

Verifying That a Compiler Preserves Concurrent Value-Dependent Information-Flow Security

Robert Sison 

Data61, CSIRO, Australia
UNSW Sydney, Australia
Robert.Sison@data61.csiro.au

Toby Murray

University of Melbourne, Australia
toby.murray@unimelb.edu.au

Abstract

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

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

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

2012 ACM Subject Classification Security and privacy → Logic and verification; Security and privacy → Information flow control; Software and its engineering → Compilers

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

Digital Object Identifier 10.4230/LIPIcs.ITP.2019.27

Related Version A full version of the paper is available at <https://arxiv.org/abs/1907.00713>.

Supplement Material The Isabelle/HOL theories are available at <https://covern.org/itp19.html>.

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

Acknowledgements We would like to thank our anonymous reviewers, as well as Carroll Morgan, Kai Engelhardt, Gerwin Klein, Christine Rizkallah, Matthew Brecknell, Johannes Åman Pohjola, and Qian Ge, for their very helpful feedback on earlier versions of this paper.

1 Introduction

It is well known that program translations of the kind carried out by compilers can in principle break security properties like confidentiality [12, 2]. Yet source level reasoning about confidentiality remains common [20, 19, 18]. Existing verified compilers like CompCert



© Robert Sison and Toby Murray;
licensed under Creative Commons License CC-BY


10th International Conference on Interactive Theorem Proving (ITP 2019).

Editors: John Harrison, John O’Leary, and Andrew Tolmach; Article No. 27; pp. 27:1–27:19

Leibniz International Proceedings in Informatics



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



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

[15] and CakeML [14] preserve semantics, but semantics preservation alone may be insufficient to preserve confidentiality, especially for shared memory concurrent programs whose threads must guard against timing leaks in order to prevent them manifesting as storage leaks [22].

Supporting secure compilation of programs that must enforce *value-dependent* security policies poses an additional challenge, because in such policies the sensitivity of a memory location can depend on the values held in other memory locations. Thus, unlike prior work on secure compilation [4], preserving security under refinement requires a refinement relation that is strong enough to preserve those memory contents on which the policy depends.

In prior work [22], we presented a definition for a notion of value-dependent security-preserving refinement that is compositional for concurrent programs: by applying it to each thread individually, one can derive a secure refinement of the concurrent composition.

The essence of this notion of security-preserving refinement (presented fully in Subsection 2.2) is in its refinement preservation obligation (**coupling-inv-pres** in Figure 1). Here, the usual square-shaped commuting diagram that is commonly used to depict (semantics-preserving) refinement (Figure 4a) has been replaced by a *cube* (Figure 1). The additional dimension of this cube reflects that it preserves a 2-safety hyperproperty [6] that compares two executions rather than examining a single one. As such, it is significantly more complicated to prove than standard notions of semantics-preserving refinement typical in verified compilation [15, 14].

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

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

In Section 3 we present our decomposition principle, which decomposes the cube (Figure 1) into three separate obligations (Figure 4). The first of these is akin to semantics-preserving refinement, while the second and third essentially ensure together that the refinement has not introduced any termination- and timing-leaks.

In Section 4 we show how the decomposition principle can almost halve the effort to prove secure refinement – in this case, of a program that is especially prone to introduced timing leaks because it branches on secrets (a feature not yet allowed by our compiler). There, we present a side-by-side comparison of the proof effort, both with and without the decomposition principle. We find that using it reduces the proof’s complexity by 44%.

In Section 5, we present our compiler and its formal verification, as an application of the decomposition principle. This compiler translates concurrent programs written in an imperative While language, with locking primitives for mediating access to shared memory, into a RISC-style assembly language. It does so by compiling each thread individually, and in doing so preserves a formal security property that remains compositional between threads. Furthermore, our compiler demonstrates a way of formalising and proving when it is safe for a compiler to perform optimisations in the presence of concurrency. To ensure that the contents of shared memory locations are preserved under compilation despite potential interference from other threads, our compiler tracks which shared memory locations are *stable* (free from any such interference). It then makes use of this tracking to avoid redundant loads from stable shared variables safely, that would otherwise be considered unsafe to omit.

All results are mechanised in Isabelle/HOL,¹ and in Section 6 we explain how, in order to validate our theory, we instantiated it so that we could execute our compiler in Isabelle. This enabled us to execute it over a While language model of the Cross Domain Desktop Compositor [5] (CDDC), a concurrent program that enforces information flow control over value-dependently classified input. To our knowledge this is the first proof of information flow security for an assembly-level model of a non-trivial concurrent program, demonstrating the power of verified secure compilation for deriving security properties of compiled code.

2 Background and example

We begin by introducing with an illustrative example (Figure 2) the challenges of verifying *value-dependent information-flow security* in the presence of *shared-variable concurrency*.

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 2b), to locations belonging to the other, which we classify Low (Figure 2c).

Here and generally, 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, this 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, and we aim to verify that as a whole it satisfies the security requirement. But to avoid a state space explosion that is exponential in the number of threads, we must do this *compositionally*: one thread at a time, then combining the results of these analyses.

We focus on a particular worker thread (Figure 2a), 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 $domain = \text{LOW}$ (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 $domain = \text{LOW}$ exactly when the phone indicates it is

¹ The *wr-compiler* totals ~7k lines, and verification + compilation of the 2-thread CDDC model totals ~1.6k lines of Isabelle proof script, excluding whitespace and comments. See “Supplement Material”.

27:4 Verifying That a Compiler Preserves Concurrent Value-Dependent Inflow Security

```


while TRUE do
  lock(workspace_lock);
  while !suspended do
    lock(source_lock);
    workspace := source;
    /* ... operations on workspace ... */
    if domain = LOW then
      low_sink := workspace
    else
      high_sink := workspace;
      workspace := 0
    fi;
    unlock(source_lock)
  od;
  unlock(workspace_lock);
  while suspended do skip od
od

```

(a) Input processing worker thread program.



(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.

■ **Figure 2** Example: Touchscreen input processing for a dual-personality smartphone.

providing the Low personality (Figure 2c), 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 – e.g. processing input from a fingerprint scanner – in such a way that it can assume $workspace$ no longer contains any sensitive values.

When we analyse one thread like this worker in terms of our compositional security property (Subsection 2.1), all of the other threads in the system are trusted to do two things:

1. They follow a *synchronisation scheme*: here, if read- or write-access to a certain variable is governed by a lock, they must hold it in order to access the variable in that manner.
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 we are proving 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.

