An Extended Account of Trace-relating Compiler Correctness and Secure Compilation

CARMINE ABATE and ROBERTO BLANCO, MPI-SP, Germany

ȘTEFAN CIOBĂCĂ, Alexandru Ioan Cuza University Iași, Romania

ADRIEN DURIER, MPI-SP, Germany

DEEPAK GARG, Max Planck Institute for Software Systems, Germany

CĂTĂLIN HRIȚCU, MPI-SP, Germany

MARCO PATRIGNANI, Stanford, USA and CISPA Helmholtz Center for Information Security, Germany

ÉRIC TANTER, University of Chile, Chile

JÉRÉMY THIBAUT, MPI-SP, Germany

Compiler correctness, in its simplest form, is defined as the inclusion of the set of traces of the compiled program in the set of traces of the original program. This is equivalent to the preservation of all trace properties. Here, traces collect, for instance, the externally observable events of each execution. However, this definition requires the set of traces of the source and target languages to be the same, which is not the case when the languages are far apart or when observations are fine-grained. To overcome this issue, we study a generalized compiler correctness definition, which uses source and target traces drawn from potentially different sets

C. Abate, R. Blanco, A. Durier, C. Hrițcu, É. Tanter, and J. Thibault part of this work was conducted while these authors were employed at or visiting Inria Paris.

This article revises and extends the work of Abate et al. [2] presented at ESOP'20 with the following additions: It contains a more complete account of the classes of properties that can be preserved by correct compilers by discussing safety and hyperproperty preservation. It discusses recent work on the preservation of noninterference through compilation [7, 52, 74] and interprets this work within the presented framework. It unifies the language presentation for the compilers that are proven correct using different relations. It provides a self-contained and in-depth analysis of the classes of properties that can be preserved by secure compilers by discussing subset-closed hyperproperties, hypersafety, 2-relational properties, 2-relational safety, and 2-relational hyperproperties. Generally, this article provides more intuition and explanation for some of the presented notions as well as for the discussed instances of our theory.

Please note that Coq symbols as well as Compiler Criteria are links: the former to Coq files in the external repository https://github.com/secure-compilation/different_traces, the latter to definitions inside this document.

This work was in part supported by the European Research Council under ERC Starting Grant SECOMP (715753), by the German Federal Ministry of Education and Research (BMBF) through funding for the CISPA-Stanford Center for Cybersecurity (FKZ: 13N1S0762), by DARPA grant SSITH/HOPE (FA8650-15-C-7558), by the Office of Naval Research for support through grant N00014-18-1-2620, Accountable Protocol Customization, and by UAIC internal grant 07/2018.

Authors' addresses: C. Abate, R. Blanco, A. Durier, C. Hrițcu, and J. Thibault, MPI-SP, Universitätsstraße 140, Bochum, Germany; emails: {carmine.abate, roberto.blanco, adrien.durier}@mpi-sp.org, catalin.hritcu@gmail.com, jeremy.thibault@mpi-sp.org; Ș. Ciobăcă, Department of Computer Science, Alexandru Ioan Cuza University Iași, Bulevardul Carol I, Nr. 11, 700506, Iași, Romania; email: stefan.ciobaca@info.uaic.ro; D. Garg, Max Planck Institute for Software Systems, Saarland Informatics Campus, Saarbrücken, Germany; email: dg@mpi-sws.org; M. Patrignani, Stanford, 343 Serra Mall USA and CISPA Helmholtz Center for Information Security, Saarland Informatics Campus, Saarbrücken, Germany; email: mp@cs.stanford.edu; É. Tanter, Computer Science Department (DCC), University of Chile, Chile; email: etanter@dcc.uchile.cl.

This work is licensed under a Creative Commons Attribution International 4.0 License.

© 2021 Copyright held by the owner/author(s).

0164-0925/2021/11-ART14 \$15.00

<https://doi.org/10.1145/3460860>

and connected by an arbitrary relation. We set out to understand what guarantees this generalized compiler correctness definition gives us when instantiated with a non-trivial relation on traces. When this trace relation is not equality, it is no longer possible to preserve the trace properties of the source program unchanged. Instead, we provide a generic characterization of the target trace property ensured by correctly compiling a program that satisfies a given source property, and dually, of the source trace property one is required to show to obtain a certain target property for the compiled code. We show that this view on compiler correctness can naturally account for undefined behavior, resource exhaustion, different source and target values, side channels, and various abstraction mismatches. Finally, we show that the same generalization also applies to many definitions of *secure* compilation, which characterize the protection of a compiled program linked against adversarial code.

CCS Concepts: • **Security and privacy** → **Formal security models**; • **Software and its engineering** → **Compilers**; *Software verification*;

Additional Key Words and Phrases: Trace properties, hyperproperties, property-preserving compilation, compiler correctness, secure compilation, Galois connection, formal languages, programming languages

ACM Reference format:

Carmine Abate, Roberto Blanco, Ștefan Ciobăcă, Adrien Durier, Deepak Garg, Cătălin Hrițcu, Marco Patrignani, Éric Tanter, and Jérémy Thibault. 2021. An Extended Account of Trace-relating Compiler Correctness and Secure Compilation. *ACM Trans. Program. Lang. Syst.* 43, 4, Article 14 (November 2021), 48 pages.

<https://doi.org/10.1145/3460860>

1 INTRODUCTION

Compiler correctness is an old idea [46, 49, 50] that has seen a significant revival in recent times. This new wave was started by the creation of the CompCert verified C compiler [41] and continued by the proposal of many significant extensions and variants of CompCert [10, 11, 15, 29, 36, 37, 51, 67, 73, 76, 80] and the success of many other milestone compiler verification projects, including Vellvm [83], Pilsner [56], CakeML [77], and CertiCoq [5]. Verification through proof assistants allows the user of a compiler to trust the proofs without diving into all of the details. Still, to clearly understand the benefits and limitations of using a verified compiler, she has to deeply understand the statement of correctness. This is true not just for correct compilation, but also for secure compilation, which is the more recent idea that a compilation chain should not just provide correctness but also security against co-linked adversarial components [4, 32].

Basic Compiler Correctness. The gold standard for compiler correctness is *semantic preservation*, which intuitively says that the semantics of a compiled program (in the target language) is compatible with the semantics of the original program (in the source language). For practical verified compilers, such as CompCert [41] and CakeML [77], semantic preservation is stated extrinsically, by referring to *traces*. In these two settings, a trace is an ordered sequence of events—such as inputs from and outputs to an external environment—that are produced by the execution of a program.

A basic definition of compiler correctness can be given by the inclusion of the set of traces of the compiled program in the set of traces of the original program. Formally [41]:

Definition 1.1 (Basic Compiler Correctness (CC)). A compiler $\downarrow$ is *correct* (CC) iff

$$\forall W \ t. \ W \downarrow \rightsquigarrow t \Rightarrow W \rightsquigarrow t.$$

This definition says that for any whole¹ source program W , if we compile it (denoted $W \downarrow$), execute it in the semantics of the target language, and observe a trace t , then the original W can

¹For simplicity, for now, we ignore separate compilation and linking, returning to it in Section 6.

produce *the same* trace t in the semantics of the source language.² This definition is simple and easy to understand, since it only references a few familiar concepts: a compiler between a source and a target language, each equipped with a trace-producing semantics (usually nondeterministic).

Beyond Basic Compiler Correctness. Definition 1.1 implicitly assumes that the source and target traces are drawn from the very same set, and requires that any target trace produced by a compiled program can be faithfully reproduced by the source program. In practice, existing verified compiler adopts a less restrictive formulation of compiler correctness:

CompCert [41] The original compiler correctness theorem of CompCert [41] can be seen as an instance of basic compiler correctness, but it does not provide any guarantees for programs that can exhibit undefined behavior [68]. As allowed by the C standard, such unsafe programs are not even considered to be in the source language, so are not quantified over. This has important practical implications, since undefined behavior often leads to exploitable security vulnerabilities [16, 30, 31] and serious confusion even among experienced C and C++ developers [40, 68, 78, 79]. As such, since 2010, CompCert provides an additional top-level correctness theorem³ that better accounts for the presence of unsafe programs by providing guarantees for them up to the point when they encounter undefined behavior [68]. This new theorem goes beyond the basic correctness definition above, as a target trace need only correspond to a source trace *up to the occurrence* of undefined behavior in the source trace.

CakeML [77] Compiler correctness for CakeML accounts for memory exhaustion in target executions. Crucially, memory exhaustion events cannot occur in source traces, only in target traces. Hence, dually to CompCert, compiler correctness only requires source and target traces to coincide up to the occurrence of a memory exhaustion event in the target trace.

Trace-relating Compiler Correctness. Generalized formalizations of compiler correctness like the ones above can be naturally expressed as instances of a uniform definition, which we call *trace-relating compiler correctness*. This generalizes basic compiler correctness by (a) considering that source and target traces belong to *possibly distinct* sets Traces_S and Traces_T , and (b) being parameterized by an arbitrary *trace relation* $\sim$.

Definition 1.2 (Trace-relating Compiler Correctness ($\text{CC}^\sim$)). A compiler $\downarrow$ is *correct* with respect to a trace relation $\sim \subseteq \text{Traces}_S \times \text{Traces}_T$ iff

$$\forall W. \forall t. W \downarrow \rightsquigarrow t \Rightarrow \exists s \sim t. W \rightsquigarrow s.$$

This definition requires that, for any target trace t produced by the compiled program $W \downarrow$, there exists a source trace s that can be produced by the original program W and is *related* to t according to $\sim$ (i.e., $s \sim t$). By choosing the trace relation appropriately, one can recover the different notions of compiler correctness presented above:

Basic CC Take $s \sim t$ to be $s = t$. Trivially, the basic CC of Definition 1.1 is $\text{CC}^\sim$.

CompCert Undefined behavior is modeled in CompCert as a trace-terminating event *Wrong* that can occur in any of its languages (source, target, and all intermediate languages), so for a given phase (or composition thereof), we have $\text{Traces}_S = \text{Traces}_T$. Nevertheless, the relation between source and target traces with which to instantiate $\text{CC}^\sim$ to obtain CompCert's current theorem is the following (note that we denote *finite* traces—or prefixes—as m):

$$s \sim t \equiv s = t \vee (\exists m \leq t. s = m \cdot \text{Wrong}).$$

²Typesetting convention [60]: We use a blue, sans-serif font for source elements, an orange, bold font for target ones, and a black, italic font for elements common to both languages.

³Stated at the top of the CompCert file `driver/Complements.v` and discussed by Regehr [68].

A compiler satisfying $CC^\sim$ for this trace relation can turn a source prefix ending in undefined behavior $m \cdot \text{Wrong}$ (where “.” is concatenation) either into the same prefix in the target (first disjunct) or into a target trace that starts with the prefix m but then continues *arbitrarily* (second disjunct, “ $\leq$ ” is the prefix relation).

CakeML Here, target traces are sequences of symbols from an alphabet Σ_T that has a specific trace-terminating event, **Resource_limit_hit**, which is not available in the source alphabet Σ_S (i.e., $\Sigma_T = \Sigma_S \cup \{\text{Resource_limit_hit}\}$). Then, the compiler correctness theorem of CakeML can be obtained by instantiating $CC^\sim$ with the following $\sim$ relation:

$$s \sim t \equiv s = t \vee (\exists m. m \leq s. t = m \cdot \text{Resource_limit_hit}).$$

The resulting $CC^\sim$ instance relates a target trace ending in **Resource_limit_hit** after executing prefix m to a source trace that first produces m and then continues in a way given by the semantics of the source program.

Beyond undefined behavior and resource exhaustion, there are many other practical uses for $CC^\sim$: In this article, we show that it also accounts for differences between source and target values, for a single source output being turned into a series of target outputs, and for side-channels.

On the flip side, the compiler correctness statement and its implications can be more difficult to understand for $CC^\sim$ than for $CC^=$. The full implications of choosing a particular $\sim$ relation can be subtle. In fact, using a bad relation can make the compiler correctness statement trivial or unexpected. For instance, it should be easy to see that if one uses the total relation, which relates all source traces to all target ones, the $CC^\sim$ property holds for every compiler, yet it might take one a bit more effort to understand that the same is true even for the following relation:

$$s \sim t \equiv \exists W. W \rightsquigarrow s \wedge W \downarrow \rightsquigarrow t.$$

Reasoning about Trace Properties. To understand more about a particular $CC^\sim$ instance, we propose to also look at how it preserves *trace properties*—defined as sets of allowed traces [39]—from the source to the target. For instance, it is well known that $CC^=$ is equivalent to the preservation of all trace properties (where $W \models \pi$ reads “ W satisfies property π ” and stands for $\forall t. W \rightsquigarrow t \Rightarrow t \in \pi$):

$$CC^= \equiv \forall \pi \in 2^{\text{Trace}}. \forall W. W \models \pi \Rightarrow W \downarrow \models \pi.$$

However, to the best of our knowledge, similar results have not been formulated for trace relations beyond equality, when it is no longer possible to preserve the trace properties of the source program unchanged. For trace-relating compiler correctness, where source and target traces can be drawn from different sets and related by an arbitrary trace relation, there are two crucial questions to ask:

- (1) For a source trace property π_S of a program—established for instance by formal verification—what is the strongest target property that any $CC^\sim$ compiler is guaranteed to ensure for the produced target program?
- (2) For a target trace property π_T , what is the weakest source property we need to show of the original source program to obtain π_T for the result of any $CC^\sim$ compiler?

Far from being mere hypothetical questions, they can help the developer of a verified compiler better understand the compiler correctness theorem they are proving, and we expect that any user of such a compiler will need to ask either one or the other if they are to make use of that theorem. In this work, we provide a simple and natural answer to these questions, for any instance of $CC^\sim$. Building upon a bijection between relations and Galois connections [6, 26, 54], we observe that any trace relation $\sim$ corresponds to two *property mappings* $\tilde{\tau}$ and $\tilde{\sigma}$, which are functions mapping source properties to target ones ($\tilde{\tau}$ standing for “to target”) and target properties to source ones

Fig. 1. The equivalent compiler correctness definitions forming our trinitarian view.

($\tilde{\sigma}$ standing for “to source”):

$$\begin{aligned}
\tilde{\tau}(\pi_S) &= \{t \mid \exists s. s \sim t \wedge s \in \pi_S\}, \\
\tilde{\sigma}(\pi_T) &= \{s \mid \forall t. s \sim t \Rightarrow t \in \pi_T\}.
\end{aligned}$$

The *existential image* of $\sim$, $\tilde{\tau}$, answers the first question above by mapping a given source property π_S to the target property that contains all target traces for which *there exists a related source trace* that satisfies π_S . Dually, the *universal image* of $\sim$, $\tilde{\sigma}$, answers the second question by mapping a given target property π_T to the source property that contains all source traces for which *all related target traces* satisfy π_T . We introduce two new correct compilation definitions in terms of **trace property preservation** (TP):

- $\text{TP}^{\tilde{\tau}}$ quantifies over all source trace properties and uses $\tilde{\tau}$ to obtain the corresponding target properties;
- $\text{TP}^{\tilde{\sigma}}$ quantifies over all target trace properties and uses $\tilde{\sigma}$ to obtain the corresponding source properties.

We prove that these two definitions are equivalent to $\text{CC}^\sim$, yielding a novel trinitarian view of compiler correctness (Figure 1).

Contributions.

- We propose a new trinitarian view of compiler correctness that accounts for non-trivial relations between source and target traces. While, as discussed above, specific instances of the $\text{CC}^\sim$ definition have already been used in practice, we seem to be the first to propose assessing the meaningfulness of $\text{CC}^\sim$ instances in terms of how properties are preserved between the source and the target, and in particular by looking at the property mappings $\tilde{\sigma}$ and $\tilde{\tau}$ induced by the trace relation $\sim$. We prove that $\text{CC}^\sim$, $\text{TP}^{\tilde{\sigma}}$, and $\text{TP}^{\tilde{\tau}}$ are equivalent for any trace relation (Section 2.2), as illustrated in Figure 1. In the opposite direction, we show that for every trace relation corresponding to a given Galois connection [26], an analogous equivalence holds.
- We extend these results from the preservation of trace properties to the larger class of subset-closed hyperproperties, e.g., noninterference (Section 3.1),⁴ and to the classes of safety properties (Section 3.2) and all hyperproperties (Section 3.3).
- We use $\text{CC}^\sim$ compilers of various complexities to illustrate that our view on compiler correctness naturally accounts for undefined behavior (Section 4.1), resource exhaustion (Section 4.2), different source and target values (Section 4.3), and differences in the granularity of data and observable events (Section 4.4). We expect these ideas to extend to other discrepancies between source and target traces. For each compiler, we show how to choose the relation between source and target traces and how the induced property mappings preserve interesting trace properties and subset-closed hyperproperties. We look at the way

⁴Given the deterministic nature of our programs, we consider notions of noninterference that are often used in deterministic languages. We leave notions of noninterference in nondeterministic languages for future work.

particular $\tilde{\sigma}$ and $\tilde{\tau}$ work on different kinds of properties and how the produced properties can be expressed for different kinds of traces.

- We analyze the impact of correct compilation on noninterference [28], showing what can still be preserved (and thus also what is lost) when target observations are finer than source ones, e.g., side-channel observations (Section 5). We formalize the guarantee obtained by correct compilation of a noninterfering program as *abstract noninterference* [27], a weakening of target noninterference. Dually, we identify a family of declassifications of target noninterference for which source reasoning is possible.
- We show that the trinitarian view also extends to a large class of *secure compilation* definitions [3], formally characterizing the protection of the compiled program against linked adversarial code (Section 6). For each secure compilation definition, we again propose both a property-free characterization in the style of $\text{CC}^\sim$ and two characterizations in terms of preserving a class of source or target properties satisfied against arbitrary adversarial contexts. The additional quantification over contexts allows for finer distinctions when considering different property classes, so we study mapping classes not only of trace properties and hyperproperties, but also of relational hyperproperties [3].
- We provide instances of secure compilers that preserve three different classes of hyperproperties (trace, safety, and hypersafety properties) when targeting a language with additional trace events that are not possible in the source (Section 7).

The results and insights that we provide often follow one's expected intuition and may be considered unsurprising. However, our framework is the first to capture such expectations formally and precisely, and as such it provides a uniform way to discuss these and to formalize future (possibly surprising) ones. The article closes with discussions of related (Section 8) and future work (Section 9). Some technical proofs can be found in the Appendix (Section A).

The traces considered in our examples are structured, usually as sequences of events. We notice, however, that unless explicitly mentioned, all our definitions and results are more general and make no assumption whatsoever about the structure of traces. Most of the theorems formally or informally mentioned in the article were mechanized in the Coq proof assistant and are marked with . This development has around 10K lines of code and is available at the following address: https://github.com/secure-compilation/different_traces.

2 TRACE-RELATING COMPILER CORRECTNESS

In this section, we start by generalizing the trace property preservation definitions at the end of the introduction to TP^σ and TP^τ , which depend on two *arbitrary* mappings σ and τ (Section 2.1). We prove that, whenever σ and τ form a Galois connection, TP^σ and TP^τ are equivalent (Theorem 2.4). We then exploit a bijective correspondence between trace relations and Galois connections to close the trinitarian view (Section 2.2), with two main benefits: First, it helps us assess the meaningfulness of a given trace relation by looking at the property mappings it induces; second, it allows us to construct new compiler correctness definitions starting from a desired mapping of properties. Finally, we generalize the classic result that compiler correctness (i.e., $\text{CC}^\sim$) is enough to preserve not just trace properties but also all subset-closed hyperproperties [18]. For this, we show that $\text{CC}^\sim$ is also equivalent to subset-closed hyperproperty preservation, for which we also define both a version in terms of $\tilde{\sigma}$ and a version in terms of $\tilde{\tau}$ (Section 3.1).

2.1 Property Mappings

As explained in Section 1, trace-relating compiler correctness $\text{CC}^\sim$, by itself, lacks a crisp description of which trace properties are preserved by compilation. Since even the syntax of traces can

differ between source and target, one can either focus on trace properties of the source (and then interpret them in the target) or on trace properties of the target (and then interpret them in the source). Formally, we need two property mappings, $\tau : 2^{\text{Traces}} \rightarrow 2^{\text{TraceT}}$ and $\sigma : 2^{\text{TraceT}} \rightarrow 2^{\text{Traces}}$, which lead us to the following generalization of trace property preservation (TP):

Definition 2.1 (TP^σ and TP^τ). Given two property mappings, $\tau : 2^{\text{Traces}} \rightarrow 2^{\text{TraceT}}$ and $\sigma : 2^{\text{TraceT}} \rightarrow 2^{\text{Traces}}$, for a compilation chain $\cdot\downarrow$, we define TP^τ and TP^σ as follows:

$$\begin{aligned}\text{TP}^\tau &\equiv \forall \pi_S. \forall W. W \models \pi_S \Rightarrow W\downarrow \models \tau(\pi_S), \\ \text{TP}^\sigma &\equiv \forall \pi_T. \forall W. W \models \sigma(\pi_T) \Rightarrow W\downarrow \models \pi_T.\end{aligned}$$

For an arbitrary source program W , τ interprets a source property π_S as the *target guarantee* for $W\downarrow$. Dually, σ defines a *source obligation* sufficient for the satisfaction of a target property π_T after compilation. Ideally:

- (i) Given π_T , the target interpretation of the source obligation $\sigma(\pi_T)$ should actually guarantee that π_T holds, i.e., $\tau(\sigma(\pi_T)) \subseteq \pi_T$;
- (ii) Dually for π_S , we would not want the source obligation for $\tau(\pi_S)$ to be harder than π_S itself, i.e., $\sigma(\tau(\pi_S)) \supseteq \pi_S$.

These requirements are satisfied when the two maps form a *Galois connection* between the posets of source and target properties ordered by inclusion. We briefly recall the definition and the characteristic property of Galois connections [20, 47].

Definition 2.2 (*Galois Connection*). Let $(X, \leq)$ and $(Y, \sqsubseteq)$ be two posets. A pair of maps, $\alpha : X \rightarrow Y$, $\gamma : Y \rightarrow X$ is a Galois connection *iff* it satisfies the *adjunction law*: $\forall x \in X. \forall y \in Y. \alpha(x) \sqsubseteq y \iff x \leq \gamma(y)$. α (respectively, γ) is the lower (upper) adjoint or abstraction (concretization) function and Y (X) the abstract (concrete) domain.

We will often write $\alpha : (X, \leq) \sqsubseteq (Y, \sqsubseteq) : \gamma$ to denote a Galois connection, or simply $\alpha : X \sqsubseteq Y : \gamma$, or even $\alpha \sqsubseteq \gamma$ when the involved posets are clear from context.

LEMMA 2.3 (CHARACTERISTIC PROPERTY OF GALOIS CONNECTIONS). *If $\alpha : (X, \leq) \sqsubseteq (Y, \sqsubseteq) : \gamma$ is a Galois connection, then α, γ are monotone and $\text{id} \leq \gamma \circ \alpha$ and $\alpha \circ \gamma \sqsubseteq \text{id}$, i.e.,*

$$\begin{aligned}\forall x \in X. x \leq \gamma(\alpha(x)) \\ \forall y \in Y. \alpha(\gamma(y)) \sqsubseteq y.\end{aligned}$$

If X, Y are complete lattices, then α is continuous, i.e., $\forall F \subseteq X. \alpha(\bigsqcup F) = \bigsqcup \alpha(F)$.

If two property mappings, τ and σ , form a Galois connection on trace properties ordered by set inclusion, then Lemma 2.3 (with $\alpha = \tau$ and $\gamma = \sigma$) tells us that they satisfy conditions (i), (ii) above, i.e., $\tau(\sigma(\pi_T)) \subseteq \pi_T$ and $\sigma(\tau(\pi_S)) \supseteq \pi_S$.⁵ These conditions on τ and σ are sufficient to show the equivalence of the criteria they define, respectively, TP^τ and TP^σ .

THEOREM 2.4 (TP^τ AND TP^σ COINCIDE). *Let $\tau : 2^{\text{Traces}} \rightrightarrows 2^{\text{TraceT}} : \sigma$ be a Galois connection, with τ and σ the lower and upper adjoints (respectively). Then $\text{TP}^\tau \iff \text{TP}^\sigma$.*

PROOF. Notice that if a program satisfies a property π , then it satisfies every less restrictive i.e., bigger property $\pi' \supseteq \pi$. Building on this:

- ($\Rightarrow$) Assume TP^τ and that W satisfies $\sigma(\pi_T)$. Apply TP^τ to W and $\sigma(\pi_T)$ and deduce that $W\downarrow$ satisfies $\tau(\sigma(\pi_T)) \subseteq \pi_T$.

⁵While target traces are often “more concrete” than source ones, trace properties 2^{Trace} (which in Coq we represent as the function type $\text{Trace} \rightarrow \text{Prop}$) are contravariant in Trace and thus target properties correspond to the *abstract domain*.

($\Leftarrow$) Assume TP^σ and that W satisfies $\pi_S \subseteq \sigma(\tau(\pi_S))$. Apply TP^σ to W and $\sigma(\tau(\pi_S))$ deducing $W \downarrow$ satisfies $\tau(\pi_S)$. $\square$

2.2 Trace Relations and Property Mappings

We now investigate the relation between $CC^\sim$, TP^τ , and TP^σ . We show that for a trace relation and its corresponding Galois connection (Lemma 2.7), the three criteria are equivalent (Theorem 2.8). This equivalence offers interesting insights for both verification and the design of a correct compiler. For a $CC^\sim$ compiler, the equivalence makes explicit both the guarantees one has after compilation ($\tilde{\tau}$) and source proof obligations to ensure the satisfaction of a given target property ($\tilde{\sigma}$). However, a compiler designer might first determine the target guarantees the compiler itself must provide, i.e., τ , and then prove an equivalent statement, $CC^\sim$, for which more convenient proof techniques exist in the literature [9, 77].

Definition 2.5 (Existential and Universal Image [26]). Given any two sets X and Y and a relation $\sim \subseteq A \times B$, define the relation's existential or direct image, $\tilde{\tau} : 2^X \rightarrow 2^Y$ and its universal image, $\tilde{\sigma} : 2^Y \rightarrow 2^X$ as follows:

$$\begin{aligned}\tilde{\tau} &= \lambda \pi \in 2^X. \{y \mid \exists x. x \sim y \wedge x \in \pi\} \\ \tilde{\sigma} &= \lambda \pi \in 2^Y. \{x \mid \forall y. x \sim y \Rightarrow y \in \pi\}.\end{aligned}$$

When trace relations are considered, the corresponding existential and universal images can be used to instantiate Definition 2.1 leading to the trinitarian view already mentioned in Section 1.

