non-compositional *may happen in parallel* analysis (e.g. Masticola & Ryder (1993)). Then, instead of demanding the explicit declaration of the sorts of guarantees implied by locking discipline (as we do), Mantel *et al.* (2015) proposed automating their inference and proof of the compatibility condition using a reachability analysis making use of dynamic pushdown networks. We leave adapting and implementing such an approach for our CVDNI setting to future work.

We note also that, like the CVDNI theory and our work of Section 3, recent work by 2518 Frumin et al. (2021) concerns compositionality of machine-checked proof efforts for non-2519 interference in a concurrent setting that are obtained potentially via a variety of proof 2520 techniques. They model more fine-grained synchronisation than we do here, via atomic 2521 compare-and-swap operations that can be used to implement mutex locking primitives. 2522 However, they do not study compilation as a means of preserving such proofs, which is the 2523 focus of our work here. We believe that the CVDNI refinement notions we presented could 2524 support certain cases of compilation between different synchronisation primitives, pro-2525 vided only new thread-private state is needed (like the registers in RISC), and the shared 2526 variable interactions can be proved as preserved. For example, we expect mutex locking 2527 primitives (with slightly different semantics to ours here) could feasibly be refined to a 2528 compare-and-swap-based implementation in this way-this we also leave to future work. 2529

2530

2485

2486

2487

2488

2489

2490

2491

2492

2493

2494

2495 2496 2497

2498 2499

2500

2501

2502

2503

2504

2505 2506 2507

2508

2509

2510

2511

2512

2513

2514

2515

2516

2531

7.2 Noninterference-preserving compilation

Tedesco *et al.* (2016) present a type-directed compilation scheme that preserves a *fault*-2532 resilient noninterference property. The compilation scheme of our wr-compiler was 2533 inspired by theirs. Like our com-secure CVDNI security property that wr-compiler pre-2534 serves, Tedesco et al.'s security property is also strong bisimulation-based (Sabelfeld & 2535 Sands, 2000). But where our property accounts (via mode states) for controlled interference 2536 by other threads, theirs instead quantifies over all possible interference by the environment 2537 with the memory contents. While this simplifies their task of proving that their security 2538 property is preserved under compilation-as it need not require the compiler to preserve 2539 the contents of memory—it means their security property cannot capture value-dependent 2540 noninterference. In contrast, our wr-compiler must obey our secure-refinement notion's 2541 requirement that memory contents are preserved. 2542

The line of work most relevant to ours is that which was conducted (concurrently) 2543 by Barthe et al. (2020), wherein they achieved the remarkable result of proving that 2544 a modification of the CompCert C compiler (Leroy, 2009) preserves the cryptographic 2545 constant-time class of noninterference (2-safety) properties. Their proof approach was to 2546 use various notions of *constant-time simulation* (CT-simulation) first presented by Barthe 2547 et al. (2018), originally intended for application to the Jasmin compiler (Almeida et al., 2548 2017). Although not targeting programs with concurrency or mixed-sensitivity reuse (as 2549 our work does), CT-simulation shares in common with the refinement notions used by this 2550 paper that it in essence rests on a simulation diagram that is cube-shaped, as it must pre-2551 serve a 2-safety hyperproperty. We submit that Barthe et al. (2020) broadly validates the 2552 argument we made in Sison & Murray (2019), that decomposing such cube-shaped dia-2553 grams into square-shaped ones is what will make them feasible to apply to the verification 2554 of fully fledged compilers like CompCert-noting that they described the only compila-2555 tion pass they proved with their non-decomposed, cube-shaped diagram as "not especially 2556 pleasant because the diagrams are difficult to exploit" (Barthe et al., 2020). 2557

Note that the refinement theory of Barthe et al. (2018, 2020) preserves security via 2558 refinement phrased in terms of forward simulation (Leroy, 2009)-that is, each step of the 2559 abstract program must be simulated by the target program. In contrast, our theory presented 2560 here is instead targeted towards preserving refinement via backward simulations,⁸ in which 2561 each step of the concrete (compiled) program must be simulated by the abstract program. 2562 This difference arises because in our setting we need to account for leakage that might 2563 occur and be visible only in intermediate states. In their setting, in contrast, leakage that 2564 occurs in intermediate states remains visible forever in the concrete program semantics via 2565 a *leakage trace*. It remains unclear whether we could have adopted a similar approach in 2566 our work, thereby enabling a (simpler) forward simulation argument. In particular, it is 2567 not clear what the semantics of leakage traces should be for a language that supports both 2568 value-dependent classification and shared-memory concurrency as ours does. 2569

