

Secure Composition of Robust and Optimising Compilers

MATTHIS KRUSE, CISPA Helmholtz Center for Information Security and Saarland University, Germany

MICHAEL BACKES, CISPA Helmholtz Center for Information Security, Germany

MARCO PATRIGNANI, University of Trento, Italy

To ensure that secure applications do not leak their secrets, they are required to uphold several security properties such as spatial and temporal memory safety as well as cryptographic constant time. Existing work shows how to enforce these properties individually, in an architecture-independent way, by using secure compiler passes that each focus on an individual property. Unfortunately, given two secure compiler passes that each preserve a possibly different security property, it is unclear what kind of security property is preserved by the composition of those secure compiler passes. This paper is the first to study what security properties are preserved across the composition of different secure compiler passes. Starting from a general theory of property composition for security-relevant properties (such as the aforementioned ones), this paper formalises a theory of composition of secure compilers. Then, it showcases this theory a secure multi-pass compiler that preserves the aforementioned security-relevant properties. Crucially, this paper derives the security of the multi-pass compiler from the composition of the security properties preserved by its individual passes, which include security-preserving as well as optimisation passes. From an engineering perspective, this is the desirable approach to building secure compilers.

This paper uses syntax highlighting accessible to both colourblind and black & white readers.

CCS Concepts: • **Security and privacy** → **Formal security models.**

Additional Key Words and Phrases: Memory-safety, Secure Compilation, Privacy

ACM Reference Format:

Matthis Kruse, Michael Backes, and Marco Patrignani. 2024. Secure Composition of Robust and Optimising Compilers. In . ACM, New York, NY, USA, 28 pages. <https://doi.org/XXXXXXXX.XXXXXXX>

1 INTRODUCTION

Memory Safety (MS) is a security property obtained by composing Spatial Memory Safety (SMS), which ensures array accesses are all within bounds, and Temporal Memory Safety (TMS), which ensures pointers are only used when they are valid [Akritidis et al. 2009; Azevedo de Amorim et al. 2018; Jim et al. 2002; Michael et al. 2023; Nagarakatte et al. 2009, 2010; Nacula et al. 2005]. Cryptographic Constant Time (CCT) is a security property that ensures sensitive data is not leaked via timing side-channels [Kocher 1996]. Together, SMS, TMS and Strict Cryptographic Constant Time (sCCT), an enforceable overapproximation of CCT, yield Memory Safety and Strict Cryptographic Constant Time (MS+sCCT), which is the gold standard of security properties for secure applications. Programs attaining MS+sCCT do not leak sensitive data either through erroneous memory accesses, nor through timing side-channels. As discussed in Example 1.1, these security properties can be enforced by compiler passes [Almeida et al. 2017; Bond et al. 2017], to ensure programmers need not be aware of the architectural details of where their code will run.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

POPL '24, January 17-19, 2024, London, UK

© 2024 Association for Computing Machinery.

ACM ISBN 978-1-4503-XXXX-X/18/06...\$15.00

<https://doi.org/XXXXXXXX.XXXXXXX>

50 *Example 1.1 (strncpy).* Consider the C function `strncpy` that copies a null-terminated string `src`
 51 into `dst` up to a length of `n`. This function is subject to a subtle SMS vulnerability: The bounds
 52 check `i < n` should happen *before* the access to memory location `x[i]`: otherwise the memory
 53 location past the last element will be leaked to an attacker.

```
54 void strncpy(size_t n, char *dst, char *src) {
55     for(size_t i = 0; src[i] != '\0' && i < n; ++i) {
56         dst[i] = src[i];
57     }
58 }
```

59 To prevent this vulnerability, one can use a compilation pass that enforces SMS, such as Soft-
 60 bounds [Nagarakatte et al. 2009] or BaggyBounds [Akritidis et al. 2009].

61 Because of timing attacks, fixing SMS is not enough to make `strncpy` secure. In fact, the loop
 62 can terminate early, as soon as the string-terminating character `'\0'` is encountered, thus making
 63 program execution time proportional to the length of the array pointed by `src`. Also in this case
 64 there exist compiler passes that can rewrite such programs into CCT ones [Cauligi et al. 2019].

65 Alas, code is not run in isolation, so a malicious attacker could supply code that intracts with
 66 `strncpy` and trigger a violation of either MS or CCT by calling `strncpy` with an argument for
 67 `src` that points to uninitialised memory. This would, in turn, triggering a series of reads from
 68 uninitialised memory, which is an immediate MS violation with devastating real-world conse-
 69 quences [Microsoft 2010a,b,c, 2015; VMWare 2023].

70 Robust compilers [Abate et al. 2019] are a form of secure compilers that preserve security
 71 properties even in the presence of arbitrary attackers interacting with compiled code. Thus, robust
 72 compilers can be used to prevent vulnerabilities resulting from uninitialised memory (as well as
 73 many other ones), e.g., by targeting capability-based languages such as CHERI [Woodruff et al.
 74 2014], Arm Morello [Arm 2022], or MSWasm [Michael et al. 2023], where the compiler relies on
 75 capabilities to check that pointers are always initialised.