THEOREM 2.6 (TRINITARIAN VIEW $\mathfrak{A}$). *For any trace relation $\sim$ and its existential and universal images $\tilde{\tau}$ and $\tilde{\sigma}$, we have:*

This result relies both on Theorem 2.4 and on the fact that the existential and universal images of a trace relation form a Galois connection ($\mathfrak{A}$). The theorem can be stated in a slightly more general form (Theorem 2.8), exploiting an isomorphism between the category of sets and relations and a subcategory of monotonic predicate transformers [26]. We specialize this isomorphism to what is of interest for our purposes and deduce a bijective correspondence between trace relations and Galois connections on properties.

LEMMA 2.7 (TRACE RELATIONS $\cong$ GALOIS CONNECTIONS ON TRACE PROPERTIES). *The function $\sim \mapsto \tilde{\tau} \sqsubseteq \tilde{\sigma}$ that maps a trace relation to its existential and universal images is a bijection between trace relations $2^{\text{Trace}_S \times \text{Trace}_T}$ and Galois connections on trace properties $2^{\text{Trace}_S} \sqsubseteq 2^{\text{Trace}_T}$. Its inverse is $\tau \sqsubseteq \sigma \mapsto \hat{\sim}$, where $\hat{\sim} \equiv \tau \in \tau(\{s\})$.*

The bijection just introduced allows us to generalize Theorem 2.6 and switch anytime between the three views of compiler correctness described earlier.

THEOREM 2.8 (CORRESPONDENCE OF CRITERIA). *For any trace relation $\sim$ and corresponding Galois connection $\tau \sqsubseteq \sigma$, we have: $TP^\tau \iff CC^\sim \iff TP^\sigma$.*

Note that sometimes the lifted properties may be trivial: The target guarantee can be the true property (the set of all traces) or the source obligation the false property (the empty set of traces). This might be the case when source observations abstract away too much information (Section 4.2 presents an example).

3 PRESERVING OTHER (HYPER)PROPERTY CLASSES

In this section, we investigate how to preserve other classes of (hyper)properties beyond trace properties: subset-closed hyperproperties (Section 3.1), safety properties (Section 3.2), and arbitrary hyperproperties that are not just subset-closed (Section 3.3). For each of these classes, we start by giving an intuition of what it means to preserve such a class in the equal-trace setting, then we study preservation of that class in the trace-relating setting. For subset-closed hyperproperties, we have to refine the Galois connection to ensure the information “ H_S is subset-closed” is not lost with the application of $\tilde{\tau}$. Similarly, when looking at safety properties, we have to preserve the information that a property is a safety property. For arbitrary hyperproperties one might instead require that no information at all is lost during the (pre or post) composition of $\tilde{\tau}$ and $\tilde{\sigma}$. The section concludes with a comparison of the criteria in terms of relative strengths (Section 3.4).

3.1 Preservation of Subset-closed Hyperproperties

Hyperproperty preservation is a strong requirement in general. Fortunately, many interesting hyperproperties are **subset-closed** (*SCH for short*) (e.g., noninterference), and these are known to be preserved by refinement [18]. When the trace semantics is common to source and target languages, a subset-closed hyperproperty is preserved if the behaviors of the compiled program refine the behaviors of the source program, which coincides with the statement of $CC^\sim$. We generalize this result to the trace-relating setting, introducing two other equivalent characterizations of $CC^\sim$ in terms of preservation of subset-closed hyperproperties (Theorem 3.3). To do so, we close under subsets the images of both $\tilde{\tau}$ and $\tilde{\sigma}$ so source subset-closed hyperproperties are mapped to target subset-closed ones and vice versa.

First, a hyperproperty H is defined as a set of sets of traces, $H \in 2^{2^{\text{Traces}}}$ (recall that Traces is the set of all traces) [18]. A program satisfies a hyperproperty when its complete set of traces, which from now on we will call its *behavior*, is a member of the hyperproperty.

Definition 3.1 (Hyperproperty Satisfaction [18]). A program W satisfies a hyperproperty H , written $W \models H$,⁶ iff $\text{beh}(W) \in H$, where $\text{beh}(W) = \{t \mid W \rightsquigarrow t\}$.

To talk about hyperproperty preservation in the trace-relating setting, we need an interpretation of source hyperproperties into the target and vice versa. The one we consider builds on top of the two trace property mappings τ and σ , which are naturally lifted to hyperproperty mappings. This way, we are able to extract two hyperproperty mappings from a trace relation similarly to Section 2.2:

Definition 3.2 (Lifting Property Mappings to Hyperproperty Mappings). Let $\tau : 2^{\text{Traces}} \rightarrow 2^{\text{Traces}_T}$ and $\sigma : 2^{\text{Traces}_T} \rightarrow 2^{\text{Traces}}$ be arbitrary property mappings. The images of $H_S \in 2^{2^{\text{Traces}}}$, $H_T \in 2^{2^{\text{Traces}_T}}$ under τ and σ are, respectively:

$$\tilde{\tau}(H_S) = \{\tau(\pi_S) \mid \pi_S \in H_S\}, \quad \tilde{\sigma}(H_T) = \{\sigma(\pi_T) \mid \pi_T \in H_T\}.$$

Formally, we are defining two new mappings, this time on hyperproperties, but with a small abuse of notation, we still denote them $\tilde{\tau}$ and $\tilde{\sigma}$.

Interestingly, it is not possible to apply the argument used for $CC^\sim$ to show that a $CC^\sim$ compiler guarantees $\mathcal{W} \downarrow \models \tilde{\tau}(H_S)$ whenever $\mathcal{W} \models H_S$. This is because direct images do not necessarily preserve subset-closure [44, 55]. We therefore close the image of $\tilde{\tau}$ and $\tilde{\sigma}$ under subsets (denoted as $Cl_\subseteq$) and obtain the following result:

⁶In case of ambiguity with property satisfaction the class of H will be made explicit.

THEOREM 3.3 (PRESERVATION OF SUBSET-CLOSED HYPERPROPERTIES $\rightsquigarrow$). *For any trace relation $\sim$ and its existential and universal images lifted to hyperproperties, $\tilde{\tau}$ and $\tilde{\sigma}$, and for $Cl_{\subseteq}(H) = \{\pi \mid \exists \pi' \in H. \pi \subseteq \pi'\}$, we have the following:*

$$SCHP^{Cl_{\subseteq} \circ \tilde{\tau}} \equiv \forall W \forall H_S \in SCH_S. W \models H_S \Rightarrow W \downarrow \models Cl_{\subseteq}(\tilde{\tau}(H_S));$$

$$SCHP^{Cl_{\subseteq} \circ \tilde{\sigma}} \equiv \forall W \forall H_T \in SCH_T. W \models Cl_{\subseteq}(\tilde{\sigma}(H_T)) \Rightarrow W \downarrow \models H_T.$$

The use of $Cl_{\subseteq}$ in Theorem 3.3 implies a loss of precision in preserving subset-closed hyperproperties through compilation. In Section 5, we focus on a specific security-relevant subset-closed hyperproperty, noninterference, and show that such a loss of precision can be seen as a declassification. Instead, now we define the trinity and the related formal machinery for safety properties preservation.

3.2 Preserving Safety Properties

The class of *Safety* properties collects all trace properties prescribing that “something bad never happens” or equivalently, all trace properties whose violation can be monitored and, once observed, no longer restored [18]. More abstractly, safety properties can be defined as the closed sets of a topology [18, 58], with no need to consider any particular structure on the traces. To ease the presentation, we consider the trace model adopted by Abate et al. [3] where traces resemble lists and streams of events. This model naturally comes with a notion of *prefixes* and a relation between a prefix m and a trace t , written $m \leq t$. Intuitively, π is a safety property if any trace t violating the property extends a “bad prefix” m that witnesses such a violation. Every safety property is therefore uniquely defined by the set of its “bad prefixes.” We recall below the definition and the characterization of safety properties in terms of sets of finite prefixes m .

Definition 3.4 (Safety Properties [18]). Let π be a trace property. Then,

$$\pi \in \text{Safety} \text{ iff } \forall t \notin \pi. \exists m \leq t. \forall t'. m \leq t' \Rightarrow t' \notin \pi.$$

Equivalently, $\pi \in \text{Safety}$ iff there exists a set of finite prefixes M , such that

$$\forall t. t \notin \pi \iff (\exists m \in M. m \leq t).$$

Due to this characterization of safety properties through finite prefixes (Definition 3.4), the preservation of all and only the safety properties is equivalent to $CC^=$ restricted to finite prefixes.

THEOREM 3.5. *The following are equivalent:*

$$SC^= \equiv \forall W, m. W \downarrow \rightsquigarrow^* m \Rightarrow W \rightsquigarrow^* m,$$

$$SP \equiv \forall \pi \in \text{Safety}. W \models \pi \Rightarrow W \downarrow \models \pi,$$

where $W \rightsquigarrow^* m$ stands for $\exists t. m \leq t \wedge W \rightsquigarrow t$.

Unfolding $\rightsquigarrow^*$, we can interpret $SC^=$ as follows: Whenever $W \downarrow$ produces a trace $t \geq m$ that violates a specific safety property, namely, the one defined by the singleton prefix set $\{m\}$, then W violates the *same* safety property, producing a trace $t' \geq m$ but possibly distinct from t .

The generalization we propose of $SC^=$ to the trace-relating setting, states that whenever $W \downarrow$ produces a trace t that violates a target safety property, then W violates the source *interpretation* of the property, i.e., its image through $\tilde{\sigma}$.⁷ The following theorem defines $SC^{\sim}$ and its two equivalent formulations:

⁷At least one other symmetric generalization is possible: For $\pi_S \in \text{Safetys}$ defined by $M = \{m\}$, if $W \downarrow$ produces a trace t that violates the target interpretation of π_S , i.e., $\tilde{\tau}(\pi_S)$, then W produces $s \geq m$, thus violating π_S .

Fig. 2. Composition of $\tilde{\tau} \sqsubseteq \tilde{\sigma}$ and $\text{Safe} \sqsubseteq \text{id}$.

THEOREM 3.6 (TRINITARIAN VIEW FOR SAFETY). *For a trace relation $\sim \subseteq \text{Traces} \times \text{Trace}_T$ and its corresponding property mappings $\tilde{\sigma}$ and $\tilde{\tau}$, the following are equivalent:*

$$\begin{aligned}
 \text{SC}^\sim &\equiv \forall W \forall t \forall m \leq t. W \downarrow \rightsquigarrow t \Rightarrow \exists t' \geq m \exists s \sim t'. W \rightsquigarrow s, \\
 \text{SP}^{\tilde{\sigma}} &\equiv \forall W \forall \pi_T \in \text{Safety}_T. W \models \tilde{\sigma}(\pi_T) \Rightarrow W \downarrow \models \pi_T, \\
 \text{TP}^{\text{Safe} \circ \tilde{\tau}} &\equiv \forall W \forall \pi_S \in 2^{\text{Traces}_S}. W \models \pi_S \Rightarrow W \downarrow \models (\text{Safe} \circ \tilde{\tau})(\pi_S).
 \end{aligned}$$

Coherent with the informal meaning we aimed to give to $\text{SC}^\sim$, $\text{SP}^{\tilde{\sigma}}$ quantifies over target safety properties, while $\text{SP}^{\tilde{\tau}}$ quantifies over *arbitrary* source properties, but imposes the composition of $\tilde{\tau}$ with Safe , which maps an arbitrary target property π_T to the target safety property that best over-approximates π_T .⁸ More precisely, Safe is a closure operator on target properties, with $\text{Safety}_T = \{\text{Safe}(\pi_T) \mid \pi_T \in 2^{\text{Trace}_T}\}$ being the class of target safety properties.

In Figure 2 the blue and red ellipses represent source and target properties properties, respectively, and are connected by $\tilde{\tau} \sqsubseteq \tilde{\sigma}$. The red ellipse is the class of all target safety properties. $\text{Safe} \sqsubseteq \text{id}$ is a Galois connection between target properties and the target safety properties, as Safe is a closure operator [21]. Finally, the composition of Galois connections is still a Galois connection [21]. Hence,

$$\text{Safe} \circ \tilde{\tau} : 2^{\text{Traces}_S} \sqsubseteq \text{Safety}_T : \tilde{\sigma}$$

is a Galois connection between source properties and target safety properties, which we used to prove the equivalence $\text{SP}^{\tilde{\tau}} \iff \text{SP}^{\tilde{\sigma}}$ (♣). We notice that this argument generalizes to arbitrary closure operators on target properties (♣). We come back to this in Section 6, where more such results will be needed when considering other classes of properties being preserved by secure compilers. Now, we define the trinity for arbitrary hyperproperties, not just the subset-closed ones.

3.3 Preserving Non-subset Closed Hyperproperties

Subset-closed hyperproperties are not expressive enough to all capture interesting properties, e.g., possibilistic notions of information-flow [18], so we aim to briefly discuss the preservation of *arbitrary* hyperproperties. In general, one cannot lift a Galois connection over trace properties to a Galois connection over arbitrary hyperproperties.

While two out of three of the criteria we introduce in this section are equivalent under no assumptions ($\text{HC}^\sim \iff \text{HP}^{\tilde{\tau}}$), for a comparison with the third one, we require that no information

⁸ $\text{Safe}(\pi_T) = \cap \{S_T \mid \pi_T \subseteq S_T \wedge S_T \in \text{Safety}_T\}$ is the topological closure in the topology where safety properties coincide with the closed sets (see, e.g., Clarkson and Schneider [18] and Pasqua and Mastroeni [58]).

Fig. 3. Generalization of Compiler Correctness and its trace-relating variations.

is lost in the pre or post composition of τ and σ . For this, we label the trinity in Theorem 3.8 as *weak*.

To start, we note that the following strengthening of $CC^=$, denoted $HC^=$, is equivalent to the preservation of arbitrary hyperproperties. Here, $\text{beh}(W)$ is the set of all traces of W :

THEOREM 3.7 ($HC^=$, HP). *The following are equivalent:*

$$HC^= \equiv \forall W. \text{beh}(W \downarrow) = \text{beh}(W),$$

$$HP \equiv \forall W \forall H \in 2^{\text{Trace}}. W \models H \iff W \downarrow \models H.$$

$HC^=$ requires that the behavior of $W \downarrow$ is exactly the same as the behavior of W . We generalize this to the trace-relating setting by requiring that the behavior of $W \downarrow$ coincide with the target interpretation of the source properties describing the behavior of W ⁹.

THEOREM 3.8 (WEAK TRINITY FOR HYPERPROPERTIES ⁹). *For a trace relation $\sim \subseteq \text{Traces}_S \times \text{Traces}_T$ and induced property mappings $\tilde{\sigma}$ and $\tilde{\tau}$, we have:*

$$HC^{\sim} \iff HP^{\tilde{\tau}};$$

if $\tilde{\tau} \sqsubseteq \tilde{\sigma}$ is a Galois insertion (i.e., $\tilde{\tau} \circ \tilde{\sigma} = \text{id}$), then $HC^{\sim} \Rightarrow HP^{\tilde{\sigma}}$,

if $\tilde{\sigma} \sqsubseteq \tilde{\tau}$ is a Galois reflection (i.e., $\tilde{\sigma} \circ \tilde{\tau} = \text{id}$), then $HP^{\tilde{\sigma}} \Rightarrow HP^{\tilde{\tau}}$,

$$HC^{\sim} \equiv \forall W. \text{beh}(W \downarrow) = \tilde{\tau}(\text{beh}(W)),$$

$$HP^{\tilde{\tau}} \equiv \forall W \forall H_S. W \models H_S \Rightarrow W \downarrow \models \tilde{\tau}(H_S),$$

$$HP^{\tilde{\sigma}} \equiv \forall W \forall H_T. W \models \tilde{\sigma}(H_T) \Rightarrow W \downarrow \models H_T.$$

In other words, it is still possible (and sound) to deduce a source obligation for a given target hyperproperty H_T ($HC^{\sim} \Rightarrow HP^{\tilde{\sigma}}$) when no information is lost in the composition $\tilde{\tau} \circ \tilde{\sigma}$. Dually, $HP^{\tilde{\tau}}$ (and hence $HC^{\sim}$) is a consequence of $HP^{\tilde{\sigma}}$ when no information is lost in composing in the other direction, $\tilde{\sigma} \circ \tilde{\tau}$.

3.4 Comparing the Presented Criteria

At this point, we have presented four trinities of criteria that preserve trace properties, subset-closed hyperproperties, safety properties, and arbitrary hyperproperties. Figure 3 sums up our trinities and orders them according to their relative strength.

In Section 6, we will also consider, in the setting of *secure* compilation, the class of safety hyperproperties or hypersafety, and relational hyperproperties. In the setting of *correct* compilation—which focuses only on whole programs—it is straightforward to show that the trinity for hypersafety coincides with the one for safety properties in the same way the trinity of trace properties

⁹At least one generalization is possible: $\tilde{\sigma}(\text{beh}(W \downarrow)) = \text{beh}(W)$. In this case, $HC^{\sim} \iff HP^{\tilde{\sigma}}$ holds unconditionally while the other two implications hold under the same, but swapped, hypotheses from Theorem 3.8.

and subset-closed hyperproperties coincide. Similarly the trinity for relational hyperproperties coincides with the one for hyperproperties.

4 INSTANCES OF TRACE-RELATING COMPILER CORRECTNESS

The trace-relating view of compiler correctness above can serve as a unifying framework for studying a range of interesting compilers. This section provides several representative instantiations of the framework: source languages with undefined behavior that compilation can turn into arbitrary target behavior (Section 4.1), target languages with resource exhaustion that cannot happen in the source (Section 4.2), changes in the representation of values (Section 4.3), and differences in the granularity of data and observable events (Section 4.4).

4.1 Undefined Behavior

We start by expanding upon the discussion of undefined behavior in Section 1. We first study the model of CompCert, where source and target alphabets are the same, including the event for undefined behavior. The trace relation weakens equality by allowing undefined behavior to be replaced with an arbitrary sequence of events.

Example 4.1 (CompCert-like Undefined Behavior Relation). Source and target traces are sequences of events drawn from Σ , where $Wrong \in \Sigma$ is a terminal event that represents an undefined behavior. We then use the trace relation defined in the introduction:

$$s \sim t \equiv s = t \vee \exists m \leq t. s = m \cdot Wrong.$$

Each trace of a target program produced by a $CC^\sim$ compiler either also is a trace of the original source program or has a finite prefix that the source program also produces, immediately before encountering undefined behavior. As explained in Section 1, one of the correctness theorems in CompCert can be rephrased as this variant of $CC^\sim$.

We proved that the property mappings induced by the relation can be written as ($\clubsuit$):

$$\begin{aligned} \tilde{\sigma}(\pi_T) &= \{s \mid s \in \pi_T \wedge s \neq m \cdot Wrong\} \cup \{m \cdot Wrong \mid \forall t. m \leq t \Rightarrow t \in \pi_T\}, \\ \tilde{\tau}(\pi_S) &= \{t \mid t \in \pi_S\} \cup \{t \mid \exists m \leq t. m \cdot Wrong \in \pi_S\}. \end{aligned}$$

These two mappings explain what a $CC^\sim$ compiler ensures for the $\sim$ relation above. The target-to-source mapping $\tilde{\sigma}$ states that to prove that a compiled program has a property π_T using source-level reasoning, one has to prove that any trace produced by the source program must either be a target trace satisfying π_T or have undefined behavior, but only provided that *any continuation* of the trace substituted for the undefined behavior satisfies π_T . The source-to-target mapping $\tilde{\tau}$ states that by compiling a program satisfying a property π_S , we obtain a program that produces traces that satisfy the same property or that extend a source trace that ends in undefined behavior.

These definitions can help us reason about programs. For instance, $\tilde{\sigma}$ specifies that, to prove that an event does not happen in the target, it is not enough to prove that it does not happen in the source: It is also necessary to prove that the source program does not have any undefined behavior (second disjunct). Indeed, if it had an undefined behavior, its continuations could exhibit the unwanted event.

This relation can be easily generalized to other settings. For instance, consider the setting in which we compile down to a low-level language like machine code. Target traces can now contain new events that cannot occur in the source: Indeed, in modern architectures like x86 a compiler typically uses only a fraction of the available instruction set. Some instructions might even perform dangerous operations, such as writing to the hard drive or controlling a device that is hidden from the source language. Formally, the source and target do not have the same events anymore. Thus,

we consider a source alphabet, $\Sigma_S = \Sigma \cup \{\text{Wrong}\}$, and a target alphabet, $\Sigma_T = \Sigma \cup \Sigma'$. The trace relation is defined in the same way and we obtain the same property mappings as above, except that target traces now have more events (some of which may be dangerous), the arbitrary continuations of target traces get more interesting. For instance, consider a new event that represents writing data on the hard drive, and suppose we want to prove that this event cannot happen for a compiled program. Then, proving this property requires exactly proving that the source program exhibits no undefined behavior [14]. More generally, what one can prove about target-only events can only be either that they cannot appear (because there is no undefined behavior) or that any of them can appear (in the case of undefined behavior).

In Section 7.1, we study a similar example, showing that even in a safe language linked adversarial contexts can cause dangerous target events that have no source correspondent.

4.2 Resource Exhaustion

Let us return to the discussion about resource exhaustion in Section 1.

Example 4.2 (Resource Exhaustion). We consider traces made of events drawn from Σ_S in the source, and $\Sigma_T = \Sigma_S \cup \{\text{Resource_Limit_Hit}\}$ in the target. Recall the trace relation for resource exhaustion:

$$s \sim t \equiv s = t \vee \exists m \leq s. t = m \cdot \text{Resource_Limit_Hit}$$

Formally, this relation is similar to the one for undefined behavior, except this time it is the target trace that is allowed to end early instead of the source trace.

The induced trace property mappings $\tilde{\sigma}$ and $\tilde{\tau}$ are the following (♣):

$$\begin{aligned} \tilde{\sigma}(\pi_T) &= \{s \mid s \in \pi_T\} \cap \{s \mid \forall m \leq s. m \cdot \text{Resource_Limit_Hit} \in \pi_T\}, \\ \tilde{\tau}(\pi_S) &= \pi_S \cup \{m \cdot \text{Resource_Limit_Hit} \mid \exists s \in \pi_S. m \leq s\}. \end{aligned}$$

These capture the following intuitions: The target-to-source mapping $\tilde{\sigma}$ states that to prove a property of the compiled program one has to show that the traces of the source program satisfy two conditions: (1) they must also satisfy the target property; and (2) the termination of every one of their prefixes by a resource exhaustion error must be allowed by the target property. This is rather restrictive: Any property that prevents resource exhaustion cannot be proved using source-level reasoning. Indeed, if π_T does not allow resource exhaustion, then $\tilde{\sigma}(\pi_T) = \emptyset$. This is to be expected, since resource exhaustion is simply not accounted for at the source level. The source-to-target mapping $\tilde{\tau}$ states that a compiled program produces traces that either belong to the same properties as the traces of the source program or end early due to resource exhaustion.

In this example, safety properties [39] are mapped (in both directions) to other safety properties (♣). This can be desirable for a relation: Since safety properties are usually easier to reason about, one interested only in safety properties at the target can reason about them using source-level reasoning tools for safety properties. To reason about safety, one would use the criteria presented in Section 3.2.

Since it focuses on traces and not just safety, the compiler correctness theorem in CakeML is an instance of $\text{CC}^\sim$ for the $\sim$ relation above. We have also implemented two small compilers that are correct for this relation. The full details can be found in the Coq development. The first compiler (♣) goes from a simple expression language (similar to the one in Section 4.3 but without inputs) to the same language except that execution is bounded by some amount of fuel: Each execution step consumes some amount of fuel and execution immediately halts when it runs out of fuel. The compiler is the identity.

The second compiler (♣) is more interesting: We proved this $\text{CC}^\sim$ instance for a variant of a compiler from a WHILE language to a simple stack machine by Xavier Leroy [43]. We enriched

the two languages with outputs and modified the semantics of the stack machine so it falls into an error state if the stack reaches a certain size. The proof uses a standard forward simulation modified to account for failure: If the source execution takes a step from a configuration to another configuration emitting some event (which can be a silent event), then there are two possibilities for a related target configuration: Either (i) it can take some steps to another configuration related to the second source configuration and emit the same event (as in a standard simulation); or (ii) it can take some steps to an error state without emitting any events. The latter corresponds to the case of a resource exhaustion error: The target execution can terminate early, producing only a prefix of the source execution trace, as allowed by the relation.

We conclude this subsection by noting that the resource exhaustion relation and the undefined behavior relation from the previous subsection can easily be combined. Indeed, given a relation $\sim_{UB}$ and a relation $\sim_{RE}$ defined as above on the same sets of traces, we can build a new relation $\sim$ that allows both refinement of undefined behavior and resource exhaustion by taking their union: $s \sim t \equiv s \sim_{UB} t \vee s \sim_{RE} t$. A compiler that is $CC^{\sim_{UB}}$ or $CC^{\sim_{RE}}$ is trivially $CC^{\sim}$, though the converse is not true.

4.3 Different Source and Target Values

This section first presents the common language formalization (Section 4.3.1) that the following (Section 4.3.2) and later instances (Section 4.4 and Section 7.1) build upon. This shared language formalization does not contain a key language feature, namely, the expressions that generate actions and thus labels. This is because each instance deals with specific ways to generate actions, so each instance will define its own extension to each of the languages defined below. Additionally, each instance will define its own compiler and the trace relation used to attain $CC^{\sim}$.

4.3.1 Shared Source and Target Language Formalization. The source language is a pure, statically typed expression language whose expressions e include naturals, Booleans, a Boolean conditional and a conditional for expressions that reduce to 0, arithmetic and relational operations, and sequencing.

$$e ::= n \mid b \mid \text{if } e \text{ then } e \text{ else } e \mid \text{ifz } e \text{ then } e \text{ else } e \mid e \text{ op } e \mid e; e'$$

$$\text{op} ::= + \mid \times \mid \leq \mid == \quad \text{ty} ::= B \mid N$$

Types ty are either N (naturals) or B (Booleans) and typing is standard.