- 2570
- 2571 2572

2573

2574

⁸ Again, as commonly referred to in the compiler verification literature from Leroy (2009) onwards. This is not to be confused with the "backward simulations" of concurrency verification (Lynch & Vaandrager, 1996) and data refinement (de Roever & Engelhardt, 1998; Cavalcanti & Naumann, 2002), where the refined program instead simulates the original, and where simulation proceeds from the end of the program back to the beginning.

7.3 Concurrency-compositional noninterference-preserving compilation

Neither of the above consider per-thread compositional compilation of concurrent, shared 2578 memory programs, nor value-dependent noninterference policies – the focus of our theory 2579 and compiler. Barthe et al. (2010, 2007a) however did aim to preserve noninterference of 2580 multithreaded programs by compilation, extending a prior (security) type-preserving com-2581 pilation approach (Barthe et al., 2004, 2007b). Their noninterference property however 2582 was termination- and timing-insensitive, so preventing internal timing leaks relied on the 2583 scheduler disallowing certain interleavings between threads. Also, their type-preservation 2584 argument was derived from a big-step semantics preservation property for their compiler. 2585 Here we instead rely on preservation of a small-step semantics (specifically memory con-2586 tents), which is necessary for us to preserve value-dependent security under compilation. 2587 as well as to avoid imposing non-standard requirements on the scheduler. 2588

7.4 Robust property preservation

2592 Other recent works have improved on *fully abstract compilation* (surveyed by Patrignani et al. (2019)) by mapping out the spectrum (Abate et al., 2019) or developing specific 2593 2594 forms (Patrignani & Garg, 2019) of robust property preservation, concerned with robustness of source program (hyper)properties to concrete adversarial contexts. Like Tedesco 2595 2596 et al. (2016), these works differ from ours in quantifying over a wider range of hostile interference. They also focus prominently on changes to data types, which we do not support. 2597 Thus, as a 2-safety hyperproperty quantifying over a lesser range of interference, we expect 2598 CVDNI-preservation to be implied by R2HSP (robust 2-hypersafety preservation), but do 2599 not expect it to imply any other secure compilation criterion on Abate et al.'s spectrum. 2600

While recently Patrignani & Garg (2019) instantiated their robustly safe compilation 2601 for shared-memory fork-join concurrent programs, it only preserves (1-)safety properties. 2602 2603 Previously however, Patrignani & Garg (2017) proved their trace-preserving compilation preserves k-safety hyperproperties (Clarkson & Schneider, 2010), including noninterfer-2604 2605 ence properties. However, it disallows the removal or addition of trace entries, which would be necessary to change the passage of time as seen in the observable trace events. Thus it 2606 2607 excludes the sorts of changes to pacing carried out by our compiler (regulated by *abs-steps*) and studied as optimisations by the two other works (Tedesco et al., 2016; Barthe et al., 2608 2020) on timing-sensitive security-preserving compilation mentioned above. 2609

7.5 Compiler verification in general

2613 Finally, there has been much work on large-scale verified compilation (Leroy, 2009; 2614 Kumar et al., 2014) some of which has also treated compilation of shared-memory con-2615 current programs (Lochbihler, 2018) including taking weak-memory consistency into 2616 account (Podkopaev et al., 2019). Our work here does not consider the effects of 2617 weak-memory models. In particular, such models are often defined axiomatically rather 2618 than operationally. Our notion of secure refinement and our decomposition principle 2619 (Definitions 2.15 and 2.18, respectively) are defined assuming an operational semantics 2620 for the source and target languages.

2621

2610 2611

2612

2577

2589 2590

Our work differs to prior work on verified concurrent compilation, in that it formalises and proves a compiler's ability to use information about the application's locking protocol, both to exclude unsafe access to shared variables, and conversely to know when it is safe to allow optimisations on shared variables that would typically be excluded.