76 Unfortunately, given secure compiler passes that each preserve a possibly different security
 77 property, there is no way to tell what kind of security property will the composition of those secure
 78 compilers preserve. Worse, without a framework for composing secure compiler passes, it is not
 79 possible to enable separation of concerns, e.g., to have a secure compilation pass that ensures MS
 80 that is developed independently of another secure pass for CCT, that is developed independently
 81 of other passes, such as optimisations.

82 This paper introduces a framework for reasoning about the composition of secure and optimising
 83 compiler passes akin to those of Example 1.1 and it showcases the power of this framework by
 84 instantiating it on a multi-pass compilation chain. To this end, this paper first discusses how to
 85 compose security properties, such as TMS and SMS into MS, and then adding sCCT to the mix
 86 to obtain MS+sCCT. Then, this paper defines compiler composition and formalises that given
 87 two passes that securely preserve two (possibly distinct) properties, their composition securely
 88 preserves the composition of those properties. The paper then defines several secure compiler passes,
 89 where each is either preserving a different security property (TMS, SMS, sCCT) or performing a
 90 security-preserving optimisation, (e.g., applying Constant Folding (CF) or Dead Code Elimination
 91 (DCE)). Finally, this paper shows that composing these secure compiler passes into a multi-pass
 92 compilation chain results in the end-to-end preservation of MS+sCCT. Crucially, this paper derives
 93 the security of the multi-pass compiler from the composition of the security properties preserved by
 94 its individual passes. This result showcases how the framework allows the kind of formal security
 95 reasoning that compiler writers already want (and already do), obtaining precise, compositional
 96 security reasoning while providing minimal (and modular) proof effort.

In summary, this paper makes the following contributions:

- This paper formalises security properties (Section 3) that are of interest for real-world compiler writers, namely TMS, SMS and CCT (as identified by the plethora of work enforcing such properties individually [Akritidis et al. 2009; Almeida et al. 2017; Bond et al. 2017; Cauligi et al. 2019; Dhumbumroong and Piromsopa 2020; Jung et al. 2021; Kuepper et al. 2023; Nagarakatte et al. 2009, 2010; Nam et al. 2019; Shankaranarayana et al. 2023; Younan et al. 2010; Zhou et al. 2023]). Starting from ways to formalise those properties individually, this paper shows how to compose their formalisation. The resulting security property is MS+sCCT, i.e., the gold standard of security properties for secure programs [LeMay et al. 2021].
- This paper takes the secure compilation framework of [Abate et al. 2019] and extends it to reason about the security of all different known forms of compiler composition (Section 4). For this, this paper studies sequential compiler composition as well as compilers with multiple input languages or multiple output ones, as used in existing compilation chains. This paper proves that starting from two compilers that preserve two (possibly distinct) properties, their composition preserves the intersection of those properties. Finally, this paper proves that the order of composition of sequential compiler passes is irrelevant for the resulting security. This is crucial for reordering optimisation passes and thus generating secure and efficient code.
- This paper presents a case-study showcasing the conjunction of the previous contributions (Sections 5 and 6). To this end, it presents a compilation chain consisting of several passes that ultimately preserves MS+sCCT by means of composing the individual, secure passes concerning TMS, SMS, and sCCT, respectively. Furthermore, the chain includes two optimisation passes: One performs DCE and the other CF. The formalisation of this case study showcases the power of the presented framework: The divide-and-conquer approach to software engineering is a viable strategy even for the development of secure compilers.
- The key contributions of this paper are formalised in the Coq proof assistant and the paper indicates this with \clubsuit .

This paper starts by introducing relevant notions of security properties and secure compilation (Section 2), and discusses related work (Section 7) before concluding (Section 8).

Open Source & Technical Report. A technical report with the omitted formal details, lemmas and proofs, as well as the Coq formalisation are available as supplementary material.

2 BACKGROUND: SECURITY PROPERTIES AND SECURE COMPILERS

To introduce the security argument of this paper, this section first presents the concepts of (security) properties, of their satisfaction, and of their robust satisfaction (i.e., satisfaction in the presence of an active attacker; Section 2.1). Then, borrowing from existing work [Abate et al. 2021a, 2019], the section introduces secure compilers as compilers that preserve robust property satisfaction (Section 2.2).

2.1 Properties and (Robust) Satisfaction

This paper employs the security model where programs are written in a language whose semantics emits events a . Events include security-relevant actions (e.g., reading from and writing to memory, as detailed in Section 3) and the unobservable event ε . As programs execute, their emitted events are concatenated in traces \bar{a} , which serve as the description of the behaviour of a program.¹

¹Throughout the paper, sequences are indicated with an overbar (i.e., \bar{a}), empty sequences with $[\cdot]$, and concatenation of sequences \bar{a}_1, \bar{a}_2 as $\bar{a}_1 \cdot \bar{a}_2$. Prepending elements to sequences uses the same notation: $a \cdot \bar{a}$.