Firstly, it is important that no thread in the system (including the thread 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 $domain$.

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

2.1 Concurrent value-dependent noninterference (CVDNI)

Having illustrated the challenges with an example, we now focus on the formalisation of our information-flow security property CVDNI, which we target with our per-thread analysis, and which our compiler preserves. It is defined in terms of two main elements:

1. a binary *strong low-bisimulation (modulo modes)* relation \mathcal{B} between program configurations, that establishes the required information-flow security property. Like Goguen & Meseguer-style noninterference [10], any states it relates must agree on their “low” portions, and it demands that lock-step execution preserve that correspondence. This section will explain how it is specialised further for shared-variable concurrency.
2. a *classification* function \mathcal{L} that determines the “low” portion of a program configuration, thus affecting \mathcal{B} ’s requirements. Unlike [10] however, \mathcal{L} here can depend on values in the program configuration itself, thus expressing dynamic and not just static classifications.

We now present definitions from Section III-2b of our previous work [22] simplified as noted. The theory is parameterised over the type of values Val , a finite set of shared variables Var , and a *deterministic evaluation step semantics* \rightsquigarrow between *local configurations* (of a thread in a concurrent program) each denoted by a triple $\langle tps, mds, mem \rangle$:

- tps is the *thread-private state*, which is permanently inaccessible to the attacker and the other threads. Note that due to this inaccessibility, we allow the user of the theory to parameterise the type of tps , and do not impose any particular structure.
- $mds :: Mode \Rightarrow Var \text{ set}$ is the (*access*) *mode state*, which is ghost state associating each $Mode = \{\mathbf{AsmNoW}, \mathbf{AsmNoRW}, \mathbf{GuarNoW}, \mathbf{GuarNoRW}\}$ with a set of shared variables. Intuitively, it identifies the set of variables for which the thread currently possesses (or respects) a kind of exclusivity of access granted (or obligated) by a synchronisation scheme. This facilitates compositional, assume-guarantee [11] style reasoning. For example, when our worker thread holds *source_lock*, it *assumes no other threads write to source* or its control variable ($\{source, domain\} \subseteq mds \mathbf{AsmNoW}$), otherwise it *guarantees it does not write* to them ($\mathbf{GuarNoW}$). Similarly, holding *workspace_lock* it assumes no other threads *read or write* to *workspace* ($workspace \in mds \mathbf{AsmNoRW}$), and at all other times it makes the corresponding guarantee ($\mathbf{GuarNoRW}$).
- $mem :: Mem$ is *shared memory* considered potentially accessible to the attacker and other threads. In order to make what is accessible amenable to analysis, we impose the structure $Mem = Var \Rightarrow Val$, a total map from shared variable names to their values.

The theory is then further parameterised by the value-dependent classification function $\mathcal{L} :: Mem \Rightarrow Var \Rightarrow \{\mathbf{High}, \mathbf{Low}\}$, and a function $Cvars :: Var \Rightarrow Var \text{ set}$ that returns all the control variables of a given variable. In our worker thread example, $\mathcal{L} mem x$ gives:

- **High** when x is *high_sink*, meaning *high_sink* is classified **High** at all times.
- when x is *source*: **Low** if $mem \ domain = \mathbf{LOW}$, and **High** otherwise.
- **Low** for all other variables x , meaning they are classified **Low** at all times.

The set $\mathcal{C} = \{y \mid \exists x. y \in Cvars x\}$ is then defined to contain all control variables in the system. Thus in our worker thread example, $Cvars \ source = \{domain\}$ and $\mathcal{C} = \{domain\}$.

To support compositionality for concurrent programs, the “low” portion demanded to be equal by the analysis is tightened up to be *modulo modes* – it includes non-control variables only if they are assumed to be *readable* by other threads according to the mode state: *readable* $mds x \equiv x \notin mds \mathbf{AsmNoRW}$. Thus intuitively, the user of the theory should model permanent untrusted output sinks of the whole concurrent program, as variables for which \mathcal{L} *always returns Low*, ungoverned by any synchronisation scheme that the attacker cannot be trusted to follow. (In our example, *low_sink* is untrusted permanently in this way, but *workspace* is untrusted only when unlocked.) The notion of observational indistinguishability used for the noninterference property is then defined over memories as follows.

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

$$\begin{aligned} mem_1 &=_{m_{ds}}^{\text{Low}} mem_2 \equiv \\ \forall x. x \in \mathcal{C} \vee \mathcal{L} mem_1 x = \text{Low} \wedge \text{readable } m_{ds} x &\longrightarrow mem_1 x = mem_2 x \end{aligned}$$

For this paper, we will use notation $lc_1 =_{m_{ds}}^{\text{Low}} lc_2$ to lift $=_{m_{ds}}^{\text{Low}}$ to local program configurations, asserting also that lc_1 and lc_2 are *modes-equal* (have the same mode state). Additionally, we will use notation $lc_1 =_{m_{ds}} lc_2$ to denote (alone) that lc_1 and lc_2 are modes-equal.

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

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

$$\begin{aligned} \text{com-secure } (tps, m_{ds}) &\equiv \forall mem_1 mem_2. mem_1 =_{m_{ds}}^{\text{Low}} mem_2 \longrightarrow \\ &(\exists \mathcal{B}. \text{strong-low-bisim-mm } \mathcal{B} \wedge (\langle tps, m_{ds}, mem_1 \rangle, \langle tps, m_{ds}, mem_2 \rangle) \in \mathcal{B}) \end{aligned}$$

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

$$\begin{aligned} \text{strong-low-bisim-mm } \mathcal{B} &\equiv \text{cg-consistent } \mathcal{B} \wedge \text{sym } \mathcal{B} \wedge \\ (\forall lc_1 lc_2. (lc_1, lc_2) \in \mathcal{B} \wedge lc_1 =_{m_{ds}} lc_2 &\longrightarrow lc_1 =_{m_{ds}}^{\text{Low}} lc_2 \wedge \\ (\forall lc'_1. lc_1 \rightsquigarrow lc'_1 \longrightarrow (\exists lc'_2. lc_2 \rightsquigarrow lc'_2 \wedge lc'_1 =_{m_{ds}} lc'_2 \wedge (lc'_1, lc'_2) \in \mathcal{B}))) & \end{aligned}$$