8 Conclusion

To our knowledge, we have presented the first mechanised verification that a compiler preserves concurrent, value-dependent noninterference. To this end, we provided a general decomposition principle for compositional, secure refinement. Although our compiler is a proof-of-concept targeting simple source and target languages, we nevertheless applied it to produce a verified assembly-level model of an input-handling system for the CDDC (Beaumont *et al.*, 2016), a nontrivial mixed-sensitivity concurrent program.

We expect this decomposition principle to remain applicable in reducing noninterference-refinement proof efforts for compilers that overcome the specific limitations of ours here. For example, a compiler that inserts padding to equalise the time taken on either side of a High-branch—which may change when it expands expressions into multiple instructions—may instantiate the decomposition principle with a more sophisticated concrete coupling invariant that does not require pc-security.

This work serves to demonstrate that verified security-preserving compilation for mixedsensitivity concurrent programs is now within reach, by augmenting traditional proof obligations for verified compilation (e.g. square-shaped semantics preservation) with those specific to security (e.g. absence of termination- and timing-leaks) as depicted in Figure 4. We hope that this work paves the way for future large-scale verified security-preserving compilation efforts.

Acknowledgements

We would like to thank our anonymous referees, and to thank again all those who provided feedback on the conference version of this paper (Sison & Murray, 2019) and on Robert Sison's PhD thesis (Sison, 2020). This paper describes research that was conducted during Robert's PhD candidature at UNSW Sydney and CSIRO's Data61, which was funded by an Australian Government Research Training Program (RTP) Scholarship and a CSIRO Data61 Research Project Award. We thank the Trustworthy Systems group at CSIRO's Data61 for cultivating an excellent working and learning environment.

Conflicts of Interest

None

Abate, C., Blanco, R., Garg, D., Hritcu, C., Patrignani, M. and Thibault, J. (2019) Journey beyond full abstraction: Exploring robust property preservation for secure compilation. *32nd IEEE Computer*

References

Security Foundations Symposium, CSF 2019, Hoboken, NJ, USA, June 25-28, 2019 pp. 256–271. IEEE.