$\frac{}{\vdash n : N} \quad \text{(Type-nat)}$	$\frac{}{\vdash b : B} \quad \text{(Type-bool)}$	$\frac{\vdash e_1 : N \quad \vdash e_2 : N}{\vdash e_1 \cdot e_2 : N} \quad \text{(Type-plus-times)} \quad \cdot = + \text{ or } \times$	$\frac{\vdash e_1 : N \quad \vdash e_2 : N}{\vdash e_1 \leq e_2 : B} \quad \text{(Type-le)}$
---	--	--	--

$\frac{\vdash e_1 : B \quad \vdash e_2 : \text{ty} \quad \vdash e_3 : \text{ty}}{\vdash \text{if } e_1 \text{ then } e_2 \text{ else } e_3 : \text{ty}} \quad \text{(Type-ite)}$	$\frac{\vdash e_1 : N \quad \vdash e_2 : \text{ty} \quad \vdash e_3 : \text{ty}}{\vdash \text{ifz } e_1 \text{ then } e_2 \text{ else } e_3 : \text{ty}} \quad \text{(Type-izte)}$
--	--

The language semantics deal with actions i , lists of actions is , and expression results r . A list of actions is is a list of individual actions i , which are instance-dependant and thus presented later; the same holds for source traces s .

$$r ::= n \mid b \quad i, s ::= \text{instance-specific} \quad is ::= i \cdot is \mid \emptyset$$

The source language has a standard big-step operational semantics ($e \rightsquigarrow \langle is, r \rangle$) that tells how an expression e generates a list of actions and a result $\langle is, r \rangle$.

(Sem-nat) $\frac{}{n \rightsquigarrow \langle \emptyset, n \rangle}$	(Sem-bool) $\frac{}{b \rightsquigarrow \langle \emptyset, b \rangle}$	(Sem-op-nat) $\frac{e_1 \rightsquigarrow \langle is_1, n_1 \rangle \quad e_2 \rightsquigarrow \langle is_2, n_2 \rangle \quad op \in \{+, \times\}}{e_1 \text{ op } e_2 \rightsquigarrow \langle is_1 \cdot is_2, (n_1 \text{ op } n_2) \rangle}$
(Sem-le) $\frac{e_1 \rightsquigarrow \langle is_1, n_1 \rangle \quad e_2 \rightsquigarrow \langle is_2, n_2 \rangle}{e_1 \leq e_2 \rightsquigarrow \langle is_1 \cdot is_2, (n_1 \leq n_2) \rangle}$	(Sem-ite) $\frac{e \rightsquigarrow \langle is, b \rangle \quad b?i = 1 : i = 2 \quad e_i \rightsquigarrow \langle is_i, r_i \rangle}{\text{if } e \text{ then } e_1 \text{ else } e_2 \rightsquigarrow \langle is \cdot is_i, r_i \rangle}$	
(Sem-izte) $\frac{e \rightsquigarrow \langle is, n \rangle \quad n == 0?i = 1 : i = 2 \quad e_i \rightsquigarrow \langle is_i, r_i \rangle}{\text{ifz } e \text{ then } e_1 \text{ else } e_2 \rightsquigarrow \langle is \cdot is_i, r_i \rangle}$	(Sem-seq) $\frac{e \rightsquigarrow \langle is, r \rangle \quad e' \rightsquigarrow \langle is', r' \rangle}{e; e' \rightsquigarrow \langle is \cdot is', r' \rangle}$	

The target language is analogous to the source one, except that it is untyped, it only has naturals n , and its only conditional is **ifz e then e else e** .

$e ::= n \mid e \text{ op } e \mid \text{ifz } e \text{ then } e \text{ else } e \mid e; e'$ $op ::= + \mid \times$ $r ::= n$
 $i, t ::= \text{instance-specific}$ $is ::= i \cdot is \mid \emptyset$

The semantics of the target language is also given in big-step style; since its rules are a subset of the source rules, they are omitted. Since we only have naturals and all expressions operate on them, no error result is possible in the target.

4.3.2 Different Source and Target Values. In this instance, we extend the source language with expressions to perform Booleans and natural inputs, while the target only has expressions to input naturals. To compile the $\leq$, the target is also extended with a conditional that checks if an expression is less than another.

$e ::= \dots \mid \text{in-b} \mid \text{in-n}$ $i ::= n \mid b$ $s ::= \langle is, r \rangle$
 $e ::= \dots \mid \text{in-n} \mid \text{if } e \leq e \text{ then } e \text{ else } e$ $i ::= n$ $t ::= \langle is, r \rangle$

Source actions are Boolean b and natural inputs n and source traces s are lists of actions is together with a final result r . Target actions are just natural inputs n .

The source extensions respect typing and thus well-typed programs never produce error ($\emptyset$). The semantics of the extensions adds elements to the traces.

(Type-in-b) $\vdash \text{in-b} : B$	(Type-in-n) $\vdash \text{in-n} : N$	(Sem-in-nat) $\text{in-n} \rightsquigarrow \langle n \cdot \emptyset, n \rangle$	(Sem-in-bool) $\text{in-b} \rightsquigarrow \langle b \cdot \emptyset, b \rangle$
(Sem-itele) $\frac{e_1 \rightsquigarrow \langle is_1, n_1 \rangle \quad e_2 \rightsquigarrow \langle is_2, n_2 \rangle \quad n_1 \leq n_2?i = 3 : i = 4 \quad e_i \rightsquigarrow \langle is_i, n_i \rangle}{\text{if } e_1 \leq e_2 \text{ then } e_3 \text{ else } e_4 \rightsquigarrow \langle is_1 \cdot is_2 \cdot is_i, n_i \rangle}$			

The compiler is homomorphic, translating a source expression to the same target expression; the only differences are natural numbers (and conditionals).

$n \downarrow = n$ $\text{true} \downarrow = 1$ $e_1 + e_2 \downarrow = e_1 \downarrow + e_2 \downarrow$
 $\text{in-n} \downarrow = \text{in-n}$ $\text{false} \downarrow = 0$ $e_1 \leq e_2 \downarrow = \text{if } e_1 \downarrow \leq e_2 \downarrow \text{ then } 1 \text{ else } 0$
 $\text{in-b} \downarrow = \text{in-n}$ $e_1 \times e_2 \downarrow = e_1 \downarrow \times e_2 \downarrow$ $\text{if } e_1 \text{ then } e_2 \text{ else } e_3 \downarrow = \text{ifz } e_1 \downarrow \text{ then } e_3 \downarrow \text{ else } e_2 \downarrow$
 $e; e' \downarrow = e \downarrow; e' \downarrow$ $\text{ifz } e_1 \text{ then } e_2 \text{ else } e_3 \downarrow = \text{ifz } e_1 \downarrow \text{ then } e_2 \downarrow \text{ else } e_3 \downarrow$

When compiling an *if-then-else* the *then* and *else* branches of the source are swapped in the target because of the compilation of Booleans.

Relating Traces. We relate basic values (naturals and Booleans) in a non-injective fashion, as noted below. Then, we extend the relation to lists of inputs pointwise (Rules **Empty** and **Cons**) and

lift that relation to traces (Rules [Nat](#) and [Bool](#)).

$$\begin{array}{c}
 \begin{array}{c} n \sim n \\ \hline \text{(Empty)} \\ \emptyset \sim \emptyset \end{array} \quad \begin{array}{c} \text{true} \sim n \quad \text{if } n > 0 \\ \hline \text{(Cons)} \\ \begin{array}{c} i \sim i \quad is \sim is \\ i \cdot is \sim i \cdot is \end{array} \end{array} \quad \left| \quad \begin{array}{c} \text{is} \sim is \quad n \sim n \\ \hline \text{(Nat)} \\ \langle is, n \rangle \sim \langle is, n \rangle \end{array} \quad \begin{array}{c} \text{false} \sim 0 \\ \hline \text{(Bool)} \\ \begin{array}{c} is \sim is \quad b \sim n \\ \langle is, b \rangle \sim \langle is, n \rangle \end{array} \end{array}
 \end{array}$$

Property mappings. The property mappings $\tilde{\sigma}$ and $\tilde{\tau}$ induced by the trace relation $\sim$ defined above capture the intuition behind encoding Booleans as naturals:

- the source-to-target mapping allows **true** to be encoded by any non-zero number;
- the target-to-source mapping requires that **0** be replaceable by *both* **0** and **false**.

Compiler correctness. With the relation above, the compiler is proven to satisfy $CC^\sim$.

THEOREM 4.3 ($\cdot \downarrow$ IS CORRECT $\mathfrak{C}$). $\cdot \downarrow$ is $CC^\sim$.

Simulations with different traces. In the settings where $\text{Traces} = \text{Trace}_T$, it is customary to prove compiler correctness showing a forward simulation (i.e., a simulation between source and target transition system); then, using determinacy [24, 48] of the target language and input totality [25, 82] (receptiveness) of the source, this forward simulation is flipped into a backward simulation (a simulation between target and source transition system), as described by Beringer et al. [9], Leroy [42]. This “flipping” is useful, because forward simulations are often much easier to prove (by induction on the transitions of the source) than backward ones. For the proof of Theorem 4.3, we had to show a *backward* simulation, as it was not possible to define a forward one and then flip it. Hereafter, we show the reason lies in the shape of trace relation itself and discuss when is possible to generalize the flipping to the trace-relating setting.

We first give the main idea of the flipping proof, when the inputs are the same in the source and the target [9, 42]. We only consider inputs, as it is the most interesting case, since with determinacy, nondeterminism only occurs on inputs. Given a forward simulation $\mathcal{R}$, and a target program W_T that simulates a source program W_S , W_T is able to perform an input iff so is W_S ; otherwise, say, for instance that W_S performs an output, by forward simulation W_T would also perform an output, which is impossible because of determinacy. By input totality of the source, W_S must be able to perform the exact same input as W_T ; using forward simulation and determinacy, the resulting programs must be related.

The trace relation from Section 4.3.2 is not injective (both **0** and **false** are mapped to **0**), therefore, these arguments do not apply: Not all possible inputs of target programs are accounted for in the forward simulation. To flip a forward simulation into a backward one it is necessary that, for any source program W_S and target program W_T related by the forward simulation $\mathcal{R}$, the following diagram is satisfied:

We say that a forward simulation for which this property holds is *flippable*. For our example compiler, a flippable forward simulation works as follows: Whenever a Boolean input occurs in the source, the target program must perform every strictly positive input $\mathbf{n}$ (and not just $\mathbf{1}$, as suggested by the compiler). Using this property, determinacy of the target, input totality of the source, as well as the fact that any target input has an inverse image through the relation, we can indeed show that the forward simulation can be turned into a backward one: Starting from W_S R W_T and an input i_{T2} , we show that there is i_{S1} and i_{T2} as in the diagram above, using the same arguments as when the inputs are the same; because the simulation is flippable, we can close the diagram and obtain the existence of an adequate i_{S2} . From this, we obtain $CC^\sim$.

In fact, we showed that the flippable hypothesis is also sufficient to flip a forward simulation into a backward one, even in the trace-relating setting, and proved it in a general (i.e., language independent) “flipping theorem” (♣). We have also shown that if the relation $\sim$ defines a bijection between the inputs of the source and the target, then any forward simulation is flippable, hence reobtaining the usual proof technique [9, 42] as a special case.

4.4 Abstraction Mismatches

We now consider how to relate traces where a single source action is compiled to multiple target ones. To illustrate this, we extend our source language to output (nested) pairs of arbitrary size and our target language to send values that have a fixed size. Concretely, the source is analogous to the language of Section 4.3, except that it does not have inputs (nor Booleans for simplicity) but it has pairs. Additionally, it has an expression `send e` that can emit a (nested) pair e of values in a single action. Given that e reduces to a pair, e.g., $\langle v1, \langle v2, v3 \rangle \rangle$, expression `send  $\langle v1, \langle v2, v3 \rangle \rangle$` emits action $\langle v1, \langle v2, v3 \rangle \rangle$. That expression is eventually compiled into a sequence of individual sends in the target language `send v1 ; send v2 ; send v3`, since in the target, `send e` sends the value that e reduces to, but the language cannot send pairs (although it has pair constructs).

The source and target languages are formally extended (respectively, in the first and second lines below) with pairs and sending constructs as follows: For reasons that we explain when the compiler is presented, we extend the target language with a let-in construct and variables. Finally, source traces are sequences of sent values i (which include nested pairs) and target traces are only sequences of natural numbers.

$$\begin{aligned}
 e &:: \dots \mid \langle e, e \rangle \mid e.1 \mid e.2 \mid \text{send } e & ty &:: N \mid ty \times ty & i &:: n \mid \langle i, i \rangle & s &:: is \\
 e &:: \dots \mid \langle e, e \rangle \mid e.1 \mid e.2 \mid \text{let } x = e \text{ in } e \mid x \mid \text{send } e & i &:: n & t &:: is
 \end{aligned}$$

The source additions are well-typed and their semantics is unsurprising; the semantics relies on the usual capture-avoiding substitution $[r/x]$ of a result r for a variable x .

$$\begin{array}{c}
 \text{(Type-send)} \\
 \frac{\vdash e : \tau \times \tau'}{\vdash \text{send } e}
 \end{array}
 \quad
 \begin{array}{c}
 \text{(Type-pair)} \\
 \frac{\vdash e : \tau \quad \vdash e' : \tau'}{\vdash \langle e, e' \rangle : \tau \times \tau'}
 \end{array}
 \quad
 \begin{array}{c}
 \text{(Type-p1)} \\
 \frac{\vdash e : \tau \times \tau'}{\vdash e.1 : \tau}
 \end{array}
 \quad
 \begin{array}{c}
 \text{(Type-p2)} \\
 \frac{\vdash e : \tau \times \tau'}{\vdash e.2 : \tau'}
 \end{array}$$

$$\begin{array}{c}
 \text{(Type-send)} \\
 \frac{\vdash e : N}{\vdash \text{send } e}
 \end{array}$$

$$\begin{array}{c|c|c}
\begin{array}{c} \text{(Eval-P1)} \\ \hline e \rightsquigarrow \langle \text{is}, \langle r_1, r_2 \rangle \rangle \\ \hline e.1 \rightsquigarrow \langle \text{is}, r_1 \rangle \\ \hline \text{(Eval-Send)} \\ e \rightsquigarrow \langle \text{is}, r \rangle \\ \hline \text{send } e \rightsquigarrow \langle \text{is} \cdot r, r \rangle \end{array} &
\begin{array}{c} \text{(Eval-P2)} \\ \hline e \rightsquigarrow \langle \text{is}, \langle r_1, r_2 \rangle \rangle \\ \hline e.1 \rightsquigarrow \langle \text{is}, r_2 \rangle \\ \hline \text{(Sem)} \\ e \rightsquigarrow \langle \text{is}, r \rangle \\ \hline e \rightsquigarrow \text{is} \end{array} &
\begin{array}{c} \text{(Eval-Pair)} \\ \hline e \rightsquigarrow \langle \text{is}, r \rangle \quad e' \rightsquigarrow \langle \text{is}', r' \rangle \\ \hline \langle e, e' \rangle \rightsquigarrow \langle \text{is} \cdot \text{is}', \langle r, r' \rangle \rangle \\ \hline \text{(Eval-letin)} \\ e \rightsquigarrow \langle \text{is}, r \rangle \quad e'[r/x] \rightsquigarrow \langle \text{is}', r' \rangle \\ \hline \text{let } x = e \text{ in } e' \rightsquigarrow \langle \text{is} \cdot \text{is}', r' \rangle \end{array}
\end{array}$$

The compiler is defined inductively on the type derivation of a source expression ($\cdot \downarrow : \vdash e : \tau \rightarrow e$). The only interesting case is when compiling a **send** e , where we use the source type information concerning the message (i.e., a pair) being sent to deconstruct that pair into a sequence of natural numbers, which is what is sent in the target. This is the reason we need the let-in construct in the target, since we run the pair once (as the argument of the let-in) and then we send all of its projection to avoid duplicating side effects. Technically, since it is defined on the type derivations of terms, the compiler is defined inductively on type derivations (and not simply on terms). Thus, compiling $e; e'$ would look like the following (using D as a metavariable to range over derivations):

$$\left(\frac{\frac{D}{\vdash e} \quad \frac{D'}{\vdash e'}}{\vdash e; e'} \right) \downarrow = \left(\frac{D}{\vdash e} \right) \downarrow ; \left(\frac{D'}{\vdash e'} \right) \downarrow.$$

However, note that each judgment uniquely identifies which typing rule is being applied and the underlying derivation. Thus, for compactness, we only write the judgment in the compilation and implicitly apply the related typing rule to obtain the underlying judgments for recursive calls. To differentiate this from the compiler of Section 4.3.2, this compiler has parentheses over its input.

$$\begin{array}{ll}
(\vdash n : N) \downarrow = n & (\vdash e.1 : \tau) \downarrow = (\vdash e : \tau \times \tau') \downarrow.1 \\
(\vdash e \oplus e' : N) \downarrow = (\vdash e : N) \downarrow \oplus (\vdash e' : N) \downarrow & (\vdash e.2 : \tau') \downarrow = (\vdash e : \tau \times \tau') \downarrow.2 \\
(\vdash \langle e, e' \rangle : \tau \times \tau') \downarrow = \langle (\vdash e : \tau) \downarrow, (\vdash e' : \tau') \downarrow \rangle & \\
\left(\vdash \begin{array}{l} \text{if } e \\ \text{then } e \text{ else } e' \end{array} \right) \downarrow = \begin{array}{l} \text{if } (\vdash e : N) \downarrow \\ \text{then } (\vdash e) \downarrow \text{ else } (\vdash e') \downarrow \end{array} & (\vdash \text{send } e) \downarrow = \begin{array}{l} \text{let } x = (\vdash e : \tau \times \tau') \downarrow \\ \text{in gensend } (x, \tau \times \tau') \end{array} \\
\left(\vdash \begin{array}{l} \text{ifz } e \\ \text{then } e \text{ else } e' \end{array} \right) \downarrow = \begin{array}{l} \text{ifz } (\vdash e : N) \downarrow \\ \text{then } (\vdash e) \downarrow \text{ else } (\vdash e') \downarrow \end{array} & (\vdash e; e') \downarrow = (\vdash e) \downarrow ; (\vdash e') \downarrow
\end{array}$$

$$\text{gensend } (x, \tau) = \begin{cases} \text{send } x & \text{if } \tau = N \\ \text{gensend } (x, \tau').1; \text{gensend } (x, \tau'').2 & \text{if } \tau = \tau' \times \tau'' \end{cases}$$

Relating Traces. We start with the trivial relation between numbers: $n \sim^0 n$, i.e., numbers are related when they are the same. We cannot build a relation between single actions, since a single source action is related to multiple target ones. Therefore, we define a relation between a source action i and a target trace t (a list of numbers) inductively on the structure of i .

$$\begin{array}{c|c|c|c}
\begin{array}{c} \text{(Trace-Rel-N-N)} \\ \hline n \sim^0 n \quad n' \sim^0 n' \\ \hline \langle n, n' \rangle \sim n \cdot n' \end{array} &
\begin{array}{c} \text{(Trace-Rel-N-M)} \\ \hline n \sim^0 n \quad i \sim t \\ \hline \langle n, i \rangle \sim n \cdot t \end{array} &
\begin{array}{c} \text{(Trace-Rel-M-N)} \\ \hline i \sim t \quad n \sim^0 n \\ \hline \langle i, n \rangle \sim t \cdot n \end{array} &
\begin{array}{c} \text{(Trace-Rel-M-M)} \\ \hline i \sim t \quad i' \sim t' \\ \hline \langle i, i' \rangle \sim t \cdot t' \end{array}
\end{array}$$

A pair of naturals is related to the two actions that send each element of the pair (Rule **Trace-Rel-N-N**). If a pair is made of sub-pairs, then we require all such sub-pairs to be related (Rules **Trace-Rel-N-M** to **Trace-Rel-M-M**).

We build on these rules to define the $s \sim t$ relation between source and target traces for which the compiler is correct (Theorem 4.5). Trivially, traces are related when they are both empty. Alternatively, given related traces, we can concatenate a source action and a second target trace provided that they are related (Rule [Trace-Rel-Single](#)). Before proving that the compiler is correct, we need Lemma 4.4. Intuitively, that lemma tells us that the way we break down a source sent value r into multiple target sends is correct.

$$\begin{array}{c} \text{(Trace-Rel-Single)} \\ \frac{s \sim t \quad i \sim t'}{s \cdot i \sim t \cdot t'} \end{array}$$

LEMMA 4.4 (gensend $(\cdot, \cdot)$ WORKS). *if gensend $(x, \tau \times \tau')[(\vdash r : \tau \times \tau')\downarrow/x] \rightsquigarrow t$ then $r \sim t$ (since r is necessarily a sent value i , that can be related to t).*

THEOREM 4.5 $((\cdot)\downarrow$ IS CORRECT). $(\cdot)\downarrow$ is $CC^\sim$.

With our trace relation, the trace property mappings capture the following intuitions:

- The target-to-source mapping states that a source property can reconstruct target action as it sees fit. For example, trace $4 \cdot 6 \cdot 5 \cdot 7$ is related to $\langle 4, 6 \rangle \cdot \langle 5, 7 \rangle$ and $\langle \langle 4, 6 \rangle, \langle 5, 7 \rangle \rangle$ (and many more variations). This gives freedom to the source implementation of a target behavior, which follows from the non-injectivity of $\sim$.¹⁰
- The source-to-target mapping “forgets” about the way pairs are nested, but is faithful w.r.t. the values v_i contained in a message. Notice that source safety properties are always mapped to target safety properties. For instance, if $\pi_S \in \text{Safety}_S$ prescribes that some bad number is never sent, then $\tilde{\tau}(\pi_S)$ prescribes the same number is never sent in the target and $\tilde{\tau}(\pi_S) \in \text{Safety}_T$. Of course if $\pi_S \in \text{Safety}_S$ prescribes that a particular nested pairing like $\langle 4, \langle 6, \langle 5, 7 \rangle \rangle \rangle$ never happens, then $\tilde{\tau}(\pi_S)$ is still a target safety property, but the trivial one, since $\tilde{\tau}(\pi_S) = \top \in \text{Safety}_T$.

5 TRACE-RELATING COMPILATION AND NONINTERFERENCE PRESERVATION

We now study the relation between trace-relating compilation and noninterference preservation. As mentioned earlier (Section 3.1), in the particular case where source and target observations are drawn from the same set, a correct compiler ($CC^\sim$) is enough to ensure the preservation of all subset-closed hyperproperties, in particular of *noninterference* (NI) [28]. But in the scenario where target observations are strictly more informative than source observations, this is not the case. In fact, as we will show, the best guarantee one may expect from a correct trace-relating compiler ($CC^\sim$) in such a setting is a *weakening* (or *declassification*) of target noninterference that matches the noninterference property satisfied in the source. In certain scenarios, it turns out that the noninterference property of interest in the target comes “for free,” while in others, it does not, and therefore establishing noninterference requires an additional proof effort beyond $CC^\sim$. To formalize this reasoning, this section applies the trinitarian view of trace-relating compilation to the general framework of **abstract noninterference** (ANI) [27], clarifying the kind of noninterference preservation that follows from a given trace relation and correct compilation.

We first define NI and explain the issue of preserving source NI via a $CC^\sim$ compiler (Section 5.1). We then introduce ANI, which allows characterizing various forms of noninterference (Section 5.2), and formulate a theory of ANI preservation via $CC^\sim$, both with respect to a *timing insensitive* declassification (Section 5.3) and in general (Section 5.4). We also study how to deal with cases such as undefined behavior in the target (Section 5.5). We then answer the dual question, i.e., which source NI should be satisfied to guarantee that compiled programs are noninterfering with respect to target observers (Section 5.6). Finally, we use this formal development to analyze recent work

¹⁰Making $\sim$ injective is a matter of adding open and close parentheses actions in target traces.

on correct compilers with interesting noninterference guarantees [7, 74], clarifying whether these guarantees follow from correctness alone or not (Section 5.7).

5.1 Noninterference and Trace-relating Compilation

Intuitively, **noninterference (NI)** requires that publicly observable outputs do not reveal information about private inputs. To define this formally, we need a few additions to our setup. We indicate the (disjoint) *input* and *output* projections of a trace t as t° and $t^\star$, respectively.¹¹ Denote with $[t]_{\text{low}}$ the equivalence class of a trace t , obtained using a standard low-equivalence relation that relates low (public) events only if they are equal, and ignores any difference between private events. Then, NI for source traces can be defined as:

$$\text{NI}_S = \{ \pi_S \mid \forall s_1 s_2 \in \pi_S. [s_1^\circ]_{\text{low}} = [s_2^\circ]_{\text{low}} \Rightarrow [s_1^\star]_{\text{low}} = [s_2^\star]_{\text{low}} \}.$$

That is, source NI comprises the sets of traces that have equivalent low output projections as long as their low input projections are equivalent.

When additional observations are possible in the target, it is unclear whether a noninterfering source program is compiled to a noninterfering target program or not, and if so, whether the notion of NI in the target is the expected (or desired) one. We illustrate this issue by considering a scenario where target traces extend source traces by exposing the execution time. While source noninterference NI_S requires that private inputs do not affect public outputs, NI_T additionally requires that the execution time is not affected by varying private inputs.

To model the scenario described, we represent target traces as pairs of a source trace and a natural number that denotes the time spent to produce the trace (using ω for infinite time units). Formally, if Traces_S denotes the set of source traces, then $\text{Trace}_T = \text{Traces}_S \times \mathbb{N}^\omega$ is the set of target traces, where $\mathbb{N}^\omega \triangleq \mathbb{N} \cup \{\omega\}$.

Notice that if two source traces s_1, s_2 are low-equivalent, then $\{s_1, s_2\} \in \text{NI}_S$ and $\{(s_1, 42), (s_1, 42)\} \in \text{NI}_T$, but $\{(s_1, 42), (s_2, 43)\} \notin \text{NI}_T$ and $\{(s_1, 42), (s_2, 42), (s_1, 43), (s_2, 43)\} \notin \text{NI}_T$.