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

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

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

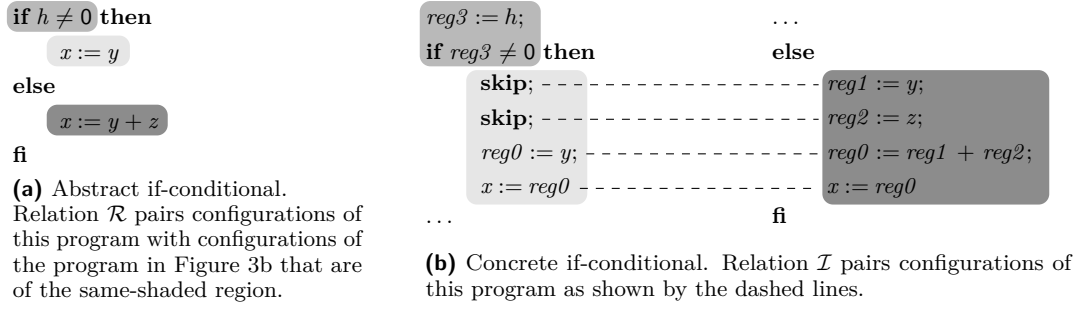
► **Definition 3** (Closedness under globally consistent changes).

$$\begin{aligned} \text{cg-consistent } \mathcal{B} &\equiv \forall tps_1 mem_1 tps_2 mem_2 m_{ds}. \\ (\langle tps_1, m_{ds}, mem_1 \rangle, \langle tps_2, m_{ds}, mem_2 \rangle) \in \mathcal{B} &\longrightarrow \\ (\forall mem'_1 mem'_2. (\forall x. (mem_1 x \neq mem'_1 x \vee mem_2 x \neq mem'_2 x \vee \\ \mathcal{L} mem_1 x \neq \mathcal{L} mem'_1 x) \longrightarrow \text{writable } m_{ds} x) \wedge mem'_1 =_{m_{ds}}^{\text{Low}} mem'_2 &\longrightarrow \\ (\langle tps_1, m_{ds}, mem'_1 \rangle, \langle tps_2, m_{ds}, mem'_2 \rangle) \in \mathcal{B}) & \end{aligned}$$

Theorem 3.1 of our prior work [22] then gives us that the parallel composition of **com-secure** programs is itself a program that enforces a system-wide value-dependent noninterference property (**sys-secure**, for whose details we refer the reader to Section III-2(a) of [22]).

2.2 CVDNI-preserving refinement

Having described the formal security property that we wish to be preserved under refinement (and compilation), we now define formally a suitable notion of secure refinement that preserves it. The proof of CVDNI-preserving refinement for a thread of a concurrent program relies on two binary relations (illustrated by Figure 3) to be nominated by the user of the theory:



■ **Figure 3** Excerpts from refinement example [22] that was used to compare proof effort (Section 4).

1. a *refinement relation* \mathcal{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 efforts.
2. a *concrete coupling invariant* \mathcal{I} that allows us to use \mathcal{B} and \mathcal{R} to build a new strong low-bisimulation (modulo modes) for the concrete program, by discarding unreachable pairs of local configurations *after the refinement*. It thereby witnesses that any changes a refinement (or compiler) makes to 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 – be imposed on the nominated \mathcal{R} and \mathcal{I} in relation to a given witness relation \mathcal{B} establishing CVDNI for the abstract program. The definitions to follow are adapted from Murray et al. [22] Section V. For better readability, we present a simplified version in which no new shared variables are added by the refinement. Consequently we introduce the notation $\stackrel{\text{mem}}{\text{mds}}$ to denote that two local configurations have equal mode state and memory, regardless of whether relating configurations of the same or differing languages.

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

► **Definition 4** (Preservation of modes and memory).

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

Regarding the closedness under changes by other threads that ensures compositionality for concurrency, on \mathcal{I} we again impose **cg-consistent** (Definition 3) from Subsection 2.1. However in the case of \mathcal{R} , we instead impose **closed-others**, a simplification of **cg-consistent** considering only environmental actions that affect the memories on both sides of the relation identically. Furthermore it ensures equality of *all* shared variables, not just those judged observable:

► **Definition 5** (Closedness of refinements under changes by others).

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

The final major requirement for CVDNI-preservation is then to prove \mathcal{R} and \mathcal{I} closed simultaneously under the pairwise executions of the concrete and abstract programs, using the aforementioned cube-shaped diagram (*coupling-inv-pres*, Figure 1) whose edges are pairs in \mathcal{B} , \mathcal{R} , and \mathcal{I} . All that then remains is for the nominated concrete coupling invariant \mathcal{I} to be symmetric, and the predicate *secure-refinement* puts together all the requirements:

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

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

Theorem 5.1 of our prior work [22] gives us that under the aforementioned conditions,

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

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

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

3 Decomposition principle for CVDNI-preserving refinement

Having presented our previous work [22]’s formalisation of our security property CVDNI and its preservation by refinement, we now present our first contribution: an alternative way of proving *secure-refinement* (Definition 6) that does away with the use of the cube-shaped, two-sided refinement obligation *coupling-inv-pres* $\mathcal{B} \mathcal{R} \mathcal{I}$ (depicted by Figure 1), by decomposing its concerns into (1) proving \mathcal{R} closed under the pairwise executions of the concrete and abstract programs alone using a square-shaped diagram (depicted by Figure 4a, which is akin to ordinary semantics-preserving refinement), and (2) a number of smaller and more separable obligations gathered together under the side-condition predicate *decomp-refinement-safe*.