Properties π are sets of traces of admissible program behaviours, ascribing what said property considers valid. The set of all properties can be partitioned into different *classes* (\mathbb{C}), i.e., safety, liveness, and neither safety nor liveness [Clarkson and Schneider 2008]. A class is simply a set of properties and for the class of safety properties, it is decidable whether a trace satisfies a safety property with just a finite trace prefix. As an example, consider a trace describing an interaction with a memory where the deallocation of an address l precedes a read at that address in memory: Dealloc l · Read l 1729 · This program behaviour is insecure with respect to a canonical notion of (temporal) memory safety dictating no use-after-frees of pointers [Azevedo de Amorim et al. 2018; Nagarakatte et al. 2010], because it reads from a memory location that was freed already. The previous finite trace prefix is enough to decide that the trace does not satisfy TMS and there is no way to append events to this prefix which would result in the trace being admissible. In the following, the execution of a whole program w that terminates in state r according to the language semantics and produces trace \bar{a} is written as $w \xRightarrow{\bar{a}} r$. With this, property satisfaction is defined as follows: whole programs w satisfy a property π iff w yields a trace \bar{a} such that \bar{a} satisfies π (Definition 2.1).

Definition 2.1 (Property Satisfaction). $\vdash p : \pi \stackrel{\text{def}}{=} \text{if } w \xRightarrow{\bar{a}} r, \text{ then } \bar{a} \in \pi.$

Property satisfaction is defined on whole programs, i.e., programs without missing definitions. Thus, from a security perspective, this considers only a passive attacker model, where the attacker observes the execution and, e.g., retrieves secrets from that. To consider a stronger model similarly to what existing work does [Abate et al. 2021a, 2019; Backes et al. 2014; Bengtson et al. 2011; Fournet et al. 2007; Gordon and Jeffrey 2003; Maffei et al. 2008; Michael et al. 2023; Sammler et al. 2019; Swasey et al. 2017], the concept of satisfaction can be extended with *robustness*. Robust satisfaction considers partial programs p , i.e., components with missing imports, which cannot run until said imports are fulfilled. To remedy this, *linking* takes two partial programs p_1, p_2 and produces a whole program w , i.e., $\text{link}(p_1; p_2) = w$. As typically done in works that consider the execution of partial programs [Abate et al. 2019; Ahmed and Blume 2011; Bowman and Ahmed 2015; Devriese et al. 2017a,b; El-Korashy et al. 2021; Patrignani and Garg 2021; Patterson and Ahmed 2017; Van Strydonck et al. 2019], this paper assumes that whole programs are the result of linking partial programs referred to as *context* (ctx) and *component* (comp). The context is an arbitrary program and thus has the role of an *attacker* that can interact with the component by means of whatever features the programming language has, and the component is what is security-relevant. With this, Definition 2.1 (Property Satisfaction) can be extended as follows: for components p to robustly satisfy a property π , take an attacker context C and link it with p , the resulting whole program must satisfy π .

Definition 2.2 (Robust Satisfaction). $\vdash_R p : \pi \stackrel{\text{def}}{=} \forall C, \text{ if } \text{link}(C; p) = w, \text{ then } \vdash w : \pi.$

Example 2.3 (Double Free in Bluetooth Subsystem). Consider CVE-2021-3564 [BlockSec 2021], one of many submissions for a double-free vulnerability. The vulnerability arises due to a race condition where the context-level function `hci_cmd_work` was not expected to behave maliciously, since it resides in the same source-code repository where the vulnerability occurs. Nevertheless, the component-level code of `hci_dev_do_open` is linked with `hci_cmd_work` and does not atomically check whether a pointer has been freed already: Therefore, `hci_dev_do_open` does not satisfy the no-double-frees property robustly, since there is an implementation for `hci_cmd_work` that leads to a violation of that property when linked with `hci_dev_do_open`.

2.2 Secure Compilers

A *compiler* ($\gamma_{\mathbf{L}}^{\mathbf{T}}$) translates syntactic descriptions of programs from a *source* (\mathbf{L}) into a *target* (\mathbf{T}) programming language. This translation is considered *correct* if it is semantics-preserving. That is, for a whole program w , the compiler should relate the \mathbf{L} semantics of w with the semantics of \mathbf{T} of the compiled counterpart of w in such a way that they are „compatible“. Unfortunately, correct compilers may be insecure compilers [Abadi 1999a; Ahmed et al. 2018; Kennedy 2006; Patrignani et al. 2019] and programs translated by insecure compilers can violate security properties that the programmer assumes to hold. To define when a compiler is secure, this paper uses the robust compilation framework [Abate et al. 2019], which the following definition summarises.