Consider the following straightforward trace relation, which relates a source trace to any target trace whose first component is equal to it, irrespective of execution time:

$$s \sim t \equiv \exists n. t = (s, n).$$

A compiler is $\text{CC}^\sim$ for this trace relation if any trace that can be exhibited in the target can be simulated in the source in some amount of time. For such a compiler, Theorem 3.3 says that if W satisfies NI_S , then $W \downarrow$ satisfies $Cl_\subseteq \circ \tilde{\tau}(\text{NI}_S)$. This hyperproperty is, however, strictly weaker than NI_T , as it contains for example $\{(s_1, 42), (s_2, 42), (s_1, 43), (s_2, 43)\}$, and one cannot conclude that $W \downarrow$ is noninterfering in the target. It is easy to check that

$$Cl_\subseteq \circ \tilde{\tau}(\text{NI}_S) = Cl_\subseteq (\{ \pi_S \times \mathbb{N}^\omega \mid \pi_S \in \text{NI}_S \}) = \{ \pi_S \times I \mid \pi_S \in \text{NI}_S \wedge I \subseteq \mathbb{N}^\omega \},$$

the first equality coming from $\tilde{\tau}(\pi_S) = \pi_S \times \mathbb{N}^\omega$, and the second from NI_S being subset-closed. As we will see, this hyperproperty *can* be characterized as a form of NI, which one might call *timing-insensitive noninterference*, i.e., ensured only against attackers that cannot measure execution time. For this characterization, and to describe different forms of noninterference as well as formally analyze their preservation by a $\text{CC}^\sim$ compiler, we rely on the general framework of *abstract noninterference* [27].

¹¹The exact shape of inputs and outputs depends on the scenario. For instance, inputs can be initial memories and outputs trace semantics of programs as in Reference [27, Section 7], while for interactive programs one may want to consider streams like Clark and Hunt [17]. We only require the sets of input and output projections to be disjoint. Further information, such as the ordering of events, is part of the attacker/observer model or the declassification of noninterference itself.

5.2 Abstract Noninterference

ANI [27] is a generalization of NI whose formulation relies on *abstractions* (in the sense of Abstract Interpretation [20]) to encompass arbitrary variants of NI. ANI is parameterized by an *observer abstraction* ρ , which denotes the distinguishing power of the attacker, and a *selection abstraction* ϕ , which specifies when to check NI, and therefore captures a form of declassification [69].¹² Formally:

$$ANI_{\phi}^{\rho} = \{ \pi \mid \forall t_1 t_2 \in \pi. \phi(t_1^{\circ}) = \phi(t_2^{\circ}) \Rightarrow \rho(t_1^{\bullet}) = \rho(t_2^{\bullet}) \}.$$

By picking $\phi = \rho = [\cdot]_{low}$, we recover the standard noninterference defined above, where NI must hold for all low inputs (i.e., no declassification of private inputs), and the observational power of the attacker is limited to distinguishing low outputs. The observational power of the attacker can be weakened by choosing a more liberal relation for ρ . For instance, one may limit the attacker to observe the *parity* of output integer values. Another way to weaken ANI is to use ϕ to specify that noninterference is only required to hold for a subset of low inputs.

The operators ϕ and ρ are defined over sets of (input and output projections of) traces, explicitly $\phi : 2^{Trace^{\circ}} \rightarrow 2^{Trace^{\circ}}$ and $\rho : 2^{Trace^{\bullet}} \rightarrow 2^{Trace^{\bullet}}$. When we write $\phi(t)$ like above, this should be understood as a convenience notation for $\phi(\{t\})$. Likewise, $\phi = [\cdot]_{low}$ should be understood as $\phi = \lambda\pi. \bigcup_{t \in \pi} [t]_{low}$, i.e., the powerset lifting of $[\cdot]_{low}$. Additionally, ϕ and ρ are required to be upper-closed operators (*uco*)—i.e., monotonic, idempotent, and extensive (i.e., $\forall \pi^*. \pi^* \subseteq \rho(\pi^*)$)—on the poset that is the powerset of (input and output projections of) traces ordered by inclusion [27].

5.3 Trace-relating Compilation and ANI for Timing

We can now reformulate our example with observable execution times in target traces in terms of ANI. We have $NI_S = ANI_{\phi}^{\rho}$ with $\phi = \rho = [\cdot]_{low}$. In this case, the hyperproperty that a compiled program $W \downarrow$ satisfies whenever W satisfies NI_S can be described as an instance of ANI:

$$Cl_{\subseteq} \circ \tilde{\tau}(NI_S) = ANI_{\phi}^{\rho}$$

for $\phi = \phi$ and $\rho(\pi) = \{ (s, n) \mid \exists (s_1, n_1) \in \pi. [s^*]_{low} = [s_1^*]_{low} \}.$

The definition of ϕ tells us that the trace relation does not affect the selection abstraction, i.e., declassification is unaffected. The definition of ρ characterizes an observer that cannot distinguish execution times for noninterfering traces (notice that n_1 in the definition of ρ is discarded). For instance, $\rho(\{(s, n_1)\}) = \rho(\{(s, n_2)\})$, for any s, n_1, n_2 . Therefore, in this setting, we know explicitly through ρ that a $CC^{\sim}$ compiler degrades source noninterference to target *timing-insensitive* noninterference.

5.4 Trace-relating Compilation and ANI in General

While the particular ϕ and ρ above can be discovered by intuition, we want to know whether there is a systematic way of obtaining them in general. In other words, for *any* trace relation $\sim$ and *any* notion of source NI, what property is guaranteed on noninterfering source programs by any $CC^{\sim}$ compiler?

We can now answer this question generally (Theorem 5.1): Any source notion of noninterference expressible as an instance of ANI is mapped to a corresponding instance of ANI in the target, whenever source traces are an abstraction of target ones (i.e., when $\sim$ is a total and surjective map). For this result, we consider trace relations that can be split into input and output trace relations

¹²To be precise, the original formulation of ANI by Giacobazzi and Mastroeni [27] includes a third parameter η , which describes the maximal input variation that the attacker may control. Here, we omit η (i.e., take it to be the identity) to simplify the presentation.

(denoted as $\sim \triangleq \langle \sim, \sim \rangle$) such that $s \sim t \iff s^\circ \sim t^\circ \wedge s^\bullet \sim t^\bullet$. The trace relation $\sim$ corresponds to a Galois connection between the sets of trace properties $\tilde{\tau} \sqsubseteq \tilde{\sigma}$ as described in Section 2.2. Similarly, the pair $\sim$ and $\sim$ corresponds to a pair of Galois connections, $\tilde{\tau}^\circ \sqsubseteq \tilde{\sigma}^\circ$ and $\tilde{\tau}^\bullet \sqsubseteq \tilde{\sigma}^\bullet$, between the sets of input and output properties. In the timing example, time is an output so we have $\sim \triangleq \langle =, \sim \rangle$ and $\sim$ is defined as $s^\bullet \sim t^\bullet \equiv \exists n. t^\bullet = (s^\bullet, n)$.

THEOREM 5.1 (COMPILING ANI). *Assume traces of source and target languages are related via $\sim \subseteq \text{Traces}_S \times \text{Traces}_T$, $\sim \triangleq \langle \sim, \sim \rangle$ such that $\sim$ and $\sim$ are both total maps from target to source traces, and $\sim$ is surjective. Assume $\downarrow$ is a $CC^\sim$ compiler, and $\phi \in \text{uco}(2^{\text{Trace}_S^\circ})$, $\rho \in \text{uco}(2^{\text{Trace}_S^\bullet})$.*

If $\mathcal{W}$ satisfies ANI_ϕ^ρ , then $\mathcal{W}\downarrow$ satisfies $\text{ANI}_{\phi^\#}^{\rho^\#}$, where $\phi^\#$ and $\rho^\#$ are defined as:

$$\begin{aligned} \phi^\# &= g^\circ \circ \phi \circ f^\circ & \rho^\# &= g^\bullet \circ \rho \circ f^\bullet \\ f^\circ(\pi^\circ) &= \{s^\circ \mid \exists t^\circ \in \pi^\circ. s^\circ \sim t^\circ\} & g^\circ(\pi_S^\circ) &= \{t^\circ \mid \forall s^\circ. s^\circ \sim t^\circ \Rightarrow s^\circ \in \pi_S^\circ\} \end{aligned}$$

(and both $f^\bullet$ and $g^\bullet$ are defined analogously).

Moreover, we can prove that if $\sim$ is surjective, then $\text{ANI}_{\phi^\#}^{\rho^\#} \subseteq \text{Cl}_\subseteq \circ \tilde{\tau}(\text{ANI}_\phi^\rho)$. Therefore, the derived guarantee $\text{ANI}_{\phi^\#}^{\rho^\#}$ is at least as strong as the hyperproperty (a priori different from some noninterference) that follows by just knowing that the compiler $\downarrow$ is $CC^\sim$.

The target abstract noninterference has to be intended as the *best correct approximation* of the source one. The mappings $f^\circ \sqsubseteq g^\circ$ are the existential and universal images of the relation $\sim_{\text{swap}} \subseteq \text{Trace}_T \times \text{Trace}_S$, defined by $t^\circ \sim_{\text{swap}} s^\circ$ if and only if $s^\circ \sim t^\circ$. Therefore, f° and g° are lower and upper adjoints, respectively (Section 2). The operator $\phi^\#$ is the best correct approximation of ϕ w.r.t. to $f^\circ \sqsubseteq g^\circ$ [20] (hence, the choice of the $(_)^\#$ notation). A similar result holds for $\rho^\#$.

Coming back to our example above, we can formally recover the intuitively justified definitions, i.e., $\phi^\# = g^\circ \circ \phi \circ f^\circ = \phi$ and $\rho^\# = g^\bullet \circ \rho \circ f^\bullet = \rho$.

5.5 Noninterference and Undefined Behavior

As stated above, Theorem 5.1 does not apply to several scenarios from Section 4 such as undefined behavior (Section 4.1). Indeed, in these cases, the relation $\sim$ is not a total map. Nevertheless, we can still exploit our framework to reason about the impact of compilation on noninterference.

Let us consider $\sim \triangleq \langle \sim, \sim \rangle$ where $\sim$ is any total and surjective map from target to source inputs (e.g., equality) and $\sim$ is defined as $s^\bullet \sim t^\bullet \equiv s^\bullet = t^\bullet \vee \exists m^\bullet \leq t^\bullet. s^\bullet = m^\bullet$. *Wrong*. Intuitively, a $CC^\sim$ compiler guarantees noninterference for the compiled program, provided that the target attacker cannot exploit undefined behavior to learn private information. This intuition can be made formal by the following theorem:

THEOREM 5.2 (RELAXED COMPILING ANI). *Relax the assumptions of Theorem 5.1 by allowing $\sim$ to be any output trace relation. If $\mathcal{W}$ satisfies ANI_ϕ^ρ , then $\mathcal{W}\downarrow$ satisfies $\text{ANI}_{\phi^\#}^{\rho^\#}$ where $\phi^\#$ is defined as in Theorem 5.1, and $\rho^\#$ is such that:*

$$\forall s t. s^\bullet \sim t^\bullet \Rightarrow \rho^\#(t^\bullet) = \rho^\#(\tilde{\tau}^\bullet(\rho(s^\bullet))). \quad (1)$$

Technically, instead of giving us a *definition* of $\rho^\#$, the theorem gives a *property* of it. The property states that, given a target output trace $t^\bullet$, the attacker cannot distinguish it from any other target output traces produced by other possible compilations ($\tilde{\tau}^\bullet$) of the source trace s it relates to, up to the observational power of the source-level attacker ρ . Therefore, given a source attacker ρ , the theorem characterizes a *family* of attackers that cannot observe any interference for a correctly compiled noninterfering program. Notice that the target attacker $\rho^\top \triangleq \lambda_. \top$ satisfies the premise

of the theorem, but defines a trivial hyperproperty, so we cannot prove in general that $ANI_{\phi^\#}^{\rho^\#} \subseteq Cl_{\subseteq} \circ \tilde{\tau}(ANI_{\phi}^{\rho})$. Also, this degenerate attacker $\rho^\top$ shows that the family of attackers described in Theorem 5.2 is nonempty, which ensures the existence of a most powerful attacker among them [27].

5.6 From Target NI to Source NI

We now explore the dual question: Under what hypothesis does trace-relating compiler correctness alone allow target noninterference to be reduced to source noninterference? This is of practical interest, as one would be able to protect from target attackers by ensuring noninterference in the source. This task can be made easier if the source language has some static enforcement mechanism [1, 44].

Let us consider the languages from Section 4.4 extended with the ability to accept inputs as (pairs of) values. It is easy to show that the compiler described in Section 4.4 (extended to treat the new input expressions homomorphically) is still $CC^\sim$: Given a target trace t with the same inputs of the source one (i.e., $s^\circ = t^\circ$), the compiler of Section 4.4 ensures that t simulates the same outputs of s (i.e., $s^\circ \sim t^\circ$). Assume that we want to satisfy a given notion of target noninterference after compilation, i.e., $W \downarrow \models ANI_{\phi}^{\rho}$. Recall that the observational power of the target attacker, ρ , is expressed as a property of sequences of values. To express the same property (or attacker) in the source, we have to abstract the way pairs of values are nested. For instance, the source attacker should not distinguish $\langle v_1, \langle v_2, v_3 \rangle \rangle$ and $\langle \langle v_1, v_2 \rangle, v_3 \rangle$. In general (i.e., when $\sim$ is not the identity), this argument is valid only when ϕ can be represented in the source. More precisely, ϕ must consider as equivalent all target inputs that are related to the same source input, because in the source it is not possible to have a finer distinction of inputs. This intuitive correspondence can be formalized as follows:

THEOREM 5.3 (TARGET ANI BY SOURCE ANI). *Let $\phi \in uco(2^{\text{Trace}_T^\circ})$, $\rho \in uco(2^{\text{Trace}_T^\circ})$ and $\sim$ a total and surjective map from source outputs to target ones and assume that*

$$\forall s \ t. s^\circ \sim t^\circ \Rightarrow \phi(t^\circ) = \phi(\tilde{\tau}^\circ(s^\circ)).$$

If $\cdot \downarrow$ is a $CC^\sim$ compiler and W satisfies $ANI_{\phi^\#}^{\rho^\#}$, then $W \downarrow$ satisfies ANI_{ϕ}^{ρ} for

$$\phi^\# = \tilde{\sigma}^\circ \circ \phi \circ \tilde{\tau}^\circ \qquad \rho^\# = \tilde{\sigma}^\circ \circ \rho \circ \tilde{\tau}^\circ.$$

5.7 Analyzing Noninterference Preserving Compilers

The results presented in this section formalize and generalize some intuitive facts about compiler correctness and noninterference, clarifying which noninterference property follows “for free” from trace-relating compiler correctness. Of course, in the general case, compiler correctness alone is not a strong enough criterion for dealing with many security properties [8, 23]. This section exploits our ANI-based framework and results to analyze two compilers from the recent literature [7, 74] that are both proven to be correct and to preserve two interesting notions of noninterference: cryptographic constant time (Section 5.7.1) and value-dependent noninterference (Section 5.7.2). For each, we explain how to express compiler correctness as an instance of $CC^\sim$, describe the noninterference property that is implied by the trace relation and the correctness result, and compare it with the noninterference properties of interest as established by their authors.

5.7.1 A Correct Compiler Preserving Cryptographic Constant Time. Barthe et al. [7] provide a correct compiler (as an extension of CompCert) that also preserves cryptographic **constant time** (CT). CT is a security property stating that the runtime of a program does not depend on its secret,

and thus an attacker cannot extrude secrets of a program by observing its execution time. A CT-preserving compiler takes code that is CT and generates code that also is CT. Thus, a CT-preserving compiler must translate runtime-equivalent source programs into runtime-equivalent target ones. Notice that it is not necessary for the leakage of target programs to be the same of their source counterparts, rather: Source programs with the same leakage must be compiled to target programs with the same leakage.

Barthe et al. [7] prove CT preservation for 17 passes of CompCert. The authors partition the 17 steps in four categories, depending on the proof technique they use to show CT preservation. Every category proves an instance of $CC^\sim$ by improving on the existing CompCert simulation. In three out of the four cases this is sufficient to also prove CT preservation, while for the last category a further proof is necessary. In what follows, we first encode CT as an instance of abstract noninterference, i.e., show for which operators $CT = ANI_{\phi_{CT}}^{\rho_{CT}}$ and then use our framework to understand why modifying CompCert simulation is sufficient in the first three categories but not in the last one. For each category, Theorem 5.2 applies, so no ρ that respects Equation (1) can notice any interference on compiled programs that were source constant-time. In the first three categories the attacker that defines CT — ρ_{CT} —respects the equation,¹³ i.e.,

$$\forall s^* \cdot t^* . s^* \sim t^* \Rightarrow \rho_{CT}(t^*) = \rho_{CT}(\tilde{t}^*(\rho_{CT}(s^*))), \quad (2)$$

and CT preservation is therefore a consequence of $CC^\sim$. In the last category, ρ_{CT} does not respect Equation (2) and the authors have to prove an additional theorem, the *CT-diagram*.

Trace Model and CT as an instance of ANI. The formal definition of CT is given by extending the semantics of the languages in CompCert and enriching the traces of input and output events with leakages. Leakages are results of execution steps that involve conditional branching or memory access. A program is CT w.r.t. a certain relation over program states φ [7, Definition 3.2] iff for every two initial states i, i' such that $\varphi(i, i')$, the leakages that can be observed are the same. Notice that in Reference [7, Definition 3.2] the secret is stored in the program states and defined by φ , therefore to regard CT as an instance of abstract noninterference program states will be regarded as inputs and events together with their leakages as outputs. More precisely, a trace t is a sequence of of triples (i, e, j) where i and j are program states and e an event in the instrumented semantics, i.e., input/output event and associated leakage.

We consider:

- ϕ_{CT} to be (the uco corresponding to) the relation defined by $t_1^* \phi_{CT} t_2^*$ iff t_1^*, t_2^* have the same length with $t_1^* = (i_0, i_1), (i_1, i_2), \dots, t_2^* = (j_0, j_1), (j_1, j_2), \dots$ and $\forall n. \varphi(i_n, j_n)$.
- ρ_{CT} to be (the uco corresponding to) the relation defined by $t_1^* \rho_{CT} t_2^*$ iff t_1^*, t_2^* have the same length with $t_1^* = e_0, e_1, \dots, t_2^* = f_0, f_1, \dots$ and $\forall n. leak(e_n) = leak(f_n)$, where $leak(e)$ denotes the leakage in the event e (projection of e on the leak-only semantics [7]).

It is easy to check that $CT = ANI_{\phi_{CT}}^{\rho_{CT}}$ for the ϕ_{CT} and ρ_{CT} given above.

We now present more details for each of the four proof techniques adopted by Barthe et al. [7]. Since CT is defined only for *safe* programs [7, Definition 3.1], we can assume no undefined behavior is ever encountered and have a simpler presentation. We also omit $\phi^\#$ coming from the application of Theorem 5.2, as it always coincides with ϕ_{CT} .

Constant-time security preservation by leakage preservation (Barthe et al. [7, Section 5.2]). For compilation passes that belong to this category, the authors prove that the source

¹³In each compilation step, source and target traces are drawn from the same set so ρ_{CT} can be applied to both source and target traces.

leakage is preserved exactly in the target. Thus, in this simple case, the theorem proved is $CC^\sim$ where $\sim$ is point-wise equality of events together with leakages, $\tilde{\tau}^\bullet$ the identity and ρ_{CT} satisfies Equation 2 by idempotency of ρ_{CT} ,

$$\begin{aligned} \rho_{CT}(\tilde{\tau}^\bullet(\rho_{CT}(s^\bullet))) &= [s^\bullet \sim t^\bullet \Rightarrow s^\bullet = t^\bullet] \\ \rho_{CT}(\tilde{\tau}^\bullet(\rho_{CT}(t^\bullet))) &= [\tilde{\tau}^\bullet = \lambda x.x] \\ \rho_{CT}(\rho_{CT}(t^\bullet)) &= [\rho_{CT} \text{ idempotent}] \\ \rho_{CT}(t^\bullet). \end{aligned}$$

CT preservation from leakage-erasing simulation (Barthe et al. [7, Section 5.3]). In this case, $CC^\sim$ is proved for a relation that erases source leakage-only events, i.e., those events that do not contain inputs or outputs, but only the amount of leakage revealed. More precisely (see also Reference [7, Fig. 8]) for $s^\bullet = e_0, e_1, \dots$ and $t^\bullet = e_0, e_1, \dots$ of the same length, $s^\bullet \sim t^\bullet$ iff

$$\forall k, e_k = e_k \vee (e_k = \epsilon \wedge e_k \text{ is leak only}).$$

The property mapping associated to the above relation, $\tilde{\tau}^\bullet$, erases all leak-only events from the traces of a source property. If an attacker cannot notice at any point any difference in the leakages of two traces and we erase the leak-only events from them, then the attacker will still not notice any difference on leakages, therefore it is easy to check that Equation 2 holds also in this case.

CT preservation via memory injection (Barthe et al. [7, Section 5.4]). This case is analogous to the one above, save that it rests on a more complex relation $\sim$ involving a *memory injection* relation (see Barthe et al. [7, Definition 5.8]). Intuitively, $\sim$ relates source and target traces that differ at most in leakages due to memory accesses. While in the previous case, leakages were simply erased, here they are modified and crucially with some uniformity. Reasoning as in the previous case, if an attacker cannot notice a difference in the leakages of two traces and we modify equal leakages of the same factor, then the attacker will still not notice any difference on leakages, thus Equation 2 holds.

CT preservation from CT-diagram (Barthe et al. [7, Section 5.5]). In this case, ρ_{CT} does not satisfy Equation 2 because the *counting simulation* ([7, Definition 5.10]) does not necessarily relate source and target leakages but only the inputs and outputs.¹⁴ $CC^\sim$ alone does not ensure that an attacker cannot observe any interference in the target leakages, to show preservation of CT the authors need to prove an extra condition, the so-called CT diagram [8].

5.7.2 Value-dependent noninterference. Sison and Murray [74] introduce a compiler that provably preserves **value-dependent noninterference (VDNI)** for a concurrent language with shared variables. *Value-dependent* means that the secrecy level of a variable—*low* or *high*—may depend on the value of some other variable, called the control variable of the first, and therefore could change throughout its lifetime.

Preservation of VDNI for concurrent programs enjoys *compositionality*, meaning that it follows from the preservation of VDNI for each single thread [52] under certain conditions. As the compositionality result is orthogonal to our framework, we can study either (1) the preservation of VDNI for one local thread or for (2) the whole-program,

In the remainder of this section, we focus on the preservation of VDNI for a single thread, which is proven by showing a *secure refinement* relation between source and compiled threads. Similarly to the previous section, the secure refinement is expressed via a cube diagram (Reference [74], Figure 1) and can be proven directly [52] or split into more obligations [74].

¹⁴The interested reader will notice the difference from the previous category by comparing condition (1) of Definition 5.10 and condition (1) of Definition 5.8 by Barthe et al. [7].

As Sison and Murray [74] use a state transition-based semantics, we first show how to encode this semantics into a trace model by defining the $\sim$ relation based on the secure refinement relation. We then show how to encode VDNI as an instance of abstract noninterference (i.e., both $\text{VDNI}_S = \text{ANI}_\phi^\rho$ and $\text{VDNI}_T = \text{ANI}_\phi^\rho$). Finally, we apply Theorem 5.2 and conclude that if $\mathcal{W}$ satisfies VDNI_S , then $\mathcal{W} \downarrow$ satisfies VDNI_T given that the trace relation $\sim$ has properties defined in Reference [52, Theorem 5.1].

Source (WHILE) and target (RISC-like assembly) languages are equipped with a determined evaluation step semantics (i.e., a semantics where the only source of nondeterminism are external inputs; Reference [74], Section 2) between thread-local configurations, which are triples of the form $\langle \text{tps}, \text{mds}, \text{mem} \rangle$. In such a configuration, mds is the access mode state for program variables and mem is a map relating global program variables to their values. Both of these components are common to the source and target language. The tps component denotes the thread-private state. In the source language, it is the program to be executed. In the target language, tps consists of the target program (labelled assembly-language instructions), of a program counter, and of the set of thread-local registers. We denote WHILE configurations by tuples of the form: $\langle \text{tps}, \text{mds}, \text{mem} \rangle$ and RISC configurations by tuples of the form: $\langle \text{tps}, \text{mds}, \text{mem} \rangle$.

Trace Model and Trace relation. We consider traces that are (possibly infinite) sequences of configurations. The traces produced by a program are the sequences of local configurations that the program may encounter during execution, according to the evaluation semantics. Let $s = \langle \text{tps}_1, \text{mds}_1, \text{mem}_1 \rangle, \langle \text{tps}_2, \text{mds}_2, \text{mem}_2 \rangle, \dots$ be a source trace. The input projection is defined by $s^\circ = \langle \text{mds}_1, \text{mem}_1 \rangle$ (the tuple consisting of the access modes and the memory in the first state) and the output projection is defined by $s^\bullet = s$ (the trace itself). Input/output projections are defined similarly for target traces.

We take the trace relation $\sim \subseteq \text{Trace}_S \times \text{Trace}_T$ to be the point-wise lifting of a secure refinement relation $\mathcal{R}$ (Reference [74], Definition 6). Source and target configurations $\langle \text{tps}, \text{mds}, \text{mem} \rangle \mathcal{R} \langle \text{tps}', \text{mds}', \text{mem}' \rangle$ that are related coincide on the access mode and memory part (i.e., $\text{mds} = \text{mds}'$ and $\text{mem} = \text{mem}'$; Reference [74], Definition 4), so $\sim$ is simply the identity and $\sim$ coincides with $\sim$.

VDNI as abstract noninterference. A program satisfies VDNI (Reference [74], Definition 2) if any two of its executions starting in low equivalent memories are related via a *strong low bisimulation modulo modes* (strong low bisimulation mm). Intuitively, a strong low bisimulation mm is a bisimulation that preserves low-equivalence. Preservation of VDNI is proved by Murray et al. [52] by showing that for every strong low-bisimulation mm $\mathcal{B}$ for source threads, there exists a target strong low bisimulation mm $\mathcal{B}$ such that if two source threads are related by $\mathcal{B}$, then the compiled threads are related by $\mathcal{B}$ (Reference [52], Theorem 5.1).