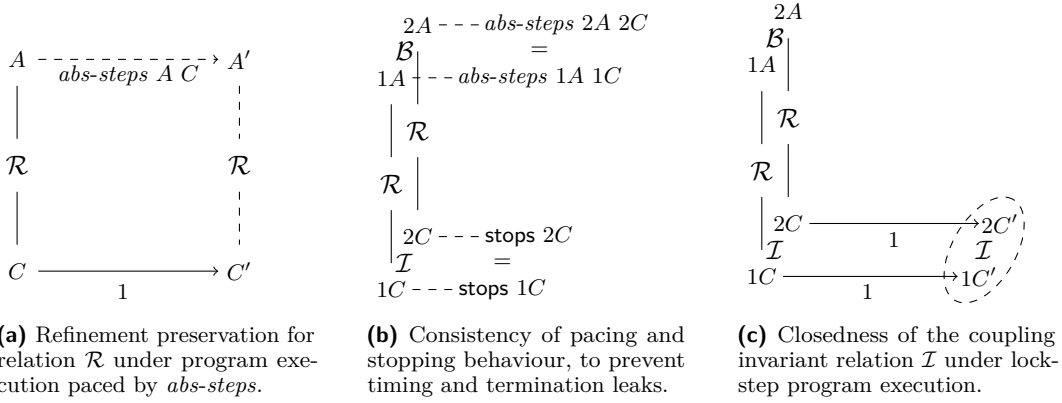
► **Definition 7** (Decomposed requirements for CVDNI-preserving secure refinement).

$$\text{secure-refinement-decomp } \mathcal{B} \mathcal{R} \mathcal{I} \text{ abs-steps} \equiv \\ \text{preserves-modes-mem } \mathcal{R} \wedge \text{closed-others } \mathcal{R} \wedge \text{cg-consistent } \mathcal{I} \wedge \text{sym } \mathcal{I} \wedge \\ \text{decomp-refinement-safe } \mathcal{B} \mathcal{R} \mathcal{I} \text{ abs-steps} \wedge (\forall lc_A lc_C. (lc_A, lc_C) \in \mathcal{R} \longrightarrow \\ (\forall lc'_C. lc_C \rightsquigarrow_C lc'_C \longrightarrow (\exists lc'_A. lc_A \rightsquigarrow_A^{(\text{abs-steps } lc_A lc_C)} lc'_A \wedge (lc'_A, lc'_C) \in \mathcal{R})))$$

The decomposition requires the provision of a new refinement parameter that we will call *abs-steps* or the *pacings function*, whose role is to dictate the pace of the refinement by returning the number of abstract steps that ought to be taken for a single concrete step, for a given abstract-concrete local configuration pair related by \mathcal{R} . The side-conditions on all of the refinement parameters (depicted by Figures 4b, 4c) are then defined as follows:

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

$$\text{decomp-refinement-safe } \mathcal{B} \mathcal{R} \mathcal{I} \text{ abs-steps} \equiv \forall lc_{1A} lc_{2A} lc_{1C} lc_{2C}. (lc_{1A}, lc_{2A}) \in \mathcal{B} \wedge \\ lc_{1A} \stackrel{\text{mds}}{=} lc_{2A} \wedge (lc_{1A}, lc_{1C}) \in \mathcal{R} \wedge (lc_{2A}, lc_{2C}) \in \mathcal{R} \wedge (lc_{1C}, lc_{2C}) \in \mathcal{I} \wedge lc_{1C} \stackrel{\text{mds}}{=} lc_{2C} \\ \longrightarrow \text{stops } lc_{1C} = \text{stops } lc_{2C} \wedge \text{abs-steps } lc_{1A} lc_{1C} = \text{abs-steps } lc_{2A} lc_{2C} \wedge \\ (\forall lc'_{1C} lc'_{2C}. lc_{1C} \rightsquigarrow_C lc'_{1C} \wedge lc_{2C} \rightsquigarrow_C lc'_{2C} \longrightarrow (lc'_{1C}, lc'_{2C}) \in \mathcal{I} \wedge lc'_{1C} \stackrel{\text{mds}}{=} lc'_{2C})$$



■ **Figure 4** Graphical depictions of refinement decomposition obligations.

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

- **stops** $lc_{1C} = \text{stops } lc_{2C}$ ensures that the refinement has not introduced any termination leaks, by asserting *consistent stopping behaviour* for \mathcal{I} -related concrete program configurations, which we know to be observationally indistinguishable.
- **abs-steps** $lc_{1A} lc_{1C} = \text{abs-steps } lc_{2A} lc_{2C}$ ensures that the refinement has not introduced any timing leaks, by asserting *consistency of the pace of the refinement* for \mathcal{R} -related program configurations, which we again know to be observationally indistinguishable.
- The final \forall -quantified clause asserts \mathcal{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 the \mathcal{B} - and \mathcal{R} -edges in Figure 4c may capture useful facts about a particular program verification technique and compiler, so their availability as assumptions is intended to reduce greatly the effort needed to specify a coupling invariant \mathcal{I} and prove it satisfies the condition.

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

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

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

In the interests of brevity we relegate proof sketches for all results to the extended version of the paper, and for fuller details we refer the reader to our Isabelle/HOL formalisation.

We now devote our attention to two instantiations of this new decomposition principle: (Section 4) for a proof of CVDNI-preservation for the refinement of a program that branches on a secret, and (Subsection 5.5) for the proof of CVDNI-preservation by a compiler.

4 Proof effort comparison

To demonstrate how the decomposition principle reduces proof complexity and effort, we returned to the example refinement discussed in Section V-E of our previous work [22], an excerpt of which is shown in Figure 3. 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 linearly correlated with the latter [28]. Formalised in Isabelle/HOL as `EgHighBranchRevC.thy` [21], the proof line count for that theory stood at about 4.6K lines of definitions and proof, of which approx. 3.6K line were proofs. Adapting the proof instead to use the decomposition principle `secure-refinement-decomp` (Definition 7), the proof line count drops from 3.6K to approx. 2K, a 44% reduction. Regarding definition changes, the new proof makes <10 lines of adaptations to a coupling invariant and pacing function used by the old proof, and adds about 30 lines worth of new helper definitions, for use with the decomposition principle. The rest of the theory and its external dependencies remain in common between the two versions.

As would be expected, the bulk of the deletions are from the full cube-shaped refinement diagram proof (Figure 1) of `secure-refinement` (Definition 6) for the refinement relation. The surviving parts of that proof just become the square-shaped refinement diagram proof (Figure 4a) of `secure-refinement-decomp` without much modification. The deletions are replaced by newly added proofs of the three sub-obligations of `decomp-refinement-safe` (Definition 8).

5 The COVERN `wr-compiler`

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

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

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

5.1 Source language

The COVERN `wr-compiler` – short for *While-to-RISC compiler* – takes the simple imperative language with while-looping and lock-based synchronisation targeted by the COVERN program logic [20], which we will refer to as `While`, consisting of the commands `cmd`:

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

The language is parameterised over a type of values Val , and binary operators $\oplus :: Val \Rightarrow Val \Rightarrow Val$. Constants $n :: Val$; $v :: Var$ and $k :: Lock$ are (resp.) shared program- and lock-variables. The semantics of the locking primitives `lock(k)` and `unlock(k)` is informed by a locking discipline provided by the user of the theory as a parameter (see Subsection 5.3). We leave for future work adding support for pointers and arrays, which we believe will be straightforward because our assume-guarantee framework already provides the means to encode the memory footprint of a command in a way that depends on values in memory.