- Almeida, J. B., Barbosa, M., Barthe, G., Blot, A., Grégoire, B., Laporte, V., Oliveira, T., Pacheco,
 H., Schmidt, B. and Strub, P.-Y. (2017) Jasmin: High-assurance and high-speed cryptography.
 Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security.
 CCS '17, pp. 1807–1823. ACM.
- Barthe, G., Basu, A. and Rezk, T. (2004) Security types preserving compilation: (extended abstract).
 Steffen, B. and Levi, G. (eds), *Verification, Model Checking, and Abstract Interpretation, 5th International Conference, VMCAI 2004, Venice, Italy, January 11-13, 2004, Proceedings*. Lecture
 Notes in Computer Science 2937, pp. 2–15. Springer.
- Barthe, G., Rezk, T., Russo, A. and Sabelfeld, A. (2007a) Security of multithreaded programs by compilation. Biskup, J. and López, J. (eds), *Computer Security - ESORICS 2007, 12th European Symposium On Research In Computer Security, Dresden, Germany, September 24-26, 2007, Proceedings.* Lecture Notes in Computer Science 4734, pp. 2–18. Springer.
- Barthe, G., Rezk, T. and Basu, A. (2007b) Security types preserving compilation. *Comput. Lang. Syst. Struct.* 33(2):35–59.
- Barthe, G., Rezk, T., Russo, A. and Sabelfeld, A. (2010) Security of multithreaded programs by
 compilation. ACM Trans. Inf. Syst. Secur. 13(3):21:1–21:32.
- Barthe, G., Grégoire, B. and Laporte, V. (2018) Secure compilation of side-channel countermea sures: The case of cryptographic "constant-time". *31st IEEE Computer Security Foundations Symposium, CSF 2018, Oxford, United Kingdom, July 9-12, 2018* pp. 328–343. IEEE Computer
 Society.
- Barthe, G., Blazy, S., Grégoire, B., Hutin, R., Laporte, V., Pichardie, D. and Trieu, A. (2020) Formal
 verification of a constant-time preserving C compiler. *Proc. ACM Program. Lang.* 4(POPL):7:1–7:30.
- Beaumont, M., McCarthy, J. and Murray, T. (2016) The cross domain desktop compositor: Using hardware-based video compositing for a multi-level secure user interface. Schwab, S., Robertson, W. K. and Balzarotti, D. (eds), *Proceedings of the 32nd Annual Conference on Computer Security Applications, ACSAC 2016, Los Angeles, CA, USA, December 5-9, 2016* pp. 533–545. ACM.
- Cavalcanti, A. and Naumann, D. A. (2002) Forward simulation for data refinement of classes.
 Eriksson, L.-H. and Lindsay, P. A. (eds), *FME 2002:Formal Methods—Getting IT Right* pp. 471–490. Springer Berlin Heidelberg.
- Clarkson, M. R. and Schneider, F. B. (2010) Hyperproperties. J. Comput. Secur. 18(6):1157–1210.
- de Roever, W. P. and Engelhardt, K. (1998) *Data Refinement: Model-oriented Proof Theories and their Comparison*. Cambridge Tracts in Theoretical Computer Science, vol. 46. Cambridge
 University Press.
- Focardi, R., Gorrieri, R. and Panini, V. (1995) The security checker: a semantics-based tool for the verification of security properties. *Proceedings The Eighth IEEE Computer Security Foundations Workshop* pp. 60–69.
- Frumin, D., Krebbers, R. and Birkedal, L. (2021) Compositional non-interference for fine-grained concurrent programs. *42nd IEEE Symposium on Security and Privacy (S&P'21), to appear; CoRR* abs/1910.00905.
- Jones, C. B. (1981) Development Methods for Computer Programs including a Notion of Interference. D.Phil. thesis, University of Oxford.
- Kaufmann, T., Pelletier, H., Vaudenay, S. and Villegas, K. (2016) When constant-time source yields variable-time binary: Exploiting curve25519-donna built with msvc 2015. *Cryptology and Network Security* pp. 573–582. Springer International Publishing.
- Klein, G., Andronick, J., Elphinstone, K., Murray, T., Sewell, T., Kolanski, R. and Heiser, G.
 (2014) Comprehensive formal verification of an OS microkernel. *ACM Transactions on Computer Systems* 32(1):2:1–2:70.
- Kumar, R., Myreen, M., Norrish, M. and Owens, S. (2014) CakeML: A verified implementation of ML. Peter Sewell (ed), ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages pp. 179–191. ACM Press.

2714

Leroy, X. (2009) A formally verified compiler back-end. J. Autom. Reason. 43(4):363-446.