The intuition for the encoding of VDNI as an instance of abstract noninterference is to model low equivalence through the operator ϕ , and bisimilarity through ρ . More rigorously, $\text{VDNI}_S = \text{ANI}_\phi^\rho$, where ϕ and ρ are defined as following:

For $s^\circ = \langle \text{mds}_1, \text{mem}_1 \rangle$,

$$\phi(s^\circ) = \left\{ \langle \text{mds}_1, \text{mem}'_1 \rangle \mid \text{mem}_1 =_{\text{mds}_1}^{\text{Low}} \text{mem}'_1 \right\},$$

where $=_{\text{mds}}^{\text{Low}}$ is the low-equivalence modulo mds (Reference [74], Definition 1).

For $s^\bullet = \langle \text{tps}_1, \text{mds}_1, \text{mem}_1 \rangle, \langle \text{tps}_2, \text{mds}_2, \text{mem}_2 \rangle, \dots$,

$$\rho(s^\bullet) = \{ \langle \text{tps}'_1, \text{mds}'_1, \text{mem}'_1 \rangle, \langle \text{tps}'_2, \text{mds}'_2, \text{mem}'_2 \rangle, \dots \mid \\ \forall i. \exists \mathcal{B}_i. (\langle \text{tps}_i, \text{mds}_i, \text{mem}_i \rangle, \langle \text{tps}'_i, \text{mds}'_i, \text{mem}'_i \rangle) \in \mathcal{B}_i \},$$

where $\mathcal{B}_i$ denotes a strong low bisimulation modulo modes. Similarly $\text{VDNI}_{\mathbf{T}} = \text{ANI}_{\phi}^{\rho}$ where

$$\begin{aligned} \phi(\mathbf{t}^*) &= \{ \langle \text{mds}_1, \text{mem}'_1 \rangle \mid \text{mem}_1 =_{\text{mds}_1}^{\text{Low}} \text{mem}'_1 \}, \\ \rho(\mathbf{t}^*) &= \{ \langle \text{tps}'_1, \text{mds}'_1, \text{mem}'_1 \rangle, \langle \text{tps}'_2, \text{mds}'_2, \text{mem}'_2 \rangle, \dots \mid \\ &\quad \forall i. \exists \mathcal{B}_i. (\langle \text{tps}_i, \text{mds}_i, \text{mem}_i \rangle, \langle \text{tps}'_i, \text{mds}'_i, \text{mem}'_i \rangle) \in \mathcal{B}_i \}. \end{aligned}$$

The relation $\mathcal{R}$ is a simulation, and therefore $\text{CC}^{\sim}$ holds. To apply Theorem 5.2 and conclude that whenever a source program $\mathbf{W}$ satisfies $\text{VDNI}_{\mathbf{S}} = \text{ANI}_{\phi}^{\rho}$, then $\mathbf{W}\downarrow$ satisfies $\text{VDNI}_{\mathbf{T}} = \text{ANI}_{\phi}^{\rho}$, it is sufficient for ρ to satisfy Equation 1, that is,

$$\rho(\mathbf{t}^*) = \rho(\tilde{\tau}^*(\rho(\mathbf{s}^*)))$$

for $\mathbf{s}^* \sim \mathbf{t}^*$. If one is willing to unfold all definitions, then this amounts to show the set of traces “bismilar” to $\mathbf{t}^*$ coincides with the set of traces that are bisimilar to some $\mathbf{t}''$ and $\mathbf{s}'' \sim \mathbf{t}''$ for some $\mathbf{s}''$ bisimilar to $\mathbf{s}^*$. Splitting the “coincides” (set equality) into the two directions of inclusion, the “*subsetq*” direction is immediate, while for the “*supseteq*” direction one has to prove some properties of $\mathcal{R}$, the ones in the definition of secure – refinement (Murray et al. [52, inlined above Theorem 5.1]) which entails preservation of low-equivalence as shown in Murray et al. [52, Theorem 5.1].

In summary, our framework makes it possible to precisely characterize the target noninterference properties that are implied by (trace-relating) correct compilation of source noninterfering programs. As we have shown, such properties are not necessarily as strong as desired. Crucially, the target noninterference property one gets *for free* for a given trace-relating correct compiler is a function of the trace relation under consideration. By considering more sophisticated trace relations, one could be able to get more interesting noninterference properties in the target *for free*—but this would likely come at the expense of a more challenging trace-relating compiler correctness proof.

6 TRACE-RELATING SECURE COMPILATION

So far, we have studied compiler correctness criteria for whole, standalone programs. However, in practice, programs do not exist in isolation, but in a context where they interact with other programs, libraries, etc. In many cases, this context cannot be assumed to be benign and could instead behave maliciously to try to disrupt a compiled program.

Hence, in this section, we consider the following *secure compilation* scenario: A source program is compiled and linked with an arbitrary target-level context, i.e., one that may not be expressible as the compilation of a source context. Compiler correctness does not address this case, as it does not consider arbitrary target contexts, looking instead at whole programs (empty context [41]) or well-behaved target contexts that behave like source ones (as in compositional compiler correctness [33, 37, 56, 76]).

Summary of the work of Abate et al. [3]. To account for this scenario, Abate et al. [3] describe several secure compilation criteria based on the preservation of classes of (hyper)properties (e.g., trace properties, safety, hypersafety, hyperproperties) against arbitrary target contexts. For each of these criteria, they give an equivalent “property-free” criterion, analogous to the equivalence between TP and $\text{CC}^{\sim}$. For instance, their **robust trace property preservation criterion** (RTP) states that, for any trace property π , if a source *partial* program $\mathbf{P}$ plugged into any context $\mathbf{C}_{\mathbf{S}}$ satisfies π , then the compiled program $\mathbf{P}\downarrow$ plugged into any target context $\mathbf{C}_{\mathbf{T}}$ satisfies π . Their equivalent criterion to RTP is RTC, which states that for any trace produced by the compiled program, when linked with any target context, there is a source context that produces the same

trace. Formally (writing $C[P]$ to mean the whole program that results from linking partial program P with context C) they define:

$$\begin{aligned} \text{RTP} &\equiv \forall P. \forall \pi. (\forall C_S. \forall t. C_S[P] \rightsquigarrow t \Rightarrow t \in \pi) \Rightarrow (\forall C_T. \forall t. C_T[P\downarrow] \rightsquigarrow t \Rightarrow t \in \pi), \\ \text{RTC} &\equiv \forall P. \forall C_T. \forall t. C_T[P\downarrow] \rightsquigarrow t \Rightarrow \exists C_S. C_S[P] \rightsquigarrow t. \end{aligned}$$

In the following, we adopt the notation $P \models_R \pi$ to mean “ P robustly satisfies π ,” i.e., P satisfies π irrespective of the contexts (C) it is linked with. Formally, $P \models_R \pi \stackrel{\text{def}}{=} \forall C. C[P] \models \pi$, where $\models$ is the same as before. Thus, we write more compactly:

$$\text{RTP} \equiv \forall \pi. \forall P. P \models_R \pi \Rightarrow P\downarrow \models_R \pi.$$

All the criteria of Abate et al. [3] share this flavor of stating the existence of some source context that simulates the behavior of any given target context, with some variations depending on the class of (hyper)properties under consideration. For trace properties, they also have criteria that preserve safety properties plus their version of liveness properties. For hyperproperties, they have criteria that preserve hypersafety properties, subset-closed hyperproperties, and arbitrary hyperproperties. Finally, they define *relational* hyperproperties, which are relations between the behaviors of multiple programs for expressing, e.g., that a program *always* runs faster than another. For relational hyperproperties, they have criteria that preserve arbitrary relational properties, relational safety properties, relational hyperproperties, and relational subset-closed hyperproperties.

Each category of criteria provides different kinds of security guarantees (confidentiality or integrity) for the code and data segments of programs. Roughly speaking, the security guarantees due to robust preservation of trace properties regard only protecting the integrity of the program from the context, the guarantees of hyperproperties also regard data confidentiality, and the guarantees of relational hyperproperties may even regard code confidentiality. Naturally, these stronger guarantees are increasingly harder to enforce and prove.

All the criteria of Abate et al. [3] are stated in a setting where source and target traces are the same. In this section, we extend their results to the trace-relating setting, obtaining trintarian views for secure compilation. There are many similarities with Section 2 that show up in the secure compilation setting, too, but also some crucial differences. As in Section 2, the application of $\tilde{\sigma}$ or $\tilde{\tau}$, may lose the information that a property belongs to the class *Safety*, or that a hyperproperty is subset-closed, which are both crucial for the equivalence with the property-free criterion of Abate et al. [3]. As in Section 2, we solve this problem by interpreting classes of properties as an *abstraction* of another class of properties induced by a closure operator. Differently from Section 2, the presence of adversarial contexts makes the criteria for subset-closed hyperproperties and trace properties distinct. Abate et al. [3] show that the criterion for robust preservation of hypersafety is distinct from robust safety preservation, and all criteria about classes of trace properties are distinct from their relational counterparts, e.g., robust preservation of relational safety and robust preservation of safety properties are different. We therefore further generalize the argument from Section 3.2 to safety hyperproperties as well as to relational hyperproperties.

Specifically, we provide a trinity for the preservation of trace properties and subset-closed hyperproperties (Section 6.1), of safety properties and hypersafety hyperproperties (Section 6.2), of hyperproperties (Section 6.3), and for 2-relational (hyper)properties (Section 6.4). We conclude the section by studying the relative expressiveness of these criteria (Section 6.5).

Robustness and Compositional Compilation. Before diving into the criteria for robust compilation, it is worth noting the relationship between these and compositional compiler correctness. **Compositional compiler correctness (CCC)** is a statement of compiler correctness for programs that are linked against *some* contexts. Unlike robustness, which imposes no constraints on the contexts,

CCC imposes conditions on the target contexts that compiled programs can be linked against: They need to be related (in ways that vary from work-to-work [38, 56]) to the source contexts [65]. As Patrignani and Garg [64] also point out, the notions of CCC and of robust compilation are incomparable: Neither can be proven stronger than the other. This is not surprising, since robust compilation criteria are used to prove compiler security while CCC is used to prove correctness.¹⁵

The criteria we adopt could be generalized further by adding an extra parameter that qualifies the relation between source and target contexts. Such a general statement would let us express both CCC and robust compilation by picking the correct extra parameter. However, we refrain from presenting such general statements, as the implications in terms of preservation of classes of (hyper)properties has not been studied for them.

6.1 Trace-relating Secure Compilation: Trace Properties and Subset-closed Hyperproperties

This section shows the simple generalization of RTC to the trace-relating setting ($\text{RTC}^\sim$) and its corresponding trinitarian view (Theorem 6.1). Then, it presents the trinitarian view for criteria that preserve subset-closed hyperproperties (Theorem 6.2).

THEOREM 6.1 (TRINITY FOR ROBUST TRACE PROPERTIES $\mathfrak{A}$). *For any trace relation $\sim$ and induced property mappings $\tilde{\tau}$ and $\tilde{\sigma}$, we have: $\text{RTP}^{\tilde{\tau}} \iff \text{RTC}^\sim \iff \text{RTP}^{\tilde{\sigma}}$, where*

$$\begin{aligned} \text{RTC}^\sim &\equiv \forall P \forall C_T \forall t. C_T[P \downarrow] \rightsquigarrow t \Rightarrow \exists C_S \exists s \sim t. C_S[P] \rightsquigarrow s, \\ \text{RTP}^{\tilde{\tau}} &\equiv \forall P \forall \pi_S \in 2^{\text{Traces}}. P \models_R \pi_S \Rightarrow P \downarrow \models_R \tilde{\tau}(\pi_S), \\ \text{RTP}^{\tilde{\sigma}} &\equiv \forall P \forall \pi_T \in 2^{\text{Traces}}. P \models_R \tilde{\sigma}(\pi_T) \Rightarrow P \downarrow \models_R \pi_T. \end{aligned}$$

The trinity for robust trace property preservation is the straightforward adaptation of the concepts of Section 2 to the definitions of Abate et al. [3]. Intuitively, these criteria simply deal with partial programs P instead of whole programs W . Necessarily, these criteria then consider arbitrary program contexts linked with P ; the universal quantification over C_S and C_T are tacit in the expression $\models_R$.

We can also generalize Section 2 to *robust* subset-closed hyperproperties (Theorem 6.2). However, unlike the correct compilation case of Section 2, the equivalent property-free criterion ($\text{RSCHC}^\sim$) does not coincide with $\text{RSC}^\sim$, but states the existence of a single source context for all the target traces produced by a program in a given context.

THEOREM 6.2 (TRINITY FOR ROBUST SUBSET-CLOSED HYPERPROPERTIES $\mathfrak{A}$). *Let SCH_S and SCH_T denote the sets of all subset-closed hyperproperties in the source and target languages, respectively. For any trace relation $\sim$ and its existential and universal images lifted to hyperproperties (that is, the lifting of the respective functions from Definition 2.5), $\tilde{\tau}$ and $\tilde{\sigma}$, and for $\text{Cl}_\subseteq(H) = \{\pi \mid \exists \pi' \in H. \pi \subseteq \pi'\}$, we have: $\text{RSCHP}^{\text{Cl}_\subseteq \circ \tilde{\tau}} \iff \text{RSCHC}^\sim \iff \text{RSCHP}^{\text{Cl}_\subseteq \circ \tilde{\sigma}}$, where*

$$\begin{aligned} \text{RSCHC}^\sim &\equiv \forall P \forall C_T \exists C_S \forall t. C_T[P \downarrow] \rightsquigarrow t \Rightarrow \exists s \sim t'. C_S[P] \rightsquigarrow s, \\ \text{RSCHP}^{\text{Cl}_\subseteq \circ \tilde{\tau}} &\equiv \forall P \forall H_S \in \text{SCH}_S. P \models_R H_S \Rightarrow P \downarrow \models_R \text{Cl}_\subseteq(\tilde{\tau}(H_S)), \\ \text{RSCHP}^{\text{Cl}_\subseteq \circ \tilde{\sigma}} &\equiv \forall P \forall H_T \in \text{SCH}_T. P \models_R \text{Cl}_\subseteq(\tilde{\sigma}(H_T)) \Rightarrow P \downarrow \models_R H_T. \end{aligned}$$

¹⁵ We remark CCC has been used to conclude security of compilation in the previously discussed work of Sison and Murray [74] (and in its predecessor [52]). However, there is a key difference in the “role” of contexts: In robust compilation criteria, contexts model attackers, while in Sison and Murray [74] contexts are other bits of compiled code. This treatment lets Sison and Murray [74] reason compositionally about the concurrently executing compiled code.

6.2 Trace-relating Secure Compilation: Safety and Hypersafety

In this section, we elaborate the robust preservation of safety (Theorem 6.3) and hypersafety properties (Theorem 6.4). Similar to Section 3.2, we consider the trace model adopted by Abate et al. [3] to ease the presentation. Our starting point is the two equivalent criteria for preservation of robust satisfaction of *all* and *only* the safety properties [3],

$$\begin{aligned} \text{RSP} &\equiv \forall P. \forall \pi \in \text{Safety}. P \models_{\mathbb{R}} \pi \Rightarrow P \downarrow \models_{\mathbb{R}} \pi, \\ \text{RSC} &\equiv \forall P. \forall C_T. \forall m. C_T [P \downarrow] \rightsquigarrow^* m \Rightarrow \exists C_S. C_S [P] \rightsquigarrow^* m, \end{aligned}$$

where $C_T [P \downarrow] \rightsquigarrow^* m$ is a shorthand for $\exists t \geq m. C_T [P \downarrow] \rightsquigarrow t$.

RSP differs from RTP, as it only quantifies over safety properties, and RSC differs from RTC, as it quantifies over finite prefixes m , rather than complete traces t . This comes from the fact that safety properties can be characterized in terms of sets of *bad* prefixes (as in Definition 3.4). Unfolding $\rightsquigarrow^*$, we can interpret RSC as follows: If $C_T [P \downarrow]$ produces a trace $t \geq m$ that violates a specific safety property, namely, the one defined by $M = \{m\}$, then there exists C_S in which P violates the *same* safety property, producing a trace $t' \geq m$ but possibly distinct from t .

Our generalization of RSC to the trace-relating setting states that whenever $C_T [P \downarrow]$ produces a trace t that violates a target safety property, there exists a source context C_S in which P violates the source *interpretation* of the property, i.e., its image through $\tilde{\sigma}$. The following theorem defines $\text{RSC}^\sim$ and its two equivalent formulations:

THEOREM 6.3 (TRINITY FOR ROBUST SAFETY PROPERTIES $\mathfrak{R}^3$). *For any trace relation $\sim$ and for the corresponding property mappings $\tilde{\tau}$ and $\tilde{\sigma}$, we have: $\text{RTP}^{\text{Safe} \circ \tilde{\tau}} \iff \text{RSC}^\sim \iff \text{RSP}^{\tilde{\sigma}}$, where*

$$\begin{aligned} \text{RSC}^\sim &\equiv \forall P \forall C_T \forall t \forall m \leq t. C_T [P \downarrow] \rightsquigarrow t \Rightarrow \exists C_S \exists t' \geq m \exists s \sim t'. C_S [P] \rightsquigarrow s, \\ \text{RTP}^{\text{Safe} \circ \tilde{\tau}} &\equiv \forall P \forall \pi_S \in 2^{\text{Traces}}. P \models_{\mathbb{R}} \pi_S \Rightarrow P \downarrow \models_{\mathbb{R}} (\text{Safe} \circ \tilde{\tau})(\pi_S), \\ \text{RSP}^{\tilde{\sigma}} &\equiv \forall P \forall \pi_T \in \text{Safety}_T. P \models_{\mathbb{R}} \tilde{\sigma}(\pi_T) \Rightarrow P \downarrow \models_{\mathbb{R}} \pi_T, \end{aligned}$$

where the closure operator *Safe* is the one introduced in Section 3.2.

Exactly like Section 3.2, Theorem 6.3 exploits the fact that

$$\text{Safe} \circ \tilde{\tau} : 2^{\text{Traces}} \rightleftharpoons \text{Safety}_T : \tilde{\sigma}$$

is a Galois connection between source properties and target safety properties and the argument generalizes to arbitrary closure operators on target properties ($\mathfrak{R}^3$). More interestingly, we can further generalize this idea to hypersafety. Hypersafety lifts the idea of safety with another level of sets (just like hyperproperties do w.r.t. trace properties) to talk about multiple runs of the same program. Just like for safety, hypersafety is concerned with a set of bad prefixes (called M) that no program upholding the hypersafety property should extend. Formally, a hyperproperty H is hypersafety if: $\forall \pi. \pi \notin H \Rightarrow (\exists M. M < \pi \wedge (\forall \pi' M < \pi' \Rightarrow \pi' \notin H))$. In Theorem 6.4, we indeed exploit the following Galois connection between source subset-closed hyperproperties and target:

$$\text{HSafe} \circ \tilde{\tau} : \text{SCH}_S \rightleftharpoons \text{HSafety}_T : \text{Cl}_\subseteq \circ \tilde{\sigma},$$

where $\text{HSafety}_T = \{ \text{HSafe}(\mathbf{H}_T) \mid \mathbf{H}_T \in 2^{2^{\text{Traces}_T}} \}$ and *HSafe* is the closure operator that maps an arbitrary target hyperproperty $\mathbf{H}_T$ to the target hypersafety that best over-approximates $\mathbf{H}_T$.¹⁶

¹⁶ $\text{HSafe}(\mathbf{H}_T) = \cap \{ \mathbf{H}'_T \mid \mathbf{H}_T \subseteq \mathbf{H}'_T \wedge \mathbf{H}'_T \in \text{HSafety}_T \}$. See, e.g., Clarkson and Schneider [18] and Pasqua and Mastroeni [58].

THEOREM 6.4 (TRINITY FOR ROBUST HYPERSAFETY $\mathcal{A}$). *For any trace relation $\sim$ and for the induced property mappings $\tilde{\tau}$ and $\tilde{\sigma}$, we have: $\text{RSCHP}^{\text{HSafe} \circ \tilde{\tau}} \iff \text{RHSC}^{\sim} \iff \text{RHSP}^{Cl_{\subseteq} \circ \tilde{\sigma}}$, where*

$$\begin{aligned} \text{RHSC}^{\sim} &\equiv \forall P \forall C_T \forall M \in \mathcal{M}^{\text{fin}}. M \leq \text{beh}(C_T[P \downarrow]) \Rightarrow \\ &\quad \exists C_S \forall m \in M \exists t \geq m. \exists s \sim t. C_S[P] \rightsquigarrow s, \\ \text{RSCHP}^{\text{HSafe} \circ \tilde{\tau}} &\equiv \forall P \forall H_S \in \text{SCH}_S. P \models_R H_S \Rightarrow P \downarrow \models_R \text{HSafe}(\tilde{\tau}(H_S)), \\ \text{RHSP}^{Cl_{\subseteq} \circ \tilde{\sigma}} &\equiv \forall P \forall H_T \in \text{HSafety}_T. P \models_R Cl_{\subseteq}(\tilde{\sigma}(H_T)) \Rightarrow P \downarrow \models_R H_T, \end{aligned}$$

and $\mathcal{M}^{\text{fin}}$ is the set of finite sets of prefixes.

We conclude this section with the following remark: The reader might wonder about extracting a “new” trace relation from the Galois connection $\text{Safe} \circ \tilde{\tau} : 2^{\text{Traces}} \hookrightarrow \text{Safety}_T : \tilde{\sigma}$ and get another formulation of $\text{RSC}^{\sim}$. We note that this is not possible in general, as the class of safety properties, i.e., closed sets, is not necessarily a powerset and hence Lemma 2.7 cannot be applied.

6.3 Trace-relating Secure Compilation: Arbitrary Hyperproperties

We already mentioned that some properties of interest for security, e.g., possibilistic information-flow are not subset closed [18]. In this section, we lift the results from Section 3.3 to the *secure* compilation setting. Once again, the trinity is *weak*, as the equivalence to $\text{RHP}^{\tilde{\sigma}}$ requires an extra assumption.

THEOREM 6.5 (WEAK TRINITY FOR ROBUST HYPERPROPERTIES $\mathcal{A}$). *For a trace relation $\sim \subseteq \text{Traces}_S \times \text{Trace}_T$ and induced property mappings $\tilde{\sigma}$ and $\tilde{\tau}$, we have:*

$$\begin{aligned} \text{RHC}^{\sim} &\iff \text{RHP}^{\tilde{\tau}}; \\ \text{if } \tilde{\tau} \hookrightarrow \tilde{\sigma} \text{ is a Galois insertion (i.e., } \tilde{\tau} \circ \tilde{\sigma} = \text{id}), \text{ then } \text{RHC}^{\sim} &\Rightarrow \text{RHP}^{\tilde{\sigma}}, \\ \text{if } \tilde{\sigma} \hookrightarrow \tilde{\tau} \text{ is a Galois reflection (i.e., } \tilde{\sigma} \circ \tilde{\tau} = \text{id}), \text{ then } \text{RHP}^{\tilde{\sigma}} &\Rightarrow \text{RHP}^{\tilde{\tau}}, \end{aligned}$$

$$\begin{aligned} \text{where } \text{RHC}^{\sim} &\equiv \forall P \forall C_T \exists C_S \forall t. C_T[P \downarrow] \rightsquigarrow t \iff (\exists s \sim t. C_S[P] \rightsquigarrow s), \\ \text{RHP}^{\tilde{\tau}} &\equiv \forall P \forall H_S. P \models_R H_S \Rightarrow P \downarrow \models_R \tilde{\tau}(H_S), \\ \text{RHP}^{\tilde{\sigma}} &\equiv \forall P \forall H_T. P \models_R \tilde{\sigma}(H_T) \Rightarrow P \downarrow \models_R H_T. \end{aligned}$$

It is therefore possible and correct to deduce a source obligation for a given target hyperproperty H_T ($\text{RHC}^{\sim} \Rightarrow \text{RHP}^{\tilde{\sigma}}$) when no information is lost in the composition $\tilde{\tau} \circ \tilde{\sigma}$. However, $\text{RHP}^{\tilde{\tau}}$ is a consequence of $\text{RHP}^{\tilde{\sigma}}$ when no information is lost in composing in the other direction, $\tilde{\sigma} \circ \tilde{\tau}$.

6.4 Trace-relating Secure Compilation: 2-Relational Hyperproperties

Finally, we turn to *relational* properties and hyperproperties. Relational hyperproperties, as defined by Abate et al. [3], are predicates on a sequence of behaviors; a sequence of programs has the relational hyperproperty if their behaviors collectively satisfy the predicate. Depending on the arity of the sequence, there exist different subclasses of relational hyperproperties, though, for simplicity, here, we only study relational hyperproperties of arity 2. A key example of a relational hyperproperty is trace equivalence, which holds if two programs have identical behaviors.

All the trinities in this section follow the pattern of their non-relational counterparts. We first explain how one can get a Galois connection between source and target relational properties from a trace relation.