We assume that the underlying concurrent execution model (e.g. operating system, scheduler) for the `While` language prevents threads from seeing each others' current program location, and thus (as in previous work [22, 19]) the `While` program command $c :: \text{cmd}$ being executed we model as thread-private state: $\langle c, \text{mds}, \text{mem} \rangle_w$. In contrast, all program variables $v :: Var$ and lock variables $k :: Lock$ reside in the shared memory mem .

5.2 Target language

The `wr`-compiler's target is a generic RISC-style assembly language like that of Tedesco et al. [29] but with lock-based synchronisation primitives added, which we will refer to as RISC:

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

The language is parameterised over the same value type Val and binary operators \oplus , shared program variables $v :: Var$ and shared lock variables $k :: Lock$ as the `While` language. Presently, direct-addressing `Load` and `Store` instructions (referring to registers $r :: Reg$) are adequate for RISC to implement all existing `While` features, and we expect adding indirect addressing to RISC to be as straightforward as adding pointer and array support to `While`.

RISC program texts P are just lists of binary instructions I , each optionally associated with a label $l :: Lab$. We assume that the underlying concurrency model for the RISC language (e.g. OS, scheduler etc.) prevents one thread from reading the program code (instructions) of another,² as well as another's registers (including the program counter). Thus, we model the distinguished program counter register's value $pc :: nat$, program text P , and register bank $regs :: Reg \Rightarrow Val$ as thread-private state: $\langle (pc, P), regs, mds, mem \rangle_r$. Apart from this adaptation to our triple format, evaluation semantics follows that of the RISC target of [29].

Finally, like Tedesco et al. [29] we generalise over the (user-supplied) register allocation scheme, and assume there are enough registers to service the maximum depth of expressions in the source program. We leave for future work the modelling and analysis of a compiler phase that spills register contents to memory, in order to make this assumption unnecessary.

5.3 Locking discipline

Like the `COVERN` logic [20], we assume that the `While` language program being compiled follows a certain locking discipline, about which the compiler has knowledge, so as to ensure that the RISC program it produces follows the same discipline.

The user of the theory provides the details of the locking discipline in the form of a *lock interpretation* parameter: $lock\text{-}interp :: Lock \Rightarrow (Var \ set \times Var \ set)$, which for each lock gives the two non-overlapping sets of program variables over which acquiring the lock grants exclusive permission to write, (resp.) read and write. These permissions are then reflected in the way the semantics of the `While` and RISC locking primitives act on the mode state.

Regarding lock interpretations and the way they interact with the user-provided value-dependent classification function \mathcal{L} (see Subsection 2.1), we inherit a few cleanliness conditions from that earlier work [20], chief of which are that lock variables k cannot be control variables, a lock variable k governing access to a program variable v must govern the same kind of access to all of v 's control variables, and \mathcal{L} must classify all lock variables as `Low`.

² As is usual for program analyses, we omit any explicit modelling of the microarchitectural state used by superscalar processors (like CPU caches, and state relied on by speculative and out-of-order execution, on whose behaviour attacks like Spectre [13] and Meltdown [16] relied). We argue however that our present assumptions are reasonable under two circumstances: when there is no such state (e.g. on microcontrollers like AVR [7]), or when such state is correctly *partitioned* by the underlying hardware [30] or the OS [8] – if the hardware allows it [9]! In the latter case, our analysis assumes that microarchitectural state footprints are partitioned according to thread (for memory containing program text) and according to classification by \mathcal{L} (for shared memory), and furthermore that each value-dependently classified region is given a distinct partition that is flushed on reclassification.

5.4 Compiler implementation and tracking of shared variable stability

We chose as a starting point the compilation scheme of [29], on the basis of their preserving a noninterference property that like ours exhibits resilience to changes made by an environment – in their case, intended for fault-resilience. Aiming to repurpose that for shared-variable concurrency, we adapted it to Isabelle, implementing it as a primitive recursive function:

$$\begin{aligned} \text{compile-cmd} &:: \text{CompRec} \Rightarrow \text{Lab option} \Rightarrow \text{Lab} \Rightarrow \text{cmd} \Rightarrow \\ &(\text{I} \times \text{CompRec}) \text{ list} \times \text{Lab option} \times \text{Lab} \times \text{CompRec} \times \text{bool} \end{aligned}$$

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

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

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

Here, `compile-cmd` takes an *initial compilation record* C , an optional *entry label* l , and the *next available label* nl , and for the benefit of the next invocation returns an optional *exit label* l' if one is used by the program just compiled, the *new next available label* nl' , and a *final compilation record* C' . We leave details of label allocation and its impact on achieving sequential composability for compiled RISC programs to the extended version of the paper.

In addition to the output RISC program $P :: \text{I list}$ itself, a call to `compile-cmd` also outputs every CompRec associated with the state of the program just before executing every instruction in P . These are returned zipped up together with P as the *CompRec-annotated RISC program* $\text{PCs} :: (\text{I} \times \text{CompRec}) \text{ list}$. (P can trivially be recovered as `map fst PCs`.) Finally, `compile-cmd` may return `True` for *failed* to reject the input program, such as when it detects a data race (see below), or if expression depth exceeds the assumed limit (Subsection 5.2).

In the style of the compilation scheme on which it was based [29], the `wr-compiler` maintains a *register record* $\Phi :: \text{reg} \rightarrow \text{exp}$, i.e. a partial map of registers to expressions on shared variables. In addition to using it to compile away any unnecessary loads from variables in shared memory, we also use it to ensure that an expression calculated by RISC in registers is equal to the value of the expression as if it had all been calculated by `While` in one step. This is especially important when writing the result of an expression back to shared memory, because the refinement is required to maintain all shared memory values.

New to the `wr-compiler` is the responsibility of maintaining an *assumption record*, which it uses primarily to detect and reject programs with data races on shared memory, and to rule out the introduction of any new ones. Each assumption record $\mathcal{S} :: (\text{Var set} \times \text{Var set})$ is a pair tracking the set of variables on which (resp.) **AsmNoW**, **AsmNoRW** assumptions are currently active at a given point in the program being compiled. As a secondary concern we also use it to assert that the two sides of any if-conditional branches act consistently on the mode state, and that while-loops restore the original mode state on termination.