For compilers $\gamma_{\mathbf{L}}^{\mathbf{T}}$ to robustly preserve a class of properties \mathbb{C} , if for any property π of that class \mathbb{C} and programs p written in \mathbf{L} where p robustly satisfies π , then the compilation of p , $\gamma_{\mathbf{L}}^{\mathbf{T}}(p)$, must robustly satisfy π .

Definition 2.4 (Robust Compilation). $\vdash \gamma_{\mathbf{L}}^{\mathbf{T}} : \mathbb{C} \stackrel{\text{def}}{=} \forall (\pi \in \mathbb{C}) (p \in \mathbf{L}_{\text{tms}}), \text{if } \vdash_R p : \pi, \text{ then } \vdash_R \gamma_{\mathbf{L}}^{\mathbf{T}}(p) : \pi.$

Note that a class of properties \mathbb{C} can represent just one property π by lifting [Clarkson and Schneider 2008] that property to sets of properties, i.e., use the powerset of π instead of π itself. Because of this, this paper writes $\vdash \gamma_{\mathbf{L}}^{\mathbf{T}} : \pi$, even though π is a property and not a class.

Example 2.5 (Types). Suppose \mathbf{L} is a statically-typed language similar to C and \mathbf{T} is dynamically typed, where both share the same syntax up to dynamic type checks. Consider the following \mathbf{L} component and its compiled version below.

```
fn foo (char * x, int n) := ifz valid_ptr(x, n, sizeof(char)) then x[0] else - 1
fn foo (      x,      n) := ifz valid_ptr(x, n, sizeof(char)) then x[0] else - 1
```

While the compiler emits code that may look correct, the generated code does not check that the provided argument is of the right type. Even though the pointer x is checked for validity, the context $\text{foo}((\text{int}^*)y, 1)$ is able to provoke a read out of bounds. Suppose the component transferred control to the context and passed ownership of a char pointer y sized 1 cells, the context can now call the component again, casting this buffer to an int^* prior to that call. The pointer is valid for one char -sized memory cell, as expected, but the actual read operation now returns $\text{sizeof}(\text{int})$ many bytes instead of just $\text{sizeof}(\text{char})$ many. Thus, even if foo may have been robust with respect to the SMS, its compiled counterpart is not and therefore the compiler fails to attain Definition 2.4.

3 SECURITY PROPERTIES: FORMALISATION, ENFORCEMENT AND COMPOSITION

This section introduces a trace model and uses it to define the key properties of interest for this paper: TMS, SMS, MS, and sCCT (Section 3.1). These properties are of practical importance (as mentioned in Section 1) and also of interest in the case study (Sections 5 and 6) this paper presents later. Lastly, for each of the key properties, this section introduces corresponding monitors (Section 3.2) that check them.

3.1 Specification Trace Model

(Security Tag) $\sigma ::= \blacksquare \mid \blacklozenge$ (Control Tag) $t ::= \text{ctx} \mid \text{comp}$ (Event) $a ::= \varepsilon \mid \frac{1}{2} \mid a_b; t; \sigma$
 (Pre-event) $a_b ::= \text{Alloc } l \ n \mid \text{Dealloc } l \mid \text{Use } l \ n \mid \text{Branch } n \mid \text{Binop } n$

The specification trace model defines events as either the empty event (ε), a crash ($\frac{1}{2}$), or as tuples consisting of a pre-event, a control-tag, and a security-tag. The purpose of the model is to define key security properties of interest, such as MS or a stricter variant of cryptographic constant time. To this end, security-tags indicate whether an event contains sensitive information (\blacksquare) or not

(\blacksquare), while control-tags state whether the context (ctx) or the component (comp) are responsible for emitting the event. The latter is necessary to be able to ignore actions done by a spurious context that, e.g., immediately deallocates a memory location twice, thus violating TMS [Nagarakatte et al. 2010]. Lastly, pre-events describe the actual kind of event that happened. One such kind is the allocation event (Alloc l n) that fires whenever a program claims n cells of memory and stores them at address l . Dually, deallocation (Dealloc l) announces that the object at location l is freed. These two events alone are enough to provide a partial description of TMS by requiring that, e.g., there is only one deallocation event that carries a location l . To be able to express SMS, there is also an event to describe reads from and writes to memory (Use l n). Finally, for cryptographic code, there is a general guideline that secrets must not be visible on a trace. Moreover, an instruction whose timing is data-dependent must not have a secret as an operand. Typical operations with data-dependent timing are branches and certain binary operations, such as division². Both operations are also modelled in the specification trace model (Branch n and Binop n).

3.1.1 Temporal Memory Safety. TMS is a safety property that describes that an unallocated object must not be used in any way. Moreover, the property requires that all allocated objects must be deallocated at some point.