Given a trace relation $\sim \subseteq \text{Traces}_S \times \text{Trace}_T$, we can relate pairs of source traces with pairs of target traces point-wise,

$$(s_1, s_2) \sim (t_1, t_2) \iff s_1 \sim t_1 \wedge s_2 \sim t_2.$$

Formally this is $\sim^2 \subseteq \text{Traces}_S^2 \times \text{Trace}_T^2$, the product of the relation $\sim$ with itself. Therefore, by Lemma 2.7 it corresponds to a Galois connection between source and target relational properties ($\mathcal{R}^2$), that with a little abuse of notation¹⁷ we still denote by

$$\tilde{\tau} : 2^{\text{Traces}_S \times \text{Traces}_S} \hookrightarrow 2^{\text{Trace}_T \times \text{Trace}_T} : \tilde{\sigma}.$$

Explicitly, for $r_S \in 2^{\text{Traces}_S \times \text{Traces}_S}$ and $r_T \in 2^{\text{Trace}_T \times \text{Trace}_T}$,

$$\begin{aligned} \tilde{\tau}(r_S) &= \{(\mathbf{t}_1, \mathbf{t}_2) \mid \exists (s_1, s_2). s_1 \sim \mathbf{t}_1 \wedge s_2 \sim \mathbf{t}_2 \wedge (s_1, s_2) \in r_S\}, \\ \tilde{\sigma}(r_T) &= \{(s_1, s_2) \mid \forall (\mathbf{t}_1, \mathbf{t}_2). s_1 \sim \mathbf{t}_1 \wedge s_2 \sim \mathbf{t}_2 \Rightarrow (\mathbf{t}_1, \mathbf{t}_2) \in r_T\}. \end{aligned}$$

$\tilde{\tau}$ and $\tilde{\sigma}$ are then lifted to relational hyperproperties similarly to Definition 3.2. Explicitly, for $R_S \in 2^{2^{\text{Traces}_S \times \text{Traces}_S}}$ and $R_T \in 2^{2^{\text{Trace}_T \times \text{Trace}_T}}$,

$$\begin{aligned} \tilde{\tau}(R_S) &= \{\tilde{\tau}(r_S) \mid r_S \in R_S\}, \\ \tilde{\sigma}(R_T) &= \{\tilde{\sigma}(r_T) \mid r_T \in R_T\}. \end{aligned}$$

Given a relational property $r \in 2^{\text{Trace} \times \text{Trace}}$ and two programs P_1, P_2 , we write $P_1, P_2 \models_R r$ for

$$\forall C. \forall t_1 t_2. C[P_1] \rightsquigarrow t_1 \wedge C[P_2] \rightsquigarrow t_2 \Rightarrow (t_1, t_2) \in r.$$

Given a relational hyperproperty $R \in 2^{2^{\text{Trace} \times \text{Trace}}}$, by $P_1, P_2 \models_R R$, we mean

$$\forall C. (\text{beh}(C[P_1]), \text{beh}(C[P_2])) \in R.$$

THEOREM 6.6 (TRINITY FOR ROBUST 2-RELATIONAL TRACE PROPERTIES $\mathcal{R}^2$). *For any trace relation $\sim$ and for the corresponding property mappings $\tilde{\tau}$ and $\tilde{\sigma}$, we have: $\text{R2rTP}^{\tilde{\tau}} \iff \text{R2rTC}^{\sim} \iff \text{R2rTP}^{\tilde{\sigma}}$, where*

$$\begin{aligned} \text{R2rTC}^{\sim} &\equiv \forall \mathbf{C}_T \forall P_1 \forall P_2 \forall \mathbf{t}_1 \forall \mathbf{t}_2. (\mathbf{C}_T[P_1] \rightsquigarrow \mathbf{t}_1 \wedge \mathbf{C}_T[P_2] \rightsquigarrow \mathbf{t}_2) \Rightarrow \\ &\quad \exists \mathbf{C}_S \exists s_1 \sim \mathbf{t}_1 \exists s_2 \sim \mathbf{t}_2. \mathbf{C}_S[P_1] \rightsquigarrow s_1 \wedge \mathbf{C}_S[P_2] \rightsquigarrow s_2, \\ \text{R2rTP}^{\tilde{\tau}} &\equiv \forall P_1 P_2 \forall r_S \in 2^{\text{Traces}_S \times \text{Traces}_S}. P_1, P_2 \models_R r_S \Rightarrow P_1 \downarrow, P_2 \downarrow \models_R \tilde{\tau}(r_S), \\ \text{R2rTP}^{\tilde{\sigma}} &\equiv \forall P_1 P_2. \forall r_T \in 2^{\text{Trace}_T \times \text{Trace}_T}. P_1, P_2 \models_R \tilde{\sigma}(r_T) \Rightarrow P_1 \downarrow, P_2 \downarrow \models_R r_T. \end{aligned}$$

Next, we propose the trinity for 2-relational subset-closed hyperproperties, i.e., elements of $2^{2^{\text{Trace} \times \text{Trace}}}$ that are closed under subsets. Exactly as in the case of subset-closed hyperproperties, the application of $\tilde{\tau}$ and $\tilde{\sigma}$ may lose the information of being subset-closed. To guarantee the equivalence of the three criteria, we compose the two mappings with a closure operator that we still denote by $Cl_{\subseteq}$.

THEOREM 6.7 (TRINITY FOR 2-RELATIONAL ROBUST SUBSET-CLOSED HYPERPROPERTIES $\mathcal{R}^2$). *For any trace relation $\sim$ and for the corresponding property mappings $\tilde{\tau}$ and $\tilde{\sigma}$, we have $\text{R2rSCHP}^{Cl_{\subseteq} \circ \tilde{\tau}} \iff \text{R2rSCHC}^{\sim} \iff \text{R2rSCHP}^{Cl_{\subseteq} \circ \tilde{\sigma}}$, where*

$$\begin{aligned} \text{R2rSCHC}^{\sim} &\equiv \forall \mathbf{C}_T \forall P_1 \forall P_2 \exists \mathbf{C}_S \forall \mathbf{t}_1 \forall \mathbf{t}_2. (\mathbf{C}_T[P_1] \rightsquigarrow \mathbf{t}_1 \wedge \mathbf{C}_T[P_2] \rightsquigarrow \mathbf{t}_2) \Rightarrow \\ &\quad \exists s_1 \sim \mathbf{t}_1 \exists s_2 \sim \mathbf{t}_2. \mathbf{C}_S[P_1] \rightsquigarrow s_1 \wedge \mathbf{C}_S[P_2] \rightsquigarrow s_2, \\ \text{R2rSCHP}^{Cl_{\subseteq} \circ \tilde{\tau}} &\equiv \forall P_1 \forall P_2 \forall R_S \in 2^{\text{RelSCH}_S}. P_1, P_2 \models_R R_S \Rightarrow P_1 \downarrow, P_2 \downarrow \models_R \tilde{\tau}(R_S), \\ \text{R2rSCHP}^{Cl_{\subseteq} \circ \tilde{\sigma}} &\equiv \forall P_1 \forall P_2 \forall R_T \in 2^{\text{RelSCH}_T}. P_1, P_2 \models_R \tilde{\sigma}(R_T) \Rightarrow P_1 \downarrow, P_2 \downarrow \models_R R_T. \end{aligned}$$

¹⁷Technically, we should write: $\tilde{\tau}^2 \hookrightarrow \tilde{\sigma}^2$.

We move now to the class of relational safety properties, a notion that generalizes safety properties to relations on programs. Similarly to Theorem 6.3, $\text{R2rSP}^{\tilde{\sigma}}$ quantifies over target relational safety properties, while $\text{R2rTP}^{2\text{rSafe}\tilde{\tau}}$ quantifies over all source relational property and compose $\tilde{\tau}$ with 2rSafe a closure operator that best approximates a relational property with a relational safety property.

THEOREM 6.8 (TRINITY FOR ROBUST 2-RELATIONAL SAFETY PROPERTIES $\clubsuit$). *For any trace relation $\sim$ and for the corresponding property mappings $\tilde{\tau}$ and $\tilde{\sigma}$, we have: $\text{R2rTP}^{2\text{rSafe}\tilde{\tau}} \iff \text{R2rSC}^{\sim} \iff \text{R2rSP}^{\tilde{\sigma}}$, where*

$$\begin{aligned} \text{R2rSC}^{\sim} &\equiv \forall \mathbf{C}_T \forall \mathbf{P}_1 \mathbf{P}_2 \forall \mathbf{t}_1 \mathbf{t}_2 \forall \mathbf{m}_1 \leq \mathbf{t}_1 \forall \mathbf{m}_2 \leq \mathbf{t}_2. \mathbf{C}_T [\mathbf{P}_1 \downarrow] \rightsquigarrow \mathbf{t}_1 \Rightarrow \mathbf{C}_T [\mathbf{P}_2 \downarrow] \rightsquigarrow \mathbf{t}_2 \Rightarrow \\ &\quad \exists \mathbf{C}_S \exists \mathbf{t}'_1 \geq \mathbf{m}_1 \exists \mathbf{s}_1 \sim \mathbf{t}'_1 \exists \mathbf{t}'_2 \geq \mathbf{m}_2 \exists \mathbf{s}_2 \sim \mathbf{t}'_2. \mathbf{C}_S [\mathbf{P}_1] \rightsquigarrow \mathbf{s}_1 \wedge \mathbf{C}_S [\mathbf{P}_2] \rightsquigarrow \mathbf{s}_2, \\ \text{R2rTP}^{2\text{rSafe}\tilde{\tau}} &\equiv \forall \mathbf{P}_1 \mathbf{P}_2 \forall \mathbf{r}_S \in 2^{\text{Traces} \times \text{Traces}}. \mathbf{P}_1, \mathbf{P}_2 \models_{\mathbf{R}} \mathbf{r}_S \Rightarrow \mathbf{P}_1 \downarrow, \mathbf{P}_2 \downarrow \models_{\mathbf{R}} (2\text{rSafe} \circ \tilde{\tau})(\mathbf{r}_S)), \\ \text{R2rSP}^{\tilde{\sigma}} &\equiv \forall \mathbf{P}_1 \mathbf{P}_2 \forall \mathbf{r}_T \in \mathbf{2rel}\text{-}\mathbf{Safety}_T. \mathbf{P}_1, \mathbf{P}_2 \models_{\mathbf{R}} \tilde{\sigma}(\mathbf{r}_T) \Rightarrow \mathbf{P}_1 \downarrow, \mathbf{P}_2 \downarrow \models_{\mathbf{R}} \mathbf{r}_T. \end{aligned}$$

Finally, we present the most general criterion: preservation of *arbitrary* 2-relational hyperproperties. As for the preservation of arbitrary hyperproperties, this (weak) trinity requires additional assumptions to hold, namely, that the Galois connection is an insertion or a reflection.

THEOREM 6.9 (WEAK TRINITY FOR ROBUST 2-RELATIONAL HYPERPROPERTIES $\clubsuit$). *For a trace relation $\sim \subseteq \text{Traces}_S \times \text{Trace}_T$ and the corresponding property mappings $\tilde{\sigma}$ and $\tilde{\tau}$, we have:*

$$\begin{aligned} \text{R2rHC}^{\sim} &\iff \text{R2rHP}^{\tilde{\tau}}; \\ \text{if } \tilde{\tau} \sqsubseteq \tilde{\sigma} \text{ is a Galois insertion (i.e., } \tilde{\tau} \circ \tilde{\sigma} &= \text{id}), \text{ then } \text{R2rHC}^{\sim} \Rightarrow \text{R2rHP}^{\tilde{\sigma}}, \\ \text{if } \tilde{\sigma} \sqsubseteq \tilde{\tau} \text{ is a Galois reflection (i.e., } \tilde{\sigma} \circ \tilde{\tau} &= \text{id}), \text{ then } \text{R2rHP}^{\tilde{\sigma}} \Rightarrow \text{R2rHP}^{\tilde{\tau}}, \end{aligned}$$

$$\begin{aligned} \text{where } \text{R2rHC}^{\sim} &\equiv \forall \mathbf{P}_1 \mathbf{P}_2 \forall \mathbf{C}_T \exists \mathbf{C}_S. \\ &\quad (\forall \mathbf{t}. \mathbf{C}_T [\mathbf{P}_1 \downarrow] \rightsquigarrow \mathbf{t} \iff (\exists \mathbf{s} \sim \mathbf{t}. \mathbf{C}_S [\mathbf{P}_1] \rightsquigarrow \mathbf{s})) \wedge \\ &\quad (\forall \mathbf{t}. \mathbf{C}_T [\mathbf{P}_2 \downarrow] \rightsquigarrow \mathbf{t} \iff (\exists \mathbf{s} \sim \mathbf{t}. \mathbf{C}_S [\mathbf{P}_2] \rightsquigarrow \mathbf{s})), \\ \text{R2rHP}^{\tilde{\tau}} &\equiv \forall \mathbf{P}_1 \forall \mathbf{P}_2 \forall \mathbf{r}_S. \mathbf{P}_1, \mathbf{P}_2 \models_{\mathbf{R}} \mathbf{r}_S \Rightarrow \mathbf{P}_1 \downarrow, \mathbf{P}_2 \downarrow \models_{\mathbf{R}} \tilde{\tau}(\mathbf{r}_S), \\ \text{R2rHP}^{\tilde{\sigma}} &\equiv \forall \mathbf{P}_1 \forall \mathbf{P}_2 \forall \mathbf{r}_T. \mathbf{P}_1, \mathbf{P}_2 \models_{\mathbf{R}} \tilde{\sigma}(\mathbf{r}_T) \Rightarrow \mathbf{P}_1 \downarrow, \mathbf{P}_2 \downarrow \models_{\mathbf{R}} \mathbf{r}_T. \end{aligned}$$

6.5 Relating the Secure Compilation Trinities

Figure 4 orders criteria referring to the same trace relation $\sim$ according to their relative strength. If a trinity entails another (denoted by $\Rightarrow$), then the former provides stronger security for a compilation chain than the latter.

The hypotheses of insertion and reflection mentioned in Theorem 6.9 and Theorem 6.5 are highlighted with the labels “Ins” and “Refl.” Recall that when composing $\tilde{\tau}$ with Safe , we quantify over the whole class of source trace properties rather than only safety properties. This is represented by the blue background in $\text{RTP}^{\text{Safe}\tilde{\tau}}$. The trinity for the robust preservation of arbitrary trace properties is on the same blue background. Red and green backgrounds are reserved for subset-closed hyperproperties and arbitrary relational properties and serve the same purpose.

We now describe how to interpret the acronyms in Figure 4. All criteria start with R, meaning they refer to robust preservation (secure compilation criteria). Criteria for relational hyperproperties—here only arity 2 is shown for simplicity—contain 2r. Next, criteria names spell the class of hyperproperties they preserve: H for hyperproperties, SCH for subset-closed hyperproperties, HS for hypersafety, T for trace properties, and S for safety properties. Finally, property-free criteria end with a C, while property-full ones involving $\tilde{\sigma}$ and $\tilde{\tau}$ end with P. Thus, *robust (R) subset-closed hyperproperty-preserving (SCH) compilation (C)* is $\text{RSCHC}^{\sim}$, *robust (R) two-relational (2r) safety-preserving (S) compilation (C)* is $\text{R2rSC}^{\sim}$, and so on.

Fig. 4. Hierarchy of trinitarian views of secure compilation criteria preserving classes of hyperproperties and the key to read each acronym. Shorthands “Ins.” and “Refl.” stand for Galois Insertion and Reflection. The symbol denotes trinitities proven in Coq.

7 INSTANCES OF TRACE-RELATING SECURE COMPILATION

This section presents instances of compilers that adopt our framework for secure compilation purposes. We provide three illustrative cases for compilers that, respectively, robustly preserve trace properties (Section 7.1), safety properties (Section 7.2), and hypersafety properties (Section 7.3). The last two examples are not novel instances we devise but rather existing work whose results we recount as instantiations of our framework.

7.1 An Instance of Trace-relating Robust Preservation of Trace Properties

This subsection illustrates trace-relating secure compilation when the target events are strictly more events than the source ones.

The source and target languages used here extend the syntax of the source language of Section 4.3.1. Both languages have *outputs of naturals*, and the expressions that generate them: `outsS n` and `outsS e`. Additionally, the target has a different output action and its related expression `outT n`; this is the only difference between the languages. The extra events in the target model the ability of target language to perform potentially dangerous operations (e.g., writing to the hard drive), which cannot be performed by the source language, and against which source-level reasoning can therefore offer no protection.

Both languages and compilation chains now deal with partial programs P , contexts C , and linking of those two to produce whole programs $C[P]$. In this setting, a whole program W is the combination of a *main expression* to be evaluated and a set of *function definitions* fs (with distinct names) that can refer to their argument (*arg*) symbolically and can be called by the main expression and by other functions ($f(e)$). The set of functions of a whole program is the union of

the functions of a partial program and a context; the latter also contains the main expression.

$$\begin{array}{ll}
 \mathbf{e} ::= \dots \mid \mathbf{f}(\mathbf{e}) \mid \mathbf{outs}_S \mathbf{n} \mid \mathbf{arg} & \mathbf{e} ::= \dots \mid \mathbf{f}(\mathbf{e}) \mid \mathbf{outs}_S \mathbf{n} \mid \mathbf{arg} \mid \mathbf{out}_T \mathbf{n} \\
 \mathbf{i} ::= \dots \mid \mathbf{outs}_S \mathbf{n} & \mathbf{i} ::= \dots \mid \mathbf{outs}_S \mathbf{n} \mid \mathbf{out}_T \mathbf{n} \\
 \mathbf{fs} ::= \langle f_1, e_1 \rangle, \dots, \langle f_n, e_n \rangle & P ::= \langle \mathbf{fs}, e \rangle \quad C ::= \mathbf{fs} \quad W ::= C[P]
 \end{array}$$

The extensions of the typing rules and the operational semantics for whole programs are unsurprising and therefore elided. The trace model also follows closely that of Section 4.3: It consists of a list of *regular events* (including the new outputs) terminated by a *result event*.¹⁸ A partial program and a context can be linked into a whole program when their functions satisfy the requirements mentioned above.

We define the homomorphic compiler $(\cdot \downarrow)$ that translates each source construct into its target correspondent. Thus, the compiler never introduces the additional target instruction $\mathbf{out}_T \mathbf{n}$. Since it is straightforward, the formalization of the compiler is elided.

Relating Traces. In the present model, source and target traces differ only in the fact that the target draws (regular) events from a strictly larger set than the source, i.e., $\Sigma_T \supset \Sigma_S$. A natural relation between source and target traces essentially maps a given target trace $\mathbf{t}$ the source trace that erases from $\mathbf{t}$ those events that exist only at the target level. This is reasonable, because only target contexts $\mathbf{C}$ (not compiled programs $\mathbf{P} \downarrow$) can perform the extra target actions, as the compiler does not introduce them. Let $\mathbf{t}|_{\Sigma_S}$ indicate trace $\mathbf{t}$ filtered to retain only those elements included in alphabet Σ_S . We define the trace relation as:

$$\mathbf{s} \sim \mathbf{t} \quad \equiv \quad \mathbf{s} = \mathbf{t}|_{\Sigma_S}.$$

In the opposite direction, a source trace $\mathbf{s}$ is related to many target ones, as any target-only events can be inserted at any point in $\mathbf{s}$. The induced mappings for this relation are:

$$\begin{aligned}
 \tilde{\tau}(\pi_S) &= \{ \mathbf{t} \mid \exists \mathbf{s}. \mathbf{s} = \mathbf{t}|_{\Sigma_S} \wedge \mathbf{s} \in \pi_S \}, \\
 \tilde{\sigma}(\pi_T) &= \{ \mathbf{s} \mid \forall \mathbf{t}. \mathbf{s} = \mathbf{t}|_{\Sigma_S} \Rightarrow \mathbf{t} \in \pi_T \}.
 \end{aligned}$$

That is, the target guarantee of a source property is that the target has the same source-level behavior, sprinkled with arbitrary target-level behavior. Conversely, the source-level obligation of a target property is the aggregate of those source traces, all of whose target-level enrichments are in the target property.

Since the languages are very similar, it is simple to prove that our compiler is secure according to the trace relation $\sim$ defined above.

THEOREM 7.1 ($\cdot \downarrow$ IS SECURE $\mathcal{R}$). $\cdot \downarrow$ is $\text{RTC}^\sim$.

7.2 An Instance of Trace-relating Robust Preservation of Safety Properties

I/O events are not the only instance of events that compilers consider. Especially in the setting of secure compilation, where a compartmentalized partial program interacts with a context, *interaction traces* are often used [3, 35, 59, 64]. Consider a language analogous to that of the previous section, where the context C defines a set of functions F_c and the program defines a different set F_p . Interaction traces (generally) record the control flow of calls between these two sets via actions that are *call* $f \ v$ and *ret* v [34]. These actions indicate a call to function f with parameter v and a return with return value v . In case the context calls a function in F_p (or returns to a function in F_p), the action is decorated with a $?$ (i.e., those actions are *call* $f \ v?$ and *ret* $v?$). Dually, the program

¹⁸Notice that the languages are strictly terminating.

calling a function in F_c (or returning to it) generates an action decorated with a $!$ (i.e., those actions are $call\ f\ v!$ and $ret\ v!$).

Patrignani and Garg [64] consider precisely such a setting. Their languages are simple like those presented here but impure; their source has an ML-like heap and the target has a memory that is indexed by natural numbers and capabilities to protect addresses. Moreover, they define a compiler that preserves safety properties of source programs (i.e., it is $RSC^\sim$ in the sense of Theorem 6.3) by relying on the target capabilities. The interesting point, however, is that they also consider source and target traces to be distinct, since the two languages have different values. Concretely, the source has **bools** and **nats** and the target only has **nats**, plus in the source, heap addresses are abstract locations ℓ while in the target they are **nats**. Thus, to prove $RSC^\sim$, they rely on a cross-language relation on values, which is lifted to trace actions, and then lifted point-wise to traces (analogously to what we have done in Section 4.3, 4.4, and 7.1). To relate addresses, their cross-language relation is equipped with a partial bijection between source and target addresses, this bijection grows monotonically with every reduction step.

Besides defining a relation on traces (which is an instance of $\sim$), they also define a relation between source and target safety properties that supports concurrent programs.¹⁹ Thus, they really provide an instantiation of τ that maps all safe source traces to the related target ones. This ensures that no additional target trace is introduced in the target property, and source safety properties are mapped to target safety ones by τ . Thus, their compiler is proven to generate code that respects τ , so they really achieve a variation of $RTP^{Safe\circ\tau}$ from Theorem 6.3. Their proofs are based on standard techniques either for secure compilation (i.e., trace-based backtranslation [61]) and for correct compilation (i.e., forward/backward simulation [42]).

7.3 An Instance of Trace-relating Robust Preservation of Hypersafety Properties

Patrignani and Garg [63] study the preservation of hypersafety from the perspective of secure compilation. Again, their result can be interpreted in our setting. They consider reactive systems, where trace alphabets are partitioned in input actions $\alpha?$ and output actions $\alpha!$, whose concatenation generate traces $\alpha^*\alpha!$. We use the same notation as before and indicate such sequences as **s** and **t**, respectively. The set of target output actions $\alpha!$ includes an action $\checkmark$ that has no source counterpart (i.e., $\nexists \alpha? \sim \checkmark$) and whose output does not depend on internal state (thus, it cannot leak secrets).²⁰ By emitting $\checkmark$ whenever undesired inputs are fed to a compiled program (e.g., passing a **nat** when a **bool** is expected), hypersafety is preserved (as $\checkmark$ does not leak secrets) [63].

More formally, they assume a relation on actions $\sim$ that is total on the source actions and injective. From there, they define **TPC**—which here corresponds to an instance of τ —that maps the set of valid source traces to the set of valid target traces (that now mention $\checkmark$) as follows:

$$\begin{aligned} \mathbf{TPC}(\pi_S) &= \left\{ \mathbf{t} \mid \mathbf{t} \in \bigcup_{n \in \mathbb{N}} \text{int}_n(\pi_S) \right\} \text{ where } \text{int}_0(\pi_S) = \{ \mathbf{t} \mid \exists s \in \pi_S \wedge s \sim \mathbf{t} \}, \\ \text{int}_{n+1}(\pi_S) &= \{ \mathbf{t} \mid \mathbf{t} \equiv \mathbf{t}_1 \alpha? \checkmark \mathbf{t}_2 \wedge \mathbf{t}_1 \mathbf{t}_2 \in \text{int}_n(\pi_S) \wedge \text{undesired}(\alpha?) \}, \end{aligned}$$

where $\text{undesired}(\alpha?)$ indicates that $\alpha?$ is an undesired input (intuitively, this is an information that can be derived from the set of source traces [63]).

Informally, given a set of source traces π_S , **TPC** generates all target traces that are related (point-wise) to a source trace (case int_0). Then (case int_{n+1}), it adds all traces (**t**) with interleavings of

¹⁹They call those safety properties monitors, since they focus on safety [72] and indicate **s** with **M** and **t** with **M**.

²⁰Technically, they assume a set of $\checkmark$ actions, but for this analogy a single action suffices.

undesired input $\alpha?$ (third conjunct) followed by $\surd$ (first conjunct) as long as the interleavings split a trace $t_1 t_2$ that has already been mapped (second conjunct).

TPC is an instance of τ that maps source hypersafety to target hypersafety (and therefore, safety to safety), thus our theory can be instantiated for the preservation of these classes of hyperproperties as well.

8 RELATED WORK

We already discussed how our results relate to some existing work in correct compilation [41, 77] and secure compilation [3, 63, 64]. We also already mentioned that most of our definitions and results make no assumptions about the structure of traces. One result that partially relies on the structure of traces is Theorem 6.3, which refers to *finite prefix* m , suggesting traces should be some sort of sequences of events (or states), as customary when one wants to refer to safety properties [18]. Without a notion of finite prefix, only $\text{RSC}^\sim$ may look different, but both $\text{RTP}^{\text{Safe}\circ\tau}$ and $\text{RSP}^{\hat{\sigma}}$ are trace-agnostic, as in general safety properties can be defined as the closed sets of any topology on traces [58].

Even for reasoning about safety, hypersafety, or arbitrary hyperproperties, traces can therefore be values, sequences of program states, or of input output events, or even the recently proposed *interaction trees* [81]. In the latter case, we believe that the compilation from IMP to ASM proposed by Xia et al. [81] can be seen as an instance of $\text{HC}^\sim$, for the relation they call “trace equivalence.”

Compilers Where Our Work Could Be Useful. Our work should be broadly applicable to understanding the guarantees provided by many verified compilers. For instance, Wang et al. [80] recently proposed a CompCert variant that compiles all the way down to machine code, and it would be interesting to see if the model at the end of Section 4.1 applies there, too. This and many other verified compilers [15, 36, 51, 73] beyond CakeML [77] deal with resource exhaustion, and it would be interesting to also apply the ideas of Section 4.2 to them.

Hur and Dreyer [33] devised a correct compiler from an ML language to assembly using a cross-language logical relation to state their CC theorem. They do not have traces, though were one to add them, the logical relation on values would serve as the basis for the trace relation and therefore their result would attain $\text{CC}^\sim$.

Switching to more informative traces capturing the interaction between the program and the context is often used as a proof technique for secure compilation [3, 34, 62]. Most of these results consider a cross-language relation, so they probably could be proved to attain one of the criteria from Figure 4.

Generalizations of Compiler Correctness. The compiler correctness definition of Morris [50] was already general enough to account for trace relations, since it considered a translation between the semantics of the source program and that of the compiled program, which he called “decode” in his diagram, reproduced in Figure 5 (left). And even some of the more recent compiler correctness definitions preserve this kind of flexibility [65]. While $\text{CC}^\sim$ can be seen as an instance of a definition by Morris [50], we are not aware of any prior work that investigated the preservation of properties when the “decode translation” is neither the identity nor a bijection, and source properties need to be re-interpreted as target ones and vice versa.