- Lochbihler, A. (2018) Mechanising a type-safe model of multithreaded java with a verified compiler.
 Journal of Automated Reasoning 61(1):243–332.
- 2717 Lynch, N. and Vaandrager, F. (1996) Forward and backward simulations. *Inf. Comput.* 128(1):1–25.
- Mantel, H., Sands, D. and Sudbrock, H. (2011) Assumptions and guarantees for compositional noninterference. *IEEE Computer Security Foundations Symposium* pp. 218–232. IEEE.
- Mantel, H., Müller-Olm, M., Perner, M. and Wenner, A. (2015) Using dynamic pushdown networks to automate a modular information-flow analysis. 25th International Symposium on Logic Based Program Synthesis and Transformation (LOPSTR).
- Masticola, S. P. and Ryder, B. G. (1993) Non-concurrency analysis. *Proceedings of the Fourth ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*. PPOPP '93, pp. 129–138. ACM.
- Molnar, D., Piotrowski, M., Schultz, D. and Wagner, D. (2006) The program counter security model: Automatic detection and removal of control-flow side channel attacks. *Proceedings of the* 8th International Conference on Information Security and Cryptology. ICISC'05, pp. 156–168.
 Springer-Verlag.
- Murray, T. (2015) On high-assurance information-flow-secure programming languages. ACM
 SIGPLAN Workshop on Programming Languages and Analysis for Security pp. 43–48.
- Murray, T., Sison, R., Pierzchalski, E. and Rizkallah, C. (2016a) Compositional security-preserving refinement for concurrent imperative programs. Archive of Formal Proofs June. http:// isa-afp.org/entries/Dependent_SIFUM_Refinement.shtml, Formal proof development.
 Murray, T., Sison, R., Pierzchalski, E. and Rizkallah, C. (2016b) Compositional security-preserving refinement for concurrent imperative programs. Archive of Formal Proofs June. http:// isa-afp.org/entries/Dependent_SIFUM_Refinement.shtml, Formal proof development.
- Murray, T., Sison, R., Pierzchalski, E. and Rizkallah, C. (2016b) Compositional verification and refinement of concurrent value-dependent noninterference. *IEEE 29th Computer Security Foundations Symposium, CSF 2016, Lisbon, Portugal, June 27 - July 1, 2016* pp. 417–431. IEEE Computer Society.
- Murray, T., Sison, R., Pierzchalski, E. and Rizkallah, C. (2016c) A dependent security type system for concurrent imperative programs. *Archive of Formal Proofs* June. http://isa-afp.org/
 entries/Dependent_SIFUM_Type_Systems.html, Formal proof development.
- Murray, T., Sison, R. and Engelhardt, K. (2018) COVERN: A logic for compositional verification of
 information flow control. *European Symposium on Security and Privacy* pp. 16–30. IEEE.
- Patrignani, M. and Garg, D. (2017) Secure Compilation and Hyperproperty Preservation. *IEEE* 30th Computer Security Foundations Symposium, CSF 2017, Santa Barbara, USA, August 21 -25, 2017. CSF'17.
- Patrignani, M. and Garg, D. (2019) Robustly safe compilation. *Programming Languages and Systems* pp. 469–498. Springer International Publishing.
- Patrignani, M., Ahmed, A. and Clarke, D. (2019) Formal approaches to secure compilation: A survey of fully abstract compilation and related work. *ACM Comput. Surv.* 51(6):125:1–125:36.
- Podkopaev, A., Lahav, O. and Vafeiadis, V. (2019) Bridging the gap between programming languages and hardware weak memory models. *Proc. ACM Program. Lang.* **3**(POPL):69:1–69:31.
- Sabelfeld, A. and Sands, D. (2000) Probabilistic noninterference for multi-threaded programs.
 Proceedings of the 13th IEEE Workshop on Computer Security Foundations. CSFW '00, pp. 200–. IEEE Computer Society.
- Sison, R. (2020) Proving Confidentiality and Its Preservation Under Compilation for Mixed Sensitivity Concurrent Programs. PhD thesis, University of New South Wales, Sydney. http:
 //doi.org/10.26190/5fab5c0a76454.
- Sison, R. and Murray, T. (2019) Verifying That a Compiler Preserves Concurrent Value-Dependent
 Information-Flow Security. Harrison, J., O'Leary, J. and Tolmach, A. (eds), *10th International Conference on Interactive Theorem Proving (ITP 2019)*. Leibniz International Proceedings in
 Informatics (LIPIcs) 141, pp. 27:1–27:19. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik.
- Staples, M., Jeffery, R., Andronick, J., Murray, T., Klein, G. and Kolanski, R. (2014) Productivity
 for proof engineering. *Empirical Software Engineering and Measurement* p. 15.
- 2758
- 2759
- 2760

	Tedesco, F. D., Sands, D. and Russo, A. (2016) Fault-resilient non-interference. IEEE 29th Computer
2761	Security Foundations Symposium, CSF 2016, Lisbon, Portugal, June 27 - July 1, 2016 pp. 401-
2762	416. IEEE Computer Society.
2763	Siveroni I (eds) Static Analysis pp. 352–367. Springer Berlin Heidelberg
2764	Volpano, D. and Smith, G. (1998) Probabilistic noninterference in a concurrent language.
2765	Proceedings. 11th IEEE Computer Security Foundations Workshop (Cat. No.98TB100238) pp.
2766	34–43.
2767	
2760	
2709	
2771	
2772	
2773	
2774	
2775	
2776	
2777	
2778	
2779	
2780	
2781	
2782	
2783	
2784	
2785	
2786	
2787	
2788	
2789	
2790	
2791	
2792	
2793	
2794	
2795	
2796	
2797	
2798	
2800	
2800	
2802	
2803	
2804	
2805	
2806	