Definition 3.1 (TMS).

$$\text{tms} := \left\{ \bar{a} \left| \begin{array}{l} \text{Alloc } l \ n; t; \sigma \leq_{\bar{a}} \text{Dealloc } l; t; \sigma' \\ \text{Use } l \ n; t; \sigma \leq_{\bar{a}} \text{Dealloc } l; t; \sigma' \\ \text{if Alloc } l \ n; t; \sigma \text{ in } \bar{a} \text{ then Dealloc } l; t; \sigma' \text{ in } \bar{a} \\ \text{at most one Dealloc } l; t; \sigma \text{ in } \bar{a} \\ \text{at most one Alloc } l \ n; t; \sigma \text{ in } \bar{a} \end{array} \right. \right\}$$

Hereby, the notation $a_1 \leq_{\bar{a}} a_2$ means that if a_1 is in \bar{a} and if a_2 is in \bar{a} , then a_1 appears before a_2 .

3.1.2 Spatial Memory Safety. SMS prohibits out-of-bounds accesses:

Definition 3.2 (SMS).

$$\text{sms} := \{ \bar{a} \mid \text{If Alloc } l \ n; t; \sigma \leq_{\bar{a}} \text{Use } l \ m; t; \sigma', \text{ then } m < n \}$$

3.1.3 Memory Safety. Full MS (similar to earlier work [Jim et al. 2002; Michael et al. 2023; Nagarakatte et al. 2009, 2010; Necula et al. 2005]) is then described as the conjunction of Definitions 3.1 and 3.2. Note, however, that this definition says nothing about memory-safety issues introduced by side-channels, such as speculation.

Definition 3.3 (MS).

$$\text{ms} := \text{tms} \cap \text{sms}$$

3.1.4 Strict Cryptographic Constant Time. CCT is a hypersafety property [Barthe et al. 2018] and, thus, difficult to check with monitors. This is because, intuitively, hypersafety properties can relate multiple execution traces with each other, but monitors work on a single execution. To sidestep this issue, this section defines the property sCCT, a stricter variant of CCT that enforces the policy that no secret appears on a trace (inspired by earlier work [Almeida et al. 2017]).

Definition 3.4 (sCCT).

$$\text{scct} := \{ \bar{a} \mid \bar{a} = [\cdot] \quad \text{or} \quad \bar{a} = a_b; t; \blacksquare \cdot \bar{a}' \wedge \bar{a}' \in \text{scct} \}$$

²This is highly architecture-dependent, but division is an operation that serves as a classic example for a data-dependent timing instruction, e.g., [Arm 2020, p. 755].

3.1.5 *Memory Safe, Strict Cryptographic Constant Time.* The combination of MS and sCCT is the intersection of these properties, MS+sCCT. Since MS and sCCT are just sets of traces that, intuitively, contain all program behaviors that follow a security policy, the intersection of them contains all program behaviors that follow both security policies, i.e., it entails all program behaviours that are both MS and sCCT.

Definition 3.5 (MS and sCCT).

$$\text{mssct} := \text{ms} \cap \text{sct}$$

3.2 Monitors

Monitors enforce safety properties by accepting or rejecting traces, i.e., if it rejects a trace, the trace does not satisfy the property the monitor checks. Since reasoning on monitors is easier than directly on just traces, this section presents a monitor for each of the previously shown safety properties (Section 3.1). To lessen the burden when proving that a monitor accepts the trace of a program execution, each monitor uses a custom trace model that contains only the relevant information related to the property the monitor checks. To go from specification traces \bar{a} to monitor-level traces \bar{a} , each property π has an associated event agreement relation $a \cong_{\pi} \bar{a}$. Figure 1 shows how the event agreement is lifted to traces. The trace agreement is the same for all properties π up to the

$$\boxed{\bar{a} \cong_{\pi}^* \bar{a}} \text{ „Specification-level trace } \bar{a} \text{ agrees with monitor-level trace } \bar{a} \text{ with respect to property } \pi\text{.”}$$

(traceagree-empty)	(traceagree-ign-L)	(traceagree-ign-R)	(traceagree-cons)
$[\cdot] \cong_{\pi}^* [\cdot]$	$\bar{a} \cong_{\pi}^* \bar{a}$	$\bar{a} \cong_{\pi}^* \bar{a}$	$a \cong_{\pi} \mathbf{a} \quad \bar{a} \cong_{\pi}^* \bar{a}$
	$\varepsilon \cdot \bar{a} \cong_{\pi}^* \bar{a}$	$\bar{a} \cong_{\pi}^* \varepsilon \cdot \bar{a}$	$a \cdot \bar{a} \cong_{\pi}^* \mathbf{a} \cdot \bar{a}$

Fig. 1. Trace-Agreement relation that equates specification-level traces with monitor-level traces.

use of the event agreement in Rule `traceagree-cons`. With agreements, this section defines monitor satisfaction for traces and then it proves that monitor satisfaction implies property satisfaction. To this end, monitor satisfaction is defined as follows. A specification trace \bar{a} monitor-satisfies property π iff there exists a (final) monitor state T and an abstract trace \bar{a} such that the specification trace \bar{a} agrees with abstract trace \bar{a} and the initial monitor³ can step to the (final) monitor state T with abstract trace \bar{a} .