Correct Compilation and Galois Connections. Melton et al. [47] and Sabry and Wadler [70] expressed a strong variant of compiler correctness using the diagram of Figure 5 (right). They require that compiled programs *parallel* the computation steps ($\Rightarrow$) of the original source programs, which can be proven showing the existence of a *decompilation* map $\#$ that makes the diagram commute, or equivalently, the existence of an adjoint for $\downarrow$ ($W \leq W' \iff W \Rightarrow W'$ for both source and target). The “parallel” intuition can be formalized as an instance of $\text{CC}^\sim$. Take source

Fig. 5. Morris's [50] (left) and Melton et al.'s [47] and Sabry and Wadler's [70] (right) compiler correctness diagrams.

and target traces to be finite or infinite sequences of program states (maximal trace semantics [19]), and relate them exactly like Melton et al. [47] and Sabry and Wadler [70].

Translation Validation. Translation validation is an important alternative to proving that all runs of a compiler are correct, as it can be more easily applied to realistic compilers. An interesting work about translation validation of security properties has been recently proposed by Namjoshi and Tabajara [53]. They can handle many security properties expressible in terms of automata as long as source and target attackers and the observable traces are the same.

Instantiating the definition of any of the presented criteria with a particular program, one has translation validation criteria with the map $\tilde{r}$ describing the target property that is (robustly) satisfied once the translation is validated. For example, one can consider

$$(tsv\sim) CC\sim(W\downarrow) = \forall t. W\downarrow \rightsquigarrow t \Rightarrow \exists s \sim t. W \rightsquigarrow s,$$

$$(rtsv\sim) RTC\sim(C_T[P\downarrow]) = \forall t. C_T[P\downarrow] \rightsquigarrow t \Rightarrow \exists C_S. \exists s \sim t. C_S[P] \rightsquigarrow s.$$

While the proof technique proposed by Namjoshi and Tabajara [53] might be generalized for $CC\sim(W\downarrow)$ —as long as $beh(W\downarrow)$ and $beh(W)$ can be expressed as one of the automata they can handle—they do not work for $RTC\sim(C_T[P\downarrow])$ because of the existential in the conclusion.

Busi et al. [13] are instead considering translation validation criteria in the spirit of $(rtsv\sim)$, their preliminary work only allows equality as trace relation, but should be subject to a generalization to the trace-relating setting similar to the one we presented in this work.

Proof Techniques. We believe existing proof techniques (beyond the simulations discussed in Section 4.3.2) that have been devised to prove compiler correctness can also be employed to prove that a compiler attains any of the presented criteria. For example, cross-language binary logical relations can be used to relate two terms of two different languages when they “behave the same” [12, 33, 71]. Additionally, they can also be used when multiple programs “behave the same” [66] in a multilanguage semantics setting [45]. Secure compilation results (which rely on the criteria of Section 6) can be proven using variations of the *backtranslation* proof technique [22, 57, 64]. Presenting this proof techniques is beyond the scope of this article, so we refer the interested reader to the work of Patrignani et al. [61].

9 CONCLUSION AND FUTURE WORK

We have extended the property preservation view on compiler correctness to arbitrary trace relations, and we believe that this will be useful for understanding the guarantees various compilers provide. An open question is whether, given a compiler, there exists a most precise $\sim$ relation for which this compiler is correct. As mentioned in Section 1, every compiler is $CC\sim$ for some $\sim$, but under which conditions is there a most precise relation? In practice, more precision may not always be better, though, as it may be at odds with compiler efficiency and may not align with more subjective notions of usefulness, leading to tradeoffs in the selection of suitable relations. Finally, another interesting direction for future work is studying whether using the relation to Galois

connections allows to more easily compose trace relations for different purposes, say, for a compiler whose target language has undefined behavior, resource exhaustion, and side-channels. In particular, are there ways to obtain complex relations by combining simpler ones in a way that eases the compiler verification burden?

Composition for Multipass Compilers. For now, we can already informally argue about the correctness of a multipass compiler, where each step is proved correct for a possibly different trace relation. Concretely, assume $\downarrow_I^S$ is a compilation chain from a source language S to an intermediate language I and $\downarrow_T^I$ from the intermediate language I to a target language T .²¹ Assume given two relations between traces of these languages: $\sim_{S,I} \subseteq \text{Trace}_S \times \text{Trace}_I$ and $\sim_{I,T} \subseteq \text{Trace}_I \times \text{Trace}_T$, such that each compiler is proven to be CC w.r.t. the expected trace relation: $\downarrow_I^S \in \text{CC}^{\sim_{S,I}}$ and $\downarrow_T^I \in \text{CC}^{\sim_{I,T}}$.

Let us consider the source-to-target compiler $\downarrow_T^S$ that is derived of the composition of the two aforementioned compilers, so $\downarrow_T^S = \downarrow_T^I \circ \downarrow_I^S$. In this case, we obtain the expected result: The correctness of the whole compiler $\downarrow_T^S$ is derived from the individual compiler correctness proofs for each step.

$$\text{CC}^{(\sim_{I,T} \circ \sim_{S,I})} \equiv \forall W \forall t. W \downarrow_T^S \rightsquigarrow t \Rightarrow \exists s \sim_{I,T} s \circ \sim_{S,I} t. W \rightsquigarrow s,$$

where $s \sim_{I,T} s \circ \sim_{S,I} t \iff \exists i \in \text{Trace}_I. s \sim_{S,I} i \wedge i \sim_{I,T} t$.

Generalizing this kind of composition to compilers that attain different criteria is unclear. For example, if $\downarrow_I^S$ preserves arbitrary hyperproperties, but $\downarrow_T^I$ preserves 2-relational safety properties, then what can we conclude for $\downarrow_T^S$? We leave investigating these interesting matters for future work.

APPENDIX

A PROOFS

PROOF OF THEOREM 2.6 (♣). See Theorems `rel_TC_τTP` and `rel_TC_σTP` in `TraceCriterion.v`, where the $\text{TP}^{\tilde{\tau}} \iff \text{TP}^{\tilde{\sigma}}$ part follows directly from Theorem 2.4. $\square$

PROOF OF LEMMA 2.7 (TRACE RELATIONS $\cong$ GALOIS CONNECTIONS ON TRACE PROPERTIES).

Gardiner et al. [26] show that the existential image is a functor from the category of sets and relations to the category of predicate transformers, mapping a set $X \mapsto 2^X$ and a relation $\sim \subseteq X \times Y \mapsto \tilde{\tau} : 2^X \rightarrow 2^Y$. They also show that such a functor is an isomorphism—hence bijective—when one considers only monotonic predicate transformers that have a—unique—upper adjoint. The universal image of $\sim$, $\tilde{\sigma}$, is the unique adjoint of $\tilde{\tau}$ (♣), hence $\sim \mapsto \tilde{\tau} \leftrightsquigarrow \tilde{\sigma}$ is itself bijective. $\square$

PROOF OF THEOREM 2.8 (CORRESPONDENCE OF CRITERIA). For a trace relation $\sim$ and the Galois connection $\tilde{\tau} \leftrightsquigarrow \tilde{\sigma}$, the result follows from Theorem 2.6. For a Galois connection $\tau \leftrightsquigarrow \sigma$ and $\sim$, use Lemma 2.7 to conclude that the existential and universal images of $\sim$ coincide with τ and σ , respectively; the goal then follows from Theorem 2.6. $\square$

LEMMA A.1 (SPECIAL RELATIONS AND CONSEQUENCES ON THE ADJOINTS). *Let X, Y be two arbitrary sets and $\sim \subseteq X \times Y$. Assume $\sim$ is a total and surjective map from Y to X . Let $\alpha \leftrightsquigarrow \gamma$ be its existential and universal image, i.e.,*

$$\begin{aligned} \tilde{\alpha} &= \lambda \pi_X. \{y \mid \exists x \in \pi_X. x \sim y\}, \\ \tilde{\gamma} &= \lambda \pi_Y. \{x \mid \forall y. x \sim y \Rightarrow y \in \pi_Y\}. \end{aligned}$$

Then $\tilde{\gamma} = \lambda \pi_Y. \{x \mid \exists y \in \pi_Y. x \sim y\}$, and $\tilde{\gamma}$ is injective.

²¹For the intermediate language, we use a `verbatim`, `emerald` font.

PROOF OF LEMMA A.1. See Lemma `rel_total_surjective` and `rel_total_surjective_up_inj` in `Galois.v`. $\square$

PROOF OF THEOREM 4.3 (2). See Theorem `correctness` in `TypeRelationExampleInput.v`. $\square$

PROOF FOR LEMMA 4.4 (`gensend` ($\cdot, \cdot$) **WORKS**). We proceed by induction on τ and then by induction on τ' :

$\tau = \mathbf{N}$ and $\tau' = \mathbf{N}$ By canonicity, we have that $r = \langle n, n' \rangle$.

`gensend` ($\cdot, \cdot$) translates that into `send n; send n'`.

By Rule `Sem-seq`, that produces $t = n; n'$.

We need to prove that $\langle n, n' \rangle \sim n; n'$, which holds by Rule `Trace-Rel-N-N`.

$\tau = \mathbf{N}$ and $\tau' = \tau_1 \times \tau_2$ Analogous to the other cases, by IH and Rule `Trace-Rel-N-M`.

$\tau = \tau_1 \times \tau_2$ and $\tau' = \mathbf{N}$ Analogous to the other cases, by IH and Rule `Trace-Rel-M-N`.

$\tau = \tau_1 \times \tau'_1$ and $\tau' = \tau_2 \times \tau'_2$ So by canonicity $r = \langle \langle r_1, r'_1 \rangle, \langle r_2, r'_2 \rangle \rangle$.

By definition of `gensend` ($\cdot, \cdot$):

$$\begin{aligned} & \text{gensend}(\mathbf{x}, \tau \times \tau'), \\ &= \text{gensend}(\mathbf{x}, \tau).1; \text{gensend}(\mathbf{x}, \tau').2. \end{aligned}$$

By the target reductions, we know $(\text{gensend}(\mathbf{x}, \tau).1; \text{gensend}(\mathbf{x}, \tau').2)[r/x] \rightsquigarrow \text{is}_1; \text{is}_2$, so by IH, we have $\langle r_1, r'_1 \rangle \sim \text{is}_1$ and $\langle r_2, r'_2 \rangle \sim \text{is}_2$.

We need to prove that $\langle \langle r_1, r'_1 \rangle, \langle r_2, r'_2 \rangle \rangle \sim \text{is}_1; \text{is}_2$, which holds by Rule `Trace-Rel-M-M`, for $i = \langle r_1, r'_1 \rangle$ and $i' = \langle r_2, r'_2 \rangle$. $\square$

PROOF OF THEOREM 4.5. Trivial induction on the typing derivation of e , the only interesting case is the compilation of `send e` in the inductive cases.

Inductive. $e = \text{send } e$ By IH, we have that if $(\vdash e : \tau \times \tau') \downarrow \rightsquigarrow t$, then $\exists s \sim t$ and $e \rightsquigarrow t$.

By definition of $(\cdot) \downarrow$ and of $\rightsquigarrow$, we need to prove that if

$$\text{let } \mathbf{x} = (\vdash e : \tau \times \tau') \downarrow \text{ in } \text{gensend}(\mathbf{x}, \tau \times \tau') \rightsquigarrow t$$

Then $\text{send } e \rightsquigarrow s$ and $s \sim t$.

The reductions proceed as follows in the target:

$$\frac{(\vdash e : \tau \times \tau') \downarrow \rightsquigarrow \langle \text{is}, (\vdash r : \tau \times \tau') \downarrow \rangle \quad \text{gensend}(\mathbf{x}, \tau \times \tau')[(\vdash r : \tau \times \tau') \downarrow / \mathbf{x}] \rightsquigarrow \langle \text{is}', r' \rangle}{\text{let } \mathbf{x} = (\vdash e : \tau \times \tau') \downarrow \text{ in } \text{gensend}(\mathbf{x}, \tau \times \tau') \rightsquigarrow \langle \text{is} \cdot \text{is}', r' \rangle}$$

In the source, we have $\frac{e \rightsquigarrow \langle \text{is}, r \rangle}{\text{send } e \rightsquigarrow \langle \text{is} \cdot r, r \rangle}$

By IH, we have that $\text{is} \sim \text{is}$.

By Rule `Trace-Rel-Single`, to prove that $\text{is}; r \sim \text{is}; \text{is}'$, we need to prove that $\text{is} \sim \text{is}'$.

By Lemma 4.4 (`gensend` ($\cdot, \cdot$) **WORKS**), we have that $r \sim \text{is}'$, so this case holds. $\square$

PROOF OF THEOREM 5.1. First, we show that $\phi^\#$ is an *uco*, the proof for $\rho^\#$ is the same.

Monotonicity. $\phi^\#$ is composition of monotonic functions, hence it is itself monotonic.

Idempotence. We have to show that for π_T , $\phi^\#(\phi^\#(\pi_T)) = \phi^\#(\pi_T)$, that unfolding the definition means

$$g^\circ \circ \phi \circ f^\circ \circ g^\circ \circ \phi \circ f^\circ(\pi_T) = g^\circ \circ \phi \circ f^\circ(\pi_T).$$

For the inclusion “ $\subseteq$,”

$$g^\circ \circ \phi \circ f^\circ \circ g^\circ \circ \phi \circ f^\circ(\pi_T) \subseteq g^\circ \circ \phi \circ \phi \circ f^\circ(\pi_T) = g^\circ \circ \phi \circ f^\circ(\pi_T)$$

the inclusion holds, because $f^\circ \circ g^\circ(x) \subseteq x$ and the equality comes from idempotency of ϕ .

For the inclusion “ $\supseteq$,”

$$g^\circ \circ \phi \circ f^\circ \circ g^\circ \circ \phi \circ f^\circ(\pi_T) \supseteq g^\circ \circ \phi \circ f^\circ \circ g^\circ \circ f^\circ(\pi_T) = g^\circ \circ \phi \circ f^\circ(\pi_T)$$

the inclusion comes from $\phi(f^\circ(\pi_T)) \supseteq f^\circ(\pi_T)$ by extensiveness of ϕ , and the equality from $f^\circ \circ g^\circ \circ f^\circ = f^\circ$.

Extensiveness. We have to show that $\pi^\#(\pi_T) \supseteq \pi_T$.

$$\pi^\#(\pi_T) = g^\circ \circ \phi \circ f^\circ(\pi_T) \supseteq g^\circ \circ f^\circ(\pi_T) \supseteq \pi_T$$

The first inclusion is due to extensiveness of ϕ , the second by g° being the upper adjoint of f° .

For the statement of the theorem to hold, assume $\mathbb{W} \models \text{ANI}_\phi^\rho$ and $\mathbb{W} \downarrow \rightsquigarrow \mathbf{t}_1, \mathbf{t}_2$ with $\phi^\#(\mathbf{t}_1^\circ) = \phi^\#(\mathbf{t}_2^\circ)$, we have to show that $\rho^\#(\mathbf{t}_1^\circ) = \rho^\#(\mathbf{t}_2^\circ)$.

By CC $^\sim$ there exists $\mathbf{s}_1 \sim \mathbf{t}_1$ and $\mathbf{s}_2 \sim \mathbf{t}_2$ such that $\mathbb{W} \rightsquigarrow \mathbf{s}_1, \mathbf{s}_2$. As a preliminary, apply Lemma A.1 to the relations $\sim \circ \text{swap}$ and deduce g° is injective. Notice also that by functionality and totality, of $\sim$ and of f° , $f^\circ(\mathbf{t}_1^\circ) = \{\mathbf{s}_1^\circ\}$ and $f^\circ(\mathbf{t}_2^\circ) = \{\mathbf{s}_2^\circ\}$ and a similar fact holds for $\mathbf{s}_2$ and $\mathbf{t}_2$.

$$\begin{aligned} \phi^\#(\mathbf{t}_1^\circ) &= \phi^\#(\mathbf{t}_2^\circ) \Rightarrow & [\text{definition of } \phi^\#] \\ g^\circ \circ \phi \circ f^\circ(\mathbf{t}_1^\circ) &= g^\circ \circ \phi \circ f^\circ(\mathbf{t}_2^\circ) \Rightarrow & [g^\circ \text{ injective}] \\ \phi \circ f^\circ(\mathbf{t}_1^\circ) &= \phi \circ f^\circ(\mathbf{t}_2^\circ) \Rightarrow & [f^\circ(\mathbf{t}_i^\circ) = \mathbf{s}_i^\circ \ i = 1, 2] \\ \phi(\mathbf{s}_1^\circ) &= \phi(\mathbf{s}_2^\circ) \Rightarrow & [\mathbb{W} \models \text{ANI}_\phi^\rho] \\ \rho(\mathbf{s}_1^\circ) &= \rho(\mathbf{s}_2^\circ) \Rightarrow & [\mathbf{s}_i^\circ = f^\circ(\mathbf{t}_i^\circ) \ i = 1, 2] \\ \rho \circ f^\circ(\mathbf{t}_1^\circ) &= \rho \circ f^\circ(\mathbf{t}_2^\circ) \Rightarrow & [\text{functionality of } g^\circ] \\ g^\circ \circ \rho \circ f^\circ(\mathbf{t}_1^\circ) &= g^\circ \circ \rho \circ f^\circ(\mathbf{t}_2^\circ) \Rightarrow & [\text{definition of } \rho^\#] \\ \rho^\#(\mathbf{t}_1^\circ) &= \rho^\#(\mathbf{t}_2^\circ), \end{aligned}$$

so $\mathbb{W} \downarrow \models \text{ANI}_{\phi^\#}^{\rho^\#}$.

We now show that if $\sim$ is surjective, i.e., g° injective, then $\text{ANI}_{\phi^\#}^{\rho^\#} \subseteq \text{Cl}_\subseteq \circ \tilde{\tau}(\text{ANI}_\phi^\rho)$.

Let $\pi_T \in \text{ANI}_{\phi^\#}^{\rho^\#}$, we show that $\pi_T \subseteq \tilde{\tau}(\pi_S)$ for some $\pi_S \in \text{ANI}_\phi^\rho$.

The source property $\pi_S = \{s \mid \exists t \in \pi_T. s \sim t\} = f(\pi_T)$ is such that $\pi_T \subseteq \tilde{\tau}(\pi_S)$. We only need to show $\pi_S \in \text{ANI}_\phi^\rho$. Let $\mathbf{s}_1, \mathbf{s}_2 \in \pi_S$,

$$\begin{aligned} \phi(\mathbf{s}_1^\circ) &= \phi(\mathbf{s}_2^\circ) \Rightarrow & [\text{by } f^\circ(\mathbf{t}^\circ) = \mathbf{s}^\circ \text{ for some } \mathbf{t} \in \pi_T] \\ \phi(f^\circ(\mathbf{t}_1^\circ)) &= \phi(f^\circ(\mathbf{t}_2^\circ)) \Rightarrow & [g^\circ \text{ is a function}] \\ g^\circ(\phi(f^\circ(\mathbf{t}_1^\circ))) &= g^\circ(\phi(f^\circ(\mathbf{t}_2^\circ))) \Rightarrow & [\text{by definition of } \phi^\#] \\ \phi^\#(\mathbf{t}_1^\circ) &= \phi^\#(\mathbf{t}_2^\circ) \Rightarrow & [\pi_T \in \text{ANI}_{\phi^\#}^{\rho^\#}] \\ \rho^\#(\mathbf{t}_1^\circ) &= \rho^\#(\mathbf{t}_2^\circ) \Rightarrow & [\text{definition of } \rho^\#] \\ g^\circ(\rho(f^\circ(\mathbf{t}_1^\circ))) &= g^\circ(\rho(f^\circ(\mathbf{t}_2^\circ))) \Rightarrow & [\text{by injectivity of } g^\circ] \\ \rho(f^\circ(\mathbf{t}_1^\circ)) &= \rho(f^\circ(\mathbf{t}_2^\circ)) \Rightarrow & [f^\circ(\mathbf{t}_i^\circ) = \mathbf{s}_i^\circ, i = 1, 2] \\ \rho(\mathbf{s}_1^\circ) &= \rho(\mathbf{s}_2^\circ), \end{aligned}$$

that shows $\pi_S \in \text{ANI}_\phi^\rho$ and concludes the proof. $\square$

PROOF OF THEOREM 5.2. Assume $W \models \text{ANI}_{\phi}^{\rho}$ and $W \downarrow \rightsquigarrow t_1, t_2$ with $\phi^{\#}(t_1^{\circ}) = \phi^{\#}(t_2^{\circ})$. We have to show that $\rho^{\#}(t_1^{\circ}) = \rho^{\#}(t_2^{\circ})$, for an arbitrary $\rho^{\#}$ that satisfies the condition

$$H \equiv \forall s \ t. \ s^{\circ} \sim t^{\circ} \Rightarrow \rho^{\#}(\tilde{\tau}^*(\rho(s^{\circ}))) = \rho^{\#}(t^{\circ}).$$

By $\text{CC}^{\sim}$ there exists $s_1 \sim t_1$ and $s_2 \sim t_2$ such that $W \rightsquigarrow s_1, s_2$. As a preliminary, recall that Lemma A.1 ensures g° is injective. Moreover, notice that by functionality and totality, of $\sim$, $f^{\circ}(t_1^{\circ}) = \{s_1^{\circ}\}$ and $f^{\circ}(t_2^{\circ}) = \{s_2^{\circ}\}$.

$$\begin{aligned} \phi^{\#}(t_1^{\circ}) &= \phi^{\#}(t_2^{\circ}) \Rightarrow && [\text{definition of } \phi^{\#}] \\ g^{\circ} \circ \phi \circ f^{\circ}(t_1^{\circ}) &= g^{\circ} \circ \phi \circ f^{\circ}(t_2^{\circ}) \Rightarrow && [g^{\circ} \text{ injective}] \\ \phi \circ f^{\circ}(t_1^{\circ}) &= \phi \circ f^{\circ}(t_2^{\circ}) \Rightarrow && [f^{\circ}(t_i^{\circ}) = s_i^{\circ} \ i = 1, 2] \\ \phi(s_1^{\circ}) &= \phi(s_2^{\circ}) \Rightarrow && [W \models \text{ANI}_{\phi}^{\rho}] \\ \rho(s_1^{\circ}) &= \rho(s_2^{\circ}) \Rightarrow && [\text{functionality of } \tilde{\tau}^*] \\ \tilde{\tau}^*(\rho(s_1^{\circ})) &= \tilde{\tau}^*(\rho(s_2^{\circ})) \Rightarrow && [\text{by functionality of } \rho^{\#}] \\ \rho^{\#}(\tilde{\tau}^*(\rho(s_1^{\circ}))) &= \rho^{\#}(\tilde{\tau}^*(\rho(s_2^{\circ}))) \Rightarrow && [\text{by condition } H] \\ \rho^{\#}(t_1^{\circ}) &= \rho^{\#}(t_2^{\circ}) \end{aligned}$$

so $W \downarrow \models \text{ANI}_{\phi}^{\rho^{\#}}$. □

PROOF OF THEOREM 5.3. Assume $W \models \text{ANI}_{\phi^{\#}}^{\rho^{\#}}$ and $W \downarrow \rightsquigarrow t_1, t_2$ with $\phi(t_1^{\circ}) = \phi(t_2^{\circ})$ and ϕ satisfying the condition $H \equiv \forall s \ t. \ s^{\circ} \sim t^{\circ} \Rightarrow \phi(t^{\circ}) = \phi(\tilde{\tau}^{\circ}(s^{\circ}))$. We have to show that $\rho(t_1^{\circ}) = \rho(t_2^{\circ})$. By $\text{CC}^{\sim}$ there exists $s_1 \sim t_1$ and $s_2 \sim t_2$ such that $W \rightsquigarrow s_1, s_2$. As a preliminary, recall that Lemma A.1 ensures $\tilde{\sigma}^*$ is injective. Moreover, notice that by functionality and totality, of $\sim$, $\tilde{\tau}^*(s_1^{\circ}) = \{t_1^{\circ}\}$ and $\tilde{\tau}^*(s_2^{\circ}) = \{t_2^{\circ}\}$.

$$\begin{aligned} \phi(t_1^{\circ}) &= \phi(t_2^{\circ}) \Rightarrow && [\text{by } H] \\ \phi(\tilde{\tau}^{\circ}(s_1^{\circ})) &= \phi(\tilde{\tau}^{\circ}(s_2^{\circ})) \Rightarrow && [\text{functionality of } \tilde{\sigma}^{\circ}] \\ \tilde{\sigma}^{\circ}(\phi(\tilde{\tau}^{\circ}(s_1^{\circ}))) &= \tilde{\sigma}^{\circ}(\phi(\tilde{\tau}^{\circ}(s_2^{\circ}))) \Rightarrow && [\text{definition of } \phi^{\#}] \\ \phi^{\#}(s_1^{\circ}) &= \phi^{\#}(s_2^{\circ}) \Rightarrow && [W \models \text{ANI}_{\phi^{\#}}^{\rho^{\#}}] \\ \rho^{\#}(s_1^{\circ}) &= \rho^{\#}(s_2^{\circ}) \Rightarrow && [\text{by definition of } \rho^{\#}] \\ \tilde{\sigma}^*(\rho^{\#}(\tilde{\tau}^*(s_1^{\circ}))) &= \tilde{\sigma}^*(\rho^{\#}(\tilde{\tau}^*(s_2^{\circ}))) \Rightarrow && [\text{injectivity of } \tilde{\sigma}^*] \\ \rho^{\#}(\tilde{\tau}^*(s_1^{\circ})) &= \rho^{\#}(\tilde{\tau}^*(s_2^{\circ})) \Rightarrow && [\tilde{\tau}^*(s_i^{\circ}) = \{t_i^{\circ}\} \ i = 1, 2] \\ \rho^{\#}(t_1^{\circ}) &= \rho^{\#}(t_2^{\circ}) \end{aligned}$$

so $W \downarrow \models \text{ANI}_{\phi}^{\rho}$. □

PROOF OF THEOREM 6.1 (♣). Theorems `rel_RTC_τRTP` and `rel_RTC_σRTP` in `RobustTraceCriterion.v`. □

PROOF OF THEOREM 6.3 (♣). Theorems `tilde_RSC_σRSP` and `tilde_RSC_Cl_τRTP` in `RobustSafetyCriterion.v`. □

PROOF OF THEOREM 6.5 (♣). Lemmas $\sigma\text{RHP_rel_RHC}$ and $\text{rel_RHC_}\sigma\text{RHP}$ and Theorem $\text{rel_RHC_}\tau\text{RHP}$ in `RobustHyperCriterion.v`. $\square$

PROOF OF THEOREM 7.1 (♣). (See theorem `extra_target_RTCT` in `MoreTargetEventsExample.v`, mechanizing a slightly simplified model.) By definition of $\text{RTC}^\sim$, we need to find a source context and source trace given a source program, target context, and target trace related by compilation and program semantics: This instantiation is simple, since the trace relation is a *function* from target traces to source traces, and it is easy to clean target contexts to produce equivalent source context without target-only events. The proof is a trivial instance of *precise, context-based backtranslation* [3, 57, 61, 75], aided by a few straightforward lemmas and where the case of function calls is guaranteed to terminate by the language. $\square$

PROOF OF THEOREM 3.6 (♣). Theorems `tilde_SC_` σSP and `tilde_SC_Cl_` τTP in `SafetyCriterion.v`. $\square$

PROOF OF THEOREM 3.7. For the implication from left to right, assume $W \models H$. By $\text{CC}^\sim$ have $\text{beh}(W \downarrow) = \text{beh}(W)$, so $W \downarrow \models H$ as well. For the implication from right to left, instantiate HP with the hyperproperty $\{\text{beh}(W)\}$, for a given W , and deduce that $W \downarrow \models \{\text{beh}(W)\}$ i.e., $\text{beh}(W \downarrow) = \text{beh}(W)$. $\square$

PROOF OF THEOREM 3.8 (♣). Theorems `rel_HC_` τHP , `rel_HC_` σHP and $\sigma\text{HP_rel_HC}$ in `HyperCriterion.v`. $\square$

ACKNOWLEDGMENTS

We thank Akram El-Korashy and Amin Timany for participating in an early discussion about this work and the anonymous reviewers for their valuable feedback.