A compilation record $C = (\Phi, \mathcal{S}) :: \text{CompRec}$ is then just a register/assumption record pair. For readability, we use `regrec`, `asmrec` to denote (resp.) a CompRec 's `fst`, `snd` projections.

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

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

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

For register record entries to be of any help in ensuring consistency of **While** and RISC expression evaluation, we exclude expression evaluation on data 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.)

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

To ensure that an input **While** program maintains register record stability, we define the predicate `no-unstable-exprs` $c C$ to capture the requirement that a program c , if started with a configuration consistent with compilation record C , will never access a lock-protected variable without holding the relevant lock. (It also checks the secondary, mode-state consistency concerns of the assumption record mentioned earlier.) We implement it as a simple static check carried out by a primitive recursive function on the structure of **While** programs.

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

5.5 Proof of CVDNI-preserving compilation

Having covered the most significant aspects of the **COVERN** *wr*-compiler's parameters and machinery, we can now present the refinement relation \mathcal{R}_{wr} (Subsubsection 5.5.1), pacing function `abs-stepswr` (Subsubsection 5.5.2), and coupling invariant \mathcal{I}_{wr} (Subsubsection 5.5.3) that we use with our new decomposition principle (of Section 3) to prove that it preserves CVDNI (Subsubsection 5.5.4).

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

Just like our example \mathcal{R} of Figure 3, \mathcal{R}_{wr} pairs abstract with concrete configurations.

Here, we will focus on \mathcal{R}_{wr} 's most notable characteristics for understanding why it is suitable to describe a CVDNI-preserving compilation.³ We focus on the case `if_expr` of \mathcal{R}_{wr} , which relates the expression evaluation part of the **While** program `if e then c_1 else c_2 fi`, with the corresponding part (including the conditional jump **Jz** after expression evaluation) of the RISC program obtained by running `compile-cmd` on it. (Variables ignored are in gray.)

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

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

³ We provide an informal description of all of the cases, their purpose, and the invariants they maintain, along with a code listing from `compile-cmd` relevant to the part that will be presented, in the extended version of our paper. For full details, we refer the reader to the Isabelle formalisation.

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

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

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

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

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

As we will see in Theorem 17, `compiled-cmd-config-consistent` also serves 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 compilation record C .

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

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

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

5.5.2 Refinement pacing function `abs-stepswr`

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

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

Our pacing function `abs-stepswr` primarily looks at the form of the RISC program instruction about to be executed. The RISC instructions are divided into three categories:

- Instructions output by `compile-expr`: **Load**, **Op**, and **MoveK**. For these, `abs-stepswr` returns 1 if the `leftmost-cmd` 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-stepswr` 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-stepswr` 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-stepswr` returns 1.

Having nominated `abs-stepswr` and \mathcal{R}_{wr} , we now have the parameters over which we are obliged to prove refinement preservation (Figure 4a) as demanded by `secure-refinement-decomp` (Definition 7). To this end, we prove firstly (elided to the extended version) 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. Also, we must prove a correctness lemma for the expression compiler:

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

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

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

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

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

The next element needed is the concrete coupling invariant \mathcal{I}_{wr} , which we define as follows:

$$\mathcal{I}_{wr} \equiv \{ \langle \langle (pc, P), regs \rangle, mds, mem \rangle_r, \langle \langle (pc', P'), regs' \rangle, mds', mem' \rangle_r \mid (pc, P) = (pc', P') \}$$

In other words, \mathcal{I}_{wr} asserts that we only need compare local configurations that are at the same location $pc = pc'$ of the same RISC program $P = P'$. When used in concert with a `no-high-branching` \mathcal{B} (see Subsubsection 5.5.4), the effect of \mathcal{I}_{wr} is to ensure that the `wr`-compiler has not introduced any *new* branching on sensitive values.

5.5.4 Successful compilations are CVDNI-preserving refinements

We are ready to prove preservation. First we qualify that we allow only `strong-low-bisim-mm` \mathcal{B} that describe only `While`-programs with no branching on `High`-classified values, as follows:

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

$$\begin{aligned} \forall c c' mds mem mem'. (\langle c, mds, mem \rangle_w, \langle c', mds, mem' \rangle_w) \in \mathcal{B} &\longrightarrow c = c' \wedge \\ (\forall e c_1 c_2. \text{leftmost-cmd } c = \text{if } e \text{ then } c_1 \text{ else } c_2 \text{ fi} &\longrightarrow \text{ev}_{\text{exp}} mem e = \text{ev}_{\text{exp}} mem' e) \end{aligned}$$

That is, it refuses to relate configurations at different program locations. Furthermore if it is at a conditional branching point, the expression e determining which branch will be taken evaluates to the same boolean value for both configurations' memories. When imposed on a relation that already ensures `Low`-equivalent memory modulo modes, this effectively disallows any present or past branching on sensitive values. Then, for such programs:

► **Lemma 15.**
$$\frac{\text{strong-low-bisim-mm } \mathcal{B} \quad \text{no-high-branching } \mathcal{B}}{\text{secure-refinement-decomp } \mathcal{B} \mathcal{R}_{wr} \mathcal{I}_{wr} \text{ abs-steps}_{wr}}$$

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

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

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

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

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

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

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

To test the theory, we instantiated it and applied the `wr-compiler` to a `While`-language model of the Cross Domain Desktop Compositor [5] (CDDC), a non-trivial concurrent program that facilitates a trusted user’s interaction with multiple desktop machines of differing clearance.

The CDDC model to which we applied the compiler is a 2-thread program that was a precursor to the 3-thread model that was verified using the COVERN program logic [20].⁴ Each of the threads of the CDDC program (together about 150 lines of `While`) we proved satisfy the compositional security property `com-secure` (Definition 2), using a precursor to the COVERN logic that yields CVDNI-witness bisimulations that are non-High-branching.

The resulting compiler is *executable* in Isabelle, meaning that `compile-cmd` can be executed on the `While` program text for each of the two threads to obtain their compilations (together totalling about 250 RISC instructions) using the Isabelle tactic `eval`. The secure compilation theorems (Subsubsection 5.5.4), together with `strong-low-bisim-mm` preservation and compositionality for `com-secure` (Theorems 5.1, 3.1 of [22], mentioned in Section 2) then allow us to derive that the compiled program is secure when its threads are run concurrently.

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

7 Related work

The following three works, like ours, focus on compilation preserving a form of noninterference.