Definition 3.6 (Monitor Satisfaction). $\vdash_{\text{mon}} \bar{a} : \pi \stackrel{\text{def}}{=} \exists \bar{a} T, \bar{a} \cong_{\pi}^* \bar{a} \text{ and } \vdash \emptyset \xrightarrow{\bar{a}}^* T$.

3.2.1 Monitor for TMS.

(Abstract Store) $T_{\text{TMS}} ::= \{\text{allocated} : L \times t, \text{freed} : L \times t\} \quad \emptyset ::= \{\text{allocated} : \emptyset, \text{freed} : \emptyset\}$

(Abstract Events) $\mathbf{a} ::= \varepsilon \mid \text{Alloc } l \ t \mid \text{Dealloc } l \ t \mid \text{Use } l \ t \mid \downarrow$

$\vdash T_{\text{TMS}} \xrightarrow{\mathbf{a}} T_{\text{TMS}'}$ „Monitor T_{TMS} does one step to $T_{\text{TMS}'}$ given event \mathbf{a} ”

³In this paper, for all monitors, the initial monitor state is denoted as \emptyset .

$$\begin{array}{c}
\text{344} \\
\text{345} \\
\text{346} \\
\text{347} \\
\text{348} \\
\text{349} \\
\text{350} \\
\text{351} \\
\text{352} \\
\text{353} \\
\text{354} \\
\text{355} \\
\text{356}
\end{array}
\frac{
\frac{
\frac{
(l; t) \in T_{TMS}.allocated \quad (l; t) \notin T_{TMS}.freed
}{\vdash T_{TMS} \xrightarrow{\text{Use } l \ t} T_{TMS}}
\quad (tms\text{-use})
}{
\frac{
(l; t) \notin T_{TMS}.allocated \quad (l; t) \notin T_{TMS}.freed
}{
T_{TMS}' = \{allocated : T_{TMS}.allocated \cup \{(l; t)\}, freed : T_{TMS}.freed\}
}
\vdash T_{TMS} \xrightarrow{\text{Alloc } l \ t} T_{TMS}'
}
\quad (tms\text{-alloc})
}{
\frac{
(l; t) \in T_{TMS}.allocated \quad (l; t) \notin T_{TMS}.freed
}{
T_{TMS}' = \{allocated : T_{TMS}.allocated \setminus \{(l; t)\}, freed : T_{TMS}.freed \cup \{(l; t)\}\}
}
\vdash T_{TMS} \xrightarrow{\text{Dealloc } l \ t} T_{TMS}'
}
\quad (tms\text{-dealloc})
}
\vdash T_{TMS} \xrightarrow{\text{Dealloc } l \ t} T_{TMS}'$$

For TMS, the state of the monitor is a record with two sets keeping track of allocated and deallocated locations. Rule `tms-use` simply requires that a location is (i) allocated and (ii) not freed. Rules `tms-alloc` and `tms-dealloc` both require a location to not be freed already and extend the monitor state accordingly. This restriction effectively disallows reallocation to reassign the same location to an object. However, the definition can easily be adapted by, e.g., attaching a natural number serving as a counter. Contrary to other monitors in this paper, the multi-step relation of the TMS monitor is non-standard:

$$\begin{array}{c}
\text{364} \\
\text{365} \\
\text{366} \\
\text{367} \\
\text{368} \\
\text{369} \\
\text{370} \\
\text{371} \\
\text{372}
\end{array}
\frac{
\boxed{\vdash T_{TMS} \xrightarrow{\bar{a}^*} T_{TMS}'} \text{ „Monitor } T_{TMS} \text{ multi-steps to } T_{TMS}' \text{ given trace } \bar{a}.”
}{
\frac{
\frac{
T_{TMS}.allocated = \emptyset
}{\vdash T_{TMS} \xrightarrow{[\cdot]^*} T_{TMS}}
\quad (tms\text{-refl})
}{
\vdash T_{TMS} \xrightarrow{\bar{a}^*} T_{TMS}'
}
\quad (tms\text{-ign-trans})
\quad
\frac{
\vdash T_{TMS} \xrightarrow{a} T_{TMS}' \vdash T_{TMS}' \xrightarrow{\bar{a}^*} T_{TMS}''
}{
\vdash T_{TMS} \xrightarrow{a \cdot \bar{a}} T_{TMS}''
}
\quad (tms\text{-trans})
}$$