REFERENCES

- [1] Martín Abadi, Anindya Banerjee, Nevin Heintze, and Jon G. Riecke. 1999. A core calculus of dependency. In *Proceedings of the 26th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL'99)*. New York, NY, 147–160. DOI: <http://doi.org/10.1145/292540.292555>
- [2] Carmine Abate, Roberto Blanco, Ștefan Ciobăcă, Adrien Durier, Deepak Garg, Catalin Hritcu, Marco Patrignani, Éric Tanter, and Jérémy Thibault. 2020. Trace-relating compiler correctness and secure compilation. In *Proceedings of the 29th European Symposium on Programming: Programming Languages and Systems, Held as Part of the European Joint Conferences on Theory and Practice of Software*. 1–28. DOI: http://doi.org/10.1007/978-3-030-44914-8_1
- [3] Carmine Abate, Roberto Blanco, Deepak Garg, Cătălin Hrițcu, Marco Patrignani, and Jérémy Thibault. 2019. Journey beyond full abstraction: Exploring robust property preservation for secure compilation. In *Proceedings of the 32nd IEEE Computer Security Foundations Symposium (CSF'19)*. Retrieved from <https://arxiv.org/abs/1807.04603>.
- [4] Amal Ahmed, Deepak Garg, Cătălin Hrițcu, and Frank Piessens. 2018. Secure compilation (Dagstuhl seminar 18201). *Dagstuhl Rep.* 8, 5 (2018), 1–30. DOI: <http://doi.org/10.4230/DagRep.8.5.1>
- [5] Abhishek Anand, Andrew Appel, Greg Morrisett, Zoe Paraskevopoulou, Randy Pollack, Olivier Savary Belanger, Matthieu Sozeau, and Matthew Weaver. 2017. CertiCoq: A verified compiler for Coq. In *Proceedings of the CoqPL Workshop*. Retrieved from <https://www.cs.princeton.edu/~appel/papers/certicoq-coqpl.pdf>.
- [6] Kevin Backhouse and Roland Backhouse. 2004. Safety of abstract interpretations for free, via logical relations and Galois connections. *Sci. Comput. Program.* 51, 1–2 (2004), 153–196. Retrieved from <https://core.ac.uk/download/pdf/82190842.pdf>.
- [7] Gilles Barthe, Sandrine Blazy, Benjamin Grégoire, Rémi Hutin, Vincent Laporte, David Pichardie, and Alix Trieu. 2020. Formal verification of a constant-time preserving C compiler. *Proc. ACM Program. Lang.* 4, POPL (2020), 7:1–7:30. DOI: <http://doi.org/10.1145/3371075>
- [8] Gilles Barthe, Benjamin Grégoire, and Vincent Laporte. 2018. Secure compilation of side-channel countermeasures: The case of cryptographic “constant-time.” In *Proceedings of the 31st IEEE Computer Security Foundations Symposium (CSF'18)*. 328–343. DOI: <http://doi.org/10.1109/CSF.2018.00031>

- [9] Lennart Beringer, Gordon Stewart, Robert Dockins, and Andrew W. Appel. 2014. Verified compilation for shared-memory C. In *Proceedings of the 23rd European Symposium on Programming: Programming Languages and Systems, Held as Part of the European Joint Conferences on Theory and Practice of Software*. 107–127. DOI: http://doi.org/10.1007/978-3-642-54833-8_7
- [10] Frédéric Besson, Sandrine Blazy, and Pierre Wilke. 2019. A verified Comp-Cert front-end for a memory model supporting pointer arithmetic and uninitialised data. *J. Autom. Reason.* 62, 4 (2019), 433–480. DOI: <http://doi.org/10.1007/s10817-017-9439-z>
- [11] Sylvie Boldo, Jacques-Henri Jourdan, Xavier Leroy, and Guillaume Melquiond. 2015. Verified compilation of floating-point computations. *J. Autom. Reason.* 54, 2 (2015), 135–163. DOI: <http://doi.org/10.1007/s10817-014-9317-x>
- [12] William J. Bowman and Amal Ahmed. 2015. Noninterference for free. In *Proceedings of the ACM SIGPLAN International Conference on Functional Programming*.
- [13] Matteo Busi, Pierpaolo Degano, and Letterio Galletta. 2019. Translation validation for security properties. *CoRR* abs/1901.05082 (2019).
- [14] Qinxiong Cao, Lennart Beringer, Samuel Gruetter, Josiah Dodds, and Andrew W. Appel. 2018. VST-Floyd: A separation logic tool to verify correctness of C programs. *J. Autom. Reason.* 61, 1–4 (2018), 367–422. DOI: <http://doi.org/10.1007/s10817-018-9457-5>
- [15] Quentin Carbonneaux, Jan Hoffmann, Tahina Ramananandro, and Zhong Shao. 2014. End-to-end verification of stack-space bounds for C programs. In *Proceedings of the ACM SIGPLAN Conference on Programming Language Design and Implementation*, Michael F. P. O’Boyle and Keshav Pingali (Eds.). 270–281. DOI: <http://doi.org/10.1145/2594291.2594301>
- [16] Catalin Cimpanu. 2019. Microsoft: 70 percent of all security bugs are memory safety issues. ZDNet. Retrieved from <https://www.zdnet.com/article/microsoft-70-percent-of-all-security-bugs-are-memory-safety-issues/>.
- [17] David Clark and Sebastian Hunt. 2008. Non-interference for deterministic interactive programs. In *Proceedings of the International Workshop on Formal Aspects in Security and Trust*. Springer, 50–66.
- [18] Michael R. Clarkson and Fred B. Schneider. 2010. Hyperproperties. *J. Comput. Secur.* 18, 6 (2010), 1157–1210. DOI: <http://doi.org/10.3233/JCS-2009-0393>
- [19] Patrick Cousot. 2002. Constructive design of a hierarchy of semantics of a transition system by abstract interpretation. *Theoretical Computer Science* 277, 1–2 (2002), 47–103. Retrieved from <https://www.di.ens.fr/~cousot/COUSOTpapers/publications/www/Cousot-TCS-02-v277p47-103-2002.pdf>.
- [20] P. Cousot and R. Cousot. 1977. Abstract interpretation: A unified lattice model for static analysis of programs by construction or approximation of fixpoints. In *Proceedings of the 4th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*. 238–252.
- [21] Patrick Cousot and Radhia Cousot. 1979. Systematic design of program analysis frameworks. In *Proceedings of the 6th ACM SIGACT-SIGPLAN Symposium on Principles of Programming Languages*. 269–282.
- [22] Dominique Devriese, Marco Patrignani, and Frank Piessens. 2016. Fully-abstract compilation by approximate back-translation. In *Proceedings of the 43rd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*. Retrieved from http://www.mpi-sws.org/~marcopat/marcopat/Publications_files/logrel-for-facomp.pdf.
- [23] Vijay D’Silva, Mathias Payer, and Dawn Xiaodong Song. 2015. The correctness-security gap in compiler optimization. In *Proceedings of the IEEE Symposium on Security and Privacy Workshops*. 73–87. DOI: <http://doi.org/10.1109/SPW.2015.33>
- [24] Joost Engelfriet. 1985. Determinacy implies (observation equivalence = trace equivalence). *Theor. Comput. Sci.* 36 (1985), 21–25. DOI: [http://doi.org/10.1016/0304-3975\(85\)90028-3](http://doi.org/10.1016/0304-3975(85)90028-3)
- [25] Riccardo Focardi and Roberto Gorrieri. 1995. A taxonomy of security properties for process algebras. *J. Comput. Secur.* 3, 1 (1995), 5–34. DOI: <http://doi.org/10.3233/JCS-1994/1995-3103>
- [26] Paul H. B. Gardiner, Clare E. Martin, and Oege De Moor. 1994. An algebraic construction of predicate transformers. *Sci. Comput. Program.* 22, 1–2 (1994), 21–44. DOI: [https://doi.org/10.1016/0167-6423\(94\)90006-X](https://doi.org/10.1016/0167-6423(94)90006-X)
- [27] Roberto Giacobazzi and Isabella Mastroeni. 2018. Abstract non-interference: A unifying framework for weakening information-flow. *ACM Trans. Priv. Secur.* 21, 2 (2018), 9. DOI: <https://doi.org/10.1145/3175660>
- [28] Joseph A. Goguen and José Meseguer. 1982. Security policies and security models. In *Proceedings of the Symposium on Security and Privacy*. 11–20. Retrieved from https://www.cs.purdue.edu/homes/ninghui/readings/AccessControl/goguen_meseguer_82.pdf.
- [29] Ronghui Gu, Zhong Shao, Jieung Kim, Xiongnan (Newman) Wu, Jérémie Koenig, Vilhelm Sjöberg, Hao Chen, David Costanzo, and Tahina Ramananandro. 2018. Certified concurrent abstraction layers. In *Proceedings of the 39th ACM SIGPLAN Conference on Programming Language Design and Implementation*, Jeffrey S. Foster and Dan Grossman (Eds.). 646–661. DOI: <http://doi.org/10.1145/3192366.3192381>
- [30] István Haller, Yuseok Jeon, Hui Peng, Mathias Payer, Cristiano Giuffrida, Herbert Bos, and Erik van der Kouwe. 2016. TypeSan: Practical type confusion detection. In *Proceedings of the ACM SIGSAC Conference on Computer and Communications Security*. 517–528. DOI: <http://doi.org/10.1145/2976749.2978405>

- [31] Heartbleed. 2014. The Heartbleed Bug. Retrieved from <http://heartbleed.com/>.
- [32] Cătălin Hrițcu, David Chisnall, Deepak Garg, and Mathias Payer. 2019. Secure Compilation. SIGPLAN PL Perspectives Blog. Retrieved from <https://blog.sigplan.org/2019/07/01/secure-compilation/>.
- [33] Chung-Kil Hur and Derek Dreyer. 2011. A Kripke logical relation between ML and assembly. In *Proceedings of the 38th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, Thomas Ball and Mooly Sagiv (Eds.), 133–146. DOI: <http://doi.org/10.1145/1926385.1926402>
- [34] Alan Jeffrey and Julian Rathke. 2005. Java Jr: Fully abstract trace semantics for a core java language. In *Proceedings of the 14th European Symposium on Programming (Lecture Notes in Computer Science)*, Vol. 3444. 423–438. DOI: http://doi.org/10.1007/978-3-540-31987-0_29
- [35] Yannis Juglaret, Cătălin Hrițcu, Arthur Azevedo de Amorim, Boris Eng, and Benjamin C. Pierce. 2016. Beyond good and evil: Formalizing the security guarantees of compartmentalizing compilation. In *Proceedings of the IEEE 29th Computer Security Foundations Symposium*. 45–60. DOI: <http://doi.org/10.1109/CSF.2016.11>
- [36] Jeehoon Kang, Chung-Kil Hur, William Mansky, Dmitri Garbuzov, Steve Zdancewic, and Viktor Vafeiadis. 2015. A formal C memory model supporting integer-pointer casts. In *Proceedings of the 36th ACM SIGPLAN Conference on Programming Language Design and Implementation*. 326–335. DOI: <http://doi.org/10.1145/2737924.2738005>
- [37] Jeehoon Kang, Yoonseung Kim, Chung-Kil Hur, Derek Dreyer, and Viktor Vafeiadis. 2016. Lightweight verification of separate compilation. In *Proceedings of the ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*. Retrieved from <http://sf.snu.ac.kr/sepcompert/>.
- [38] Jeehoon Kang, Yoonseung Kim, Chung-Kil Hur, Derek Dreyer, and Viktor Vafeiadis. 2016. Lightweight verification of separate compilation. In *Proceedings of the 43rd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL'16)*. Association for Computing Machinery, New York, NY, 178–190. DOI: <http://doi.org/10.1145/2837614.2837642>
- [39] Leslie Lamport and Fred B. Schneider. 1985. Formal foundation for specification and verification. In *Distributed Systems: Methods and Tools for Specification, an Advanced Course*. Springer-Verlag, 203–285. DOI: http://doi.org/10.1007/3-540-15216-4_15
- [40] Chris Lattner. 2011. What Every C Programmer Should Know about Undefined Behavior #1/3. LLVM Project Blog. (May 2011). Retrieved from <http://blog.llvm.org/2011/05/what-every-c-programmer-should-know.html>.
- [41] Xavier Leroy. 2009. Formal verification of a realistic compiler. *Commun. ACM* 52, 7 (2009), 107–115. DOI: <http://doi.org/10.1145/1538788.1538814>
- [42] Xavier Leroy. 2009. A formally verified compiler back-end. *J. Autom. Reason.* 43, 4 (2009), 363–446. <http://doi.org/10.1007/s10817-009-9155-4>
- [43] Xavier Leroy. 2017. The formal verification of compilers (DeepSpec Summer School 2017). Retrieved from <https://xavierleroy.org/courses/DSSS-2017/>.
- [44] Isabella Mastroeni and Michele Pasqua. 2018. Verifying bounded subset-closed hyperproperties. In *Proceedings of the International Static Analysis Symposium*. 263–283. Retrieved from <https://iris.univr.it/retrieve/handle/11562/990895/120109/MastroeniPasqua.pdf>.
- [45] Jacob Matthews and Robert Bruce Findler. 2007. Operational semantics for multi-language programs. *SIGPLAN Not.* 42, 1 (Jan. 2007), 3–10. DOI: <http://doi.org/10.1145/1190215.1190220>
- [46] John McCarthy and James Painter. 1967. Correctness of a compiler for arithmetic expressions. *Math. Asp. Comput. Sci.* 1 19 of *Proceedings of Symposia in Applied Mathematics* (1967), 33–41. Retrieved from <http://jmc.stanford.edu/articles/mcpain/mcpain.pdf>.
- [47] A. Melton, D. A. Schmidt, and G. E. Strecker. 1986. Galois connections and computer science applications. In *Proceedings of a Tutorial and Workshop on Category Theory and Computer Programming*. 299–312. Retrieved from <http://dl.acm.org/citation.cfm?id=20081.20099>.
- [48] R. Milner. 1982. *A Calculus of Communicating Systems*. Springer-Verlag, Berlin.
- [49] Robin Milner and Richard Weyhrauch. 1972. Proving compiler correctness in a mechanized logic. In *Proceedings of 7th Annual Machine Intelligence Workshop, volume 7 of Machine Intelligence*. 51–72. Retrieved from <http://www.cs.umd.edu/~hjs/pubs/compilers/archive/mi72-mil-wey.pdf>.
- [50] F. Lockwood Morris. 1973. Advice on structuring compilers and proving them correct. In *Proceedings of the ACM Symposium on Principles of Programming Languages*, Patrick C. Fischer and Jeffrey D. Ullman (Eds.), 144–152. DOI: <http://doi.org/10.1145/512927.512941>
- [51] Eric Mullen, Daryl Zuniga, Zachary Tatlock, and Dan Grossman. 2016. Verified peephole optimizations for CompCert. In *Proceedings of the 37th ACM SIGPLAN Conference on Programming Language Design and Implementation*. 448–461. DOI: <http://doi.org/10.1145/2908080.2908109>
- [52] Toby C. Murray, Robert Sison, Edward Pierzchalski, and Christine Rizkallah. 2016. Compositional verification and refinement of concurrent value-dependent noninterference. In *Proceedings of the IEEE 29th Computer Security Foundations Symposium*. IEEE Computer Society, 417–431. DOI: <http://doi.org/10.1109/CSF.2016.36>

- [53] Kedar S. Namjoshi and Lucas M. Tabajara. 2020. Witnessing secure compilation. In *Proceedings of the International Conference on Verification, Model Checking, and Abstract Interpretation*. Springer, 1–22.
- [54] David A. Naumann. 1998. A categorical model for higher order imperative programming. *Math. Struct. Comput. Sci.* 8, 4 (1998), 351–399. Retrieved from <https://www.cs.stevens.edu/~naumann/pub/cmho.ps>.
- [55] David A. Naumann and Minh Ngo. 2019. Whither specifications as programs. In *Proceedings of the International Symposium on Unifying Theories of Programming*. Springer, 39–61. Retrieved from <https://arxiv.org/abs/1906.03557>.
- [56] Georg Neis, Chung-Kil Hur, Jan-Oliver Kaiser, Craig McLaughlin, Derek Dreyer, and Viktor Vafeiadis. 2015. Pilsner: A compositionally verified compiler for a higher-order imperative language. In *Proceedings of the 20th ACM SIGPLAN International Conference on Functional Programming*. 166–178. DOI: <http://doi.org/10.1145/2784731.2784764>
- [57] Max New, William J. Bowman, and Amal Ahmed. 2016. Fully abstract compilation via universal embedding. In *Proceedings of the ACM SIGPLAN International Conference on Functional Programming*.
- [58] Michele Pasqua and Isabella Mastroeni. 2017. On topologies for (hyper)properties. In *Joint Proceedings of the 18th Italian Conference on Theoretical Computer Science and the 32nd Italian Conference on Computational Logic co-located with the IEEE International Workshop on Measurements and Networking IEEE M&N’17*. (CEUR Workshop Proceedings), Dario Della Monica, Aniello Murano, Sasha Rubin, and Luigi Sauro (Eds.), Vol. 1949. 150–161. Retrieved from <http://ceur-ws.org/Vol-1949/ICTCSpaper13.pdf>.
- [59] Marco Patrignani. 2015. *The Tome of Secure Compilation: Fully Abstract Compilation to Protected Modules Architectures*. Ph.D. Dissertation. KU Leuven, Leuven, Belgium. Retrieved from <https://lirias.kuleuven.be/bitstream/123456789/494704/1/thesis.pdf>.
- [60] Marco Patrignani. 2020. Why Should Anyone use Colours? or, Syntax Highlighting Beyond Code Snippets. arXiv:cs.SE/2001.11334 (2020).
- [61] Marco Patrignani, Amal Ahmed, and Dave Clarke. 2019. Formal approaches to secure compilation: A survey of fully abstract compilation and related work. *Comput. Surv.* (2019). Retrieved from http://theory.stanford.edu/~mp/mp/Publications_files/main-full.pdf.
- [62] Marco Patrignani and Dave Clarke. 2015. Fully abstract trace semantics for protected module architectures. *Comput. Lang. Syst. Struct.* 42 (2015), 22–45. DOI: <http://doi.org/10.1016/j.cl.2015.03.002>
- [63] Marco Patrignani and Deepak Garg. 2017. Secure compilation and hyperproperty preservation. In *Proceedings of the 30th IEEE Computer Security Foundations Symposium*. 392–404. DOI: <http://doi.org/10.1109/CSF.2017.13>
- [64] Marco Patrignani and Deepak Garg. 2019. Robustly safe compilation. In *Proceedings of the 28th European Symposium on Programming: Programming Languages and Systems (ESOP’19)*. Retrieved from <https://arxiv.org/abs/1804.00489>.
- [65] Daniel Patterson and Amal Ahmed. 2019. The next 700 compiler correctness theorems (functional pearl). *Proc. ACM Program. Lang.* 3, ICFP (July 2019). DOI: <http://doi.org/10.1145/3341689>
- [66] James T. Perconti and Amal Ahmed. 2014. Verifying an open compiler using multi-language semantics. In *Proceedings of the 23rd European Symposium on Programming (Lecture Notes in Computer Science)*, Vol. 8410. Springer, 128–148. Retrieved from <http://www.ccs.neu.edu/home/jtpercon/multilang-verify.pdf>.
- [67] Tahina Ramananandro, Zhong Shao, Shu-Chun Weng, Jérémie Koenig, and Yuchen Fu. 2015. A compositional semantics for verified separate compilation and linking. In *Proceedings of the Conference on Certified Programs and Proofs*. 3–14. DOI: <http://doi.org/10.1145/2676724.2693167>
- [68] John Regehr. 2010. A Guide to Undefined Behavior in C and C++, Part 3. Embedded in Academia blog. (July 2010). Retrieved from <https://blog.regehr.org/archives/232>.
- [69] Andrei Sabelfeld and David Sands. 2005. Dimensions and principles of declassification. In *Proceedings of the Computer Security Foundations 18th Workshop*. 255–269. Retrieved from <http://www.cse.chalmers.se/~dave/papers/sabelfeld-sands-CSFW05.pdf>.
- [70] Amr Sabry and Philip Wadler. 1997. A reflection on call-by-value. *ACM Trans. Program. Lang. Syst.* 19, 6 (1997), 916–941.
- [71] Gabriel Scherer, Max S. New, Nick Rioux, and Amal Ahmed. 2018. FabULous interoperability for ML and a linear language. In *Proceedings of the 21st International Conference on Foundations of Software Science and Computation Structures, Held as Part of the European Joint Conferences on Theory and Practice of Software (Lecture Notes in Computer Science)*, Christel Baier and Ugo Dal Lago (Eds.), Vol. 10803. Springer, 146–162. DOI: http://doi.org/10.1007/978-3-319-89366-2_8
- [72] Fred B. Schneider. 2000. Enforceable security policies. *ACM Trans. Inf. Syst. Secur.* 3, 1 (2000), 30–50. Retrieved from <http://www.cs.cornell.edu/fbs/publications/EnfSecPols.pdf>.
- [73] Jaroslav Sevcík, Viktor Vafeiadis, Francesco Zappa Nardelli, Suresh Jagannathan, and Peter Sewell. 2013. CompCertTSO: A verified compiler for relaxed-memory concurrency. *J. ACM* 60, 3 (2013), 22:1–22:50. DOI: <http://doi.org/10.1145/2487241.2487248>

- [74] Robert Sison and Toby Murray. 2019. Verifying that a compiler preserves concurrent value-dependent information-flow security. In *Proceedings of the 10th International Conference on Interactive Theorem Proving (LIPIcs)*, John Harrison, John O’Leary, and Andrew Tolmach (Eds.), Vol. 141. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 27:1–27:19. DOI: <http://doi.org/10.4230/LIPIcs.ITP.2019.27>
- [75] Lau Skorstengaard, Dominique Devriese, and Lars Birkedal. 2019. StkTokens: Enforcing well-bracketed control flow and stack encapsulation using linear capabilities. *Proc. ACM Program. Lang.* 3, POPL (Jan. 2019), 19:1–19:28.
- [76] Gordon Stewart, Lennart Beringer, Santiago Cuellar, and Andrew W. Appel. 2015. Compositional CompCert. In *Proceedings of the 42nd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*. 275–287. DOI: <http://doi.org/10.1145/2676726.2676985>
- [77] Yong Kiam Tan, Magnus O. Myreen, Ramana Kumar, Anthony Fox, Scott Owens, and Michael Norrish. 2019. The verified CakeML compiler backend. *J. Funct. Program.* 29 (2019). DOI: <http://doi.org/10.1017/S0956796818000229>
- [78] Xi Wang, Haogang Chen, Alvin Cheung, Zhihao Jia, Nickolai Zeldovich, and M. Frans Kaashoek. 2012. Undefined behavior: What happened to my code? In *Proceedings of the Asia-Pacific Workshop on Systems*. 9. DOI: <http://doi.org/10.1145/2349896.2349905>
- [79] Xi Wang, Nickolai Zeldovich, M. Frans Kaashoek, and Armando Solar-Lezama. 2013. Towards optimization-safe systems: Analyzing the impact of undefined behavior. In *Proceedings of the ACM SIGOPS 24th Symposium on Operating Systems Principles*. 260–275. DOI: <http://doi.org/10.1145/2517349.2522728>
- [80] Yuting Wang, Pierre Wilke, and Zhong Shao. 2019. An abstract stack based approach to verified compositional compilation to machine code. *Proc. ACM Program. Lang.* 3, POPL (2019), 62:1–62:30. DOI: <http://doi.org/10.1145/3290375>
- [81] Li-yao Xia, Yannick Zakowski, Paul He, Chung-Kil Hur, Gregory Malecha, Benjamin C. Pierce, and Steve Zdancewic. 2020. Interaction trees: Representing recursive and impure programs in Coq. *Proc. ACM Program. Lang.* 4, POPL (2020), 51:1–51:32. DOI: <http://doi.org/10.1145/3371119>
- [82] Aris Zakinthinos and E. Stewart Lee. 1997. A general theory of security properties. In *Proceedings of the IEEE Symposium on Security and Privacy (Oakland S&P’97)*. 94–102. DOI: <http://doi.org/10.1109/SECPRI.1997.601322>
- [83] Jianzhou Zhao, Santosh Nagarakatte, Milo M. K. Martin, and Steve Zdancewic. 2012. Formalizing the LLVM intermediate representation for verified program transformations. In *Proceedings of the 39th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*. 427–440. Retrieved from <http://www.cis.upenn.edu/~stevez/papers/ZNMZ12.pdf>.

Received May 2020; revised April 2021; accepted April 2021