Tedesco et al. [29] present a type-directed compilation scheme that preserves a *fault-resilient* noninterference property. The compilation scheme of our `wr-compiler` was inspired by theirs. Like our `com-secure` CVDNI security property that `wr-compiler` preserves, Tedesco et al.’s security property is also *strong bisimulation*-based [27]. But where our property accounts (via mode states) for *controlled interference* by other threads, theirs instead quantifies over all possible interference by the environment with the memory contents. While this simplifies their task of proving that their security property is preserved under compilation – as it need not require the compiler to preserve the contents of memory – it means their security property cannot capture value-dependent noninterference. In contrast, our `wr-compiler` must obey our `secure-refinement` notion’s requirement that memory contents are preserved.⁵

Barthe et al. [2] consider the problem of preserving *cryptographic constant-time policies*, a class of noninterference properties similar to CVDNI in its explicit consideration for capturing timing-sensitivity. Barthe et al. consider a wider scope of common categories of compile-time optimisations (than those performed by our `wr-compiler`), and mechanise proofs in Coq that such optimisations preserve various constant-time security properties. The sharing of variables in our setting severely limits the scope of our optimisations, to those that the compiler can perform knowing that a shared variable is stable because it has been locked. At present, our `wr-compiler` avoids redundant loads during expression compilation, but other optimisations like loop hoisting and constant folding we are yet to implement. Their preservation proof technique, *constant-time simulation* was developed independently to our original cube-shaped

⁴ We leave for future work an adaptation of the refinement theory and `wr-compiler` in order to support the *shared data invariants* added by the COVERN logic, required to verify the 3-thread CDDC model.

⁵ Consequently, we found and fixed a bug in their expression compiler (acknowledged privately) whereby registers in use were incorrectly reallocated. Expressions like $v + (v + 1)$ were thus compiled incorrectly to programs yielding $(v + 1) + (v + 1)$ instead, causing a violation of memory contents preservation.

secure refinement definition [22]. Like ours, theirs is also a cube-shaped obligation and makes use of a pacing function analogous to our *abs-steps*. Unlike our work here, Barthe et al. do not give a general method for decomposing their cube-shaped simulation diagrams.

Neither of the above consider per-thread compositional compilation of concurrent, shared memory programs, nor value-dependent noninterference policies – the focus of our theory and compiler. Barthe et al. [4] however did aim to preserve noninterference of multithreaded programs by compilation, extending a prior (*security*) *type-preserving* compilation approach [3]. Their noninterference property however was termination- and timing-*insensitive*, so preventing internal timing leaks relied on the scheduler disallowing certain interleavings between threads. Also, their type-preservation argument was derived from a big-step semantics preservation property for their compiler. Here we instead rely on preservation of a small-step semantics (specifically memory contents), which is necessary for us to preserve value-dependent security under compilation, as well as to avoid imposing non-standard requirements on the scheduler.

Other recent works have improved on *fully abstract compilation* (surveyed [23]) by mapping out the spectrum [1] or developing specific forms [25] of *robust property preservation*, concerned with *robustness* of source program (hyper)properties to concrete *adversarial* contexts. Like Tedesco et al. [29], 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. Thus, as a 2-safety hyperproperty quantifying over a lesser range of interference, we expect CVDNI-preservation to be implied by R2HSP (robust 2-hypersafety preservation), but do not expect it to imply any other secure compilation criterion on Abate et al.’s [1] spectrum.

While recently Patrignani and Garg [25] instantiated their *robustly safe compilation* for shared-memory fork-join concurrent programs, it only preserves (1-)safety properties. Previously however, Patrignani et al. [24] proved their *trace-preserving compilation* preserves *k*-safety hyperproperties [6], including noninterference 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 excludes optimisations carried out by our compiler (when it permits changes to pacing regulated by *abs-steps*) and studied by the two other works [29, 2] on timing-sensitive security-preserving compilation mentioned above.

Finally, there has been much work on large-scale verified compilation [15, 14] some of which has also treated compilation of shared-memory concurrent programs [17] including taking weak-memory consistency into account [26]. Our work here does not consider the effects of weak-memory models. However, it 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, to exclude unsafe access to shared variables, and conversely to know when it is safe to allow optimisations that would typically be excluded (see Subsection 5.4).

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 the CDDC [5], a non-trivial concurrent program.

This work serves to demonstrate that verified security-preserving compilation for 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.