Rules `tms-ign-trans` and `tms-trans` are the same for all monitors, but Rule `tms-refl` has, in this case, an additional premise that no more locations should be allocated. This rejects the behavior of programs that forget to free memory.

$$\begin{array}{c}
\text{373} \\
\text{374} \\
\text{375} \\
\text{376} \\
\text{377} \\
\text{378}
\end{array}
\frac{
\boxed{a \cong_{tms} a} \text{ „Abstract event } a \text{ is equivalent to } a \text{ with respect to TMS.”}
}{
\frac{
\text{Alloc } l \ n; t; \sigma \cong_{tms} \text{Alloc } l \ t
}{
\text{Branch } n \cong_{tms} \epsilon
}
\quad (tms\text{-alloc-authentic})
\quad
\frac{
}{
\downarrow \cong_{tms} \downarrow
}
\quad (tms\text{-branch-authentic})
\quad
\frac{
}{
\downarrow \cong_{tms} \downarrow
}
\quad (tms\text{-abort-authentic})
}$$

The trace agreement is entirely straightforward, so only allocation, branch, and crash are shown.

LEMMA 3.7 (TRACES WITH MONITOR SATISFACTION ARE tms). *If $\vdash_{mon} \bar{a} : tms$, then $\bar{a} \in tms$.*

Example 3.8 (A program not satisfying TMS). Consider the following C++11 library that calls `strncpy` (Example 1.1) and prints the result to the standard output stream.

```

383 int greet() { // allocates 12 chars containing a greeting message
384   char* greetings = new char[12] { "Hello_POPL!" }; // <- address l_x
385   char* to = new char[12]; // <- address l_y
386
387   strncpy(12, to, greetings);
388   delete to;
389
390   printf("%cOPL\n", to[6]);
391 }

```


Up to the body of `printf`, the program execution yields the specification trace $\text{Alloc } l_x \ 12; \text{comp}; \blacksquare \cdot \text{Alloc } l_y \ 12; \text{comp}; \blacksquare \cdot \text{Use } l_x \ 0; \text{comp}; \blacksquare \cdot \text{Use } l_x \ 0; \text{comp}; \blacksquare \cdot \text{Use } l_y \ 0; \text{comp}; \blacksquare \cdot \text{Use } l_x \ 1; \text{comp}; \blacksquare \cdot \dots \cdot \text{Use } l_x \ 12; \text{comp}; \blacksquare \cdot \text{Dealloc } l_y; \text{comp}; \blacksquare \cdot \text{Use } l_y \ 6; \text{comp}; \blacksquare$. Relating this trace to abstract monitor events yields $\bar{a} = \text{Alloc } l_x \ \text{comp} \cdot \text{Alloc } l_y \ \text{comp} \cdot \text{Use } l_x \ \text{comp} \cdot \text{Use } l_x \ \text{comp} \cdot \text{Use } l_y \ \text{comp} \cdot \dots \cdot \text{Use } l_x \ \text{comp} \cdot \text{Dealloc } l_y \ \text{comp} \cdot \text{Use } l_y \ \text{comp}$. Remembering the definition of `strncpy` (Example 1.1), observe that it does not deallocate its arguments. Even though the trace contains an out-of-bounds access right before returning from `strncpy`, this is no concern for TMS, since the location l_x is still allocated. However, having returned from `strncpy`, the `greet` function continues and deallocates l_y whose subsequent use in the `printf` call is a use-after-free bug.

The fix would be to `delete` `greetings` instead of `to` and add a `delete` to after the `printf` call, which leads to the abstract monitor trace $\bar{a}' = \text{Alloc } l_x \ \text{comp} \cdot \text{Alloc } l_y \ \text{comp} \cdot \text{Use } l_x \ \text{comp} \cdot \text{Use } l_x \ \text{comp} \cdot \text{Use } l_y \ \text{comp} \cdot \dots \cdot \text{Use } l_x \ \text{comp} \cdot \text{Dealloc } l_x \ \text{comp} \cdot \text{Use } l_y \ \text{comp}$. It follows that $\vdash_{\text{mon}} \bar{a}' : \text{tms}$ and from Lemma 3.7 (Traces with Monitor Satisfaction are tms), it follows that the program satisfies Definition 3.1 (TMS), even though the program still violates SMS.