References

- 1 Carmine Abate, Roberto Blanco, Deepak Garg, Catalin Hritcu, Marco Patrignani, and Jérémy Thibault. Exploring Robust Property Preservation for Secure Compilation. *CoRR*, abs/1807.04603, 2018. [arXiv:1807.04603](https://arxiv.org/abs/1807.04603).
- 2 G. Barthe, B. Grégoire, and V. Laporte. Secure Compilation of Side-Channel Countermeasures: The Case of Cryptographic “Constant-Time”. In *2018 IEEE 31st Computer Security Foundations Symposium (CSF)*, pages 328–343, July 2018.
- 3 Gilles Barthe, Tamara Rezk, and Amitabh Basu. Security Types Preserving Compilation. *Comput. Lang. Syst. Struct.*, 33(2):35–59, July 2007. [doi:10.1016/j.cl.2005.05.002](https://doi.org/10.1016/j.cl.2005.05.002).
- 4 Gilles Barthe, Tamara Rezk, Alejandro Russo, and Andrei Sabelfeld. Security of Multithreaded Programs by Compilation. *ACM Trans. Inf. Syst. Secur.*, 13(3):21:1–21:32, July 2010. [doi:10.1145/1805974.1805977](https://doi.org/10.1145/1805974.1805977).
- 5 Mark Beaumont, Jim McCarthy, and Toby Murray. The Cross Domain Desktop Composer: Using Hardware-Based Video Compositing for a Multi-Level Secure User Interface. In *Annual Computer Security Applications Conference (ACSAC)*, pages 533–545, 2016.
- 6 Michael R. Clarkson and Fred B. Schneider. Hyperproperties. *J. Comput. Secur.*, 18(6):1157–1210, September 2010. URL: <http://dl.acm.org/citation.cfm?id=1891823.1891830>.
- 7 Florian Dewald, Heiko Mantel, and Alexandra Weber. AVR processors as a platform for language-based security. In *Computer Security - ESORICS 2017 - 22nd European Symposium on Research in Computer Security, Oslo, Norway, September 11-15, 2017, Proceedings, Part I*, pages 427–445, 2017. [doi:10.1007/978-3-319-66402-6_25](https://doi.org/10.1007/978-3-319-66402-6_25).
- 8 Qian Ge, Yuval Yarom, Tom Chothia, and Gernot Heiser. Time Protection: the Missing OS Abstraction. In *Eurosys19*, Dresden, Germany, March 2019. ACM.
- 9 Qian Ge, Yuval Yarom, and Gernot Heiser. No Security Without Time Protection: We Need a New Hardware-Software Contract. In *Asia-Pacific Workshop on Systems (APSys)*, Korea, August 2018. ACM SIGOPS.
- 10 Joseph Goguen and José Meseguer. Security Policies and Security Models. In *Proceedings of the IEEE Symposium on Security and Privacy*, pages 11–20, Oakland, California, USA, April 1982. IEEE Computer Society.
- 11 Cliff B. Jones. *Development Methods for Computer Programs including a Notion of Interference*. D.Phil. thesis, University of Oxford, June 1981.
- 12 Thierry Kaufmann, Hervé Pelletier, Serge Vaudenay, and Karine Villegas. When Constant-Time Source Yields Variable-Time Binary: Exploiting Curve25519-donna Built with MSVC 2015. In *Cryptology and Network Security*, pages 573–582, Cham, 2016. Springer International Publishing.
- 13 Paul Kocher, Jann Horn, Anders Fogh, , Daniel Genkin, Daniel Gruss, Werner Haas, Mike Hamburg, Moritz Lipp, Stefan Mangard, Thomas Prescher, Michael Schwarz, and Yuval Yarom. Spectre Attacks: Exploiting Speculative Execution. In *40th IEEE Symposium on Security and Privacy (S&P'19)*, 2019.
- 14 Ramana Kumar, Magnus Myreen, Michael Norrish, and Scott Owens. CakeML: A verified implementation of ML. In *ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 179–191, San Diego, January 2014. ACM Press.
- 15 Xavier Leroy. A Formally Verified Compiler Back-end. *J. Autom. Reason.*, 43(4):363–446, December 2009. [doi:10.1007/s10817-009-9155-4](https://doi.org/10.1007/s10817-009-9155-4).
- 16 Moritz Lipp, Michael Schwarz, Daniel Gruss, Thomas Prescher, Werner Haas, Anders Fogh, Jann Horn, Stefan Mangard, Paul Kocher, Daniel Genkin, Yuval Yarom, and Mike Hamburg. Meltdown: Reading Kernel Memory from User Space. In *27th USENIX Security Symposium (USENIX Security 18)*, 2018.
- 17 Andreas Lochbihler. Mechanising a Type-Safe Model of Multithreaded Java with a Verified Compiler. *Journal of Automated Reasoning*, 61(1):243–332, June 2018. [doi:10.1007/s10817-018-9452-x](https://doi.org/10.1007/s10817-018-9452-x).

- 18 Luísa Lourenço and Luís Caires. Dependent Information Flow Types. In *ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 317–328, Mumbai, India, January 2015. ACM.
- 19 Heiko Mantel, David Sands, and Henning Sudbrock. Assumptions and Guarantees for Compositional Noninterference. In *IEEE Computer Security Foundations Symposium*, pages 218–232, Cernay-la-Ville, France, June 2011. IEEE.
- 20 Toby Murray, Robert Sison, and Kai Engelhardt. COVERN: A logic for compositional verification of information flow control. In *European Symposium on Security and Privacy*, London, United Kingdom, April 2018. IEEE.
- 21 Toby Murray, Robert Sison, Edward Pierzchalski, and Christine Rizkallah. Compositional Security-Preserving Refinement for Concurrent Imperative Programs. *Archive of Formal Proofs*, June 2016. , Formal proof development. URL: http://isa-afp.org/entries/Dependent_SIFUM_Refinement.shtml.
- 22 Toby Murray, Robert Sison, Edward Pierzchalski, and Christine Rizkallah. Compositional Verification and Refinement of Concurrent Value-Dependent Noninterference. In *IEEE Computer Security Foundations Symposium*, pages 417–431, Lisbon, Portugal, June 2016.
- 23 Marco Patrignani, Amal Ahmed, and Dave Clarke. Formal Approaches to Secure Compilation: A Survey of Fully Abstract Compilation and Related Work. *ACM Comput. Surv.*, 51(6):125:1–125:36, February 2019. doi:10.1145/3280984.
- 24 Marco Patrignani and Deepak Garg. Secure Compilation and Hyperproperty Preservation. In *IEEE 30th Computer Security Foundations Symposium, CSF 2017, Santa Barbara, USA, August 21 - 25, 2017*, CSF'17, 2017.
- 25 Marco Patrignani and Deepak Garg. Robustly Safe Compilation. In *Programming Languages and Systems*, pages 469–498, Cham, 2019. Springer International Publishing.
- 26 Anton Podkopaev, Ori Lahav, and Viktor Vafeiadis. Bridging the Gap Between Programming Languages and Hardware Weak Memory Models. *Proc. ACM Program. Lang.*, 3(POPL):69:1–69:31, January 2019. doi:10.1145/3290382.
- 27 Andrei Sabelfeld and David Sands. Probabilistic Noninterference for Multi-Threaded Programs. In *Proceedings of the 13th IEEE Workshop on Computer Security Foundations, CSFW '00*, pages 200–, Washington, DC, USA, 2000. IEEE Computer Society. URL: <http://dl.acm.org/citation.cfm?id=794200.795151>.
- 28 Mark Staples, Ross Jeffery, June Andronick, Toby Murray, Gerwin Klein, and Rafal Kolanski. Productivity for Proof Engineering. In *Empirical Software Engineering and Measurement*, page 15, Turin, Italy, September 2014.
- 29 F. Del Tesesco, D. Sands, and A. Russo. Fault-Resilient Non-interference. In *2016 IEEE 29th Computer Security Foundations Symposium (CSF)*, pages 401–416, June 2016.
- 30 Danfeng Zhang, Yao Wang, G. Edward Suh, and Andrew C. Myers. A Hardware Design Language for Timing-Sensitive Information-Flow Security. In *Proceedings of the Twentieth International Conference on Architectural Support for Programming Languages and Operating Systems, ASPLOS '15*, pages 503–516, New York, NY, USA, 2015. ACM. doi:10.1145/2694344.2694372.