3.2.2 Monitor for SMS.

(Abstract Store) $T_{\text{SMS}} := L \times t \times \mathbb{N}$ (Abstract Events) $a ::= \varepsilon \mid \text{Alloc } l \ t \ n \mid \text{Use } l \ t \ n$

$$\boxed{\vdash T_{\text{SMS}} \xrightarrow{a} T_{\text{SMS}'}} \text{ „Monitor } T_{\text{SMS}} \text{ does one step to } T_{\text{SMS}'} \text{ given event } a\text{.”}$$

$$\frac{\text{(sms-use)} \quad (l; t; m) \in T_{\text{SMS}} \quad n < m}{\vdash T_{\text{SMS}} \xrightarrow{\text{Use } l \ t \ n} T_{\text{SMS}}} \quad \frac{\text{(sms-alloc)} \quad (l; t; m) \notin T_{\text{SMS}}}{\vdash T_{\text{SMS}} \xrightarrow{\text{Alloc } l \ t \ n} T_{\text{SMS}} \cup \{(l; t; n)\}}$$

The state of the monitor for SMS is a set containing tuples of locations, control-tags, and the allocation size. In comparison to the trace model of the TMS monitor, the trace model here is extended by sizing and positional information. Rule `sms-use` performs a bounds check and Rule `sms-alloc` adds bounds information to the state of the monitor. The trace agreement is entirely straightforward and similar to the one for TMS.

LEMMA 3.9 (TRACES WITH MONITOR SATISFACTION ARE SMS). *If $\vdash_{\text{mon}} \bar{a} : \text{sms}$, then $\bar{a} \in \text{sms}$.* \blacksquare

Example 3.10 (Normal invocation of `strncpy`). Consider the insecure `strncpy` function from Example 1.1 with a context `strncpy(2, x, y)`, where x and y are pointers to valid regions of memory with allocated space for exactly two cells and do not contain the null-terminating character `'\0'`. For the sake of this example, the pointers have been allocated by the component and passed to the context. The loop of `strncpy` will copy exactly two cells and then check the loop condition for the last time. At that stage, the induction variable i is equal to 2 and, unfortunately, the order of checks is such that first the cell $x[i]$ is read prior to bounds checking $i < n$. Because of this, there is an out-of-bounds memory access right before exiting the function. This is also visible on the trace, which can be sketched as $\dots \cdot \text{Alloc } l_x \ 2; \text{comp}; \blacksquare \cdot \dots \cdot \text{Alloc } l_y \ 2; \text{comp}; \blacksquare \cdot \dots \cdot \text{Use } l_x \ 0; \text{comp}; \blacksquare \cdot \text{Use } l_y \ 0; \text{comp}; \blacksquare \cdot \text{Use } l_x \ 1; \text{comp}; \blacksquare \cdot \text{Use } l_y \ 1; \text{comp}; \blacksquare \cdot \text{Use } l_x \ 2; \text{comp}; \blacksquare \cdot \dots$, where l_x and l_y are the memory addresses associated to x and y , respectively. Omitting the events for all „ \dots ” for sake of brevity, the abstract monitor trace of this is $\text{Alloc } l_x \ \text{comp} \ 2 \cdot \text{Alloc } l_y \ \text{comp} \ 2 \cdot \text{Use } l_x \ \text{comp} \ 0 \cdot \text{Use } l_y \ \text{comp} \ 0 \cdot \text{Use } l_x \ \text{comp} \ 1 \cdot \text{Use } l_y \ \text{comp} \ 1 \cdot \text{Use } l_x \ \text{comp} \ 2$.

After the allocation events, the state of the monitor is $\{(l_x; \text{comp}; 2), (l_y; \text{comp}; 2)\}$. All uses up to the last are accepted by the monitor, but the last event does not satisfy the premise $2 < 2$ in Rule `sms-use`. Therefore, the whole program (`strncpy` linked with this kind of context) is not SMS.

3.2.3 Combining TMS and SMS Monitors to obtain MS.

4 COMPOSING SECURE COMPILERS

This section presents the key meta-theoretic results of this paper concerning sequential compiler composition (and of optimisation passes) (Section 4.1) and concerning other kinds of compiler composition (Section 4.2).

4.1 Secure Sequential Composition

The main result is that secure compilers in the robust compilation framework [Abate et al. 2019] compose *sequentially*. This is not intuitive in the sense that in the security domain, composition does not work without additional generalizations [Canetti et al. 2006; Fabian et al. 2022; McCullough 1988]. The sequential composition of compilers γ_L^L and γ_L^L is defined as follows: Given an L program p and compilers γ_L^L, γ_L^L , its compiled L counterpart is obtained by plugging p into $\gamma_L^L \circ \gamma_L^L$.

Definition 4.1. $\gamma_L^L \circ \gamma_L^L \stackrel{\text{def}}{=} \text{Given } p, \text{ yield } \gamma_L^L(\gamma_L^L(p))$

Consider the compilation chain for TypeScript. First, TypeScript programs are translated to JavaScript which, e.g., V8 [Google 2008] eventually compiles in parts to IgnitionBC. The following theorem establishes what happens if all these compilation steps were robustly secure with respect to MS: The resulting IgnitionBC code would be MS regardless of the context the binary runs in.

Given γ_L^L robustly preserves C_1 and γ_L^L robustly preserves C_2 , it follows that their sequential composition $\gamma_L^L \circ \gamma_L^L$ robustly preserves the intersection of classes C_1 and C_2 .

THEOREM 4.2 (SEQUENTIAL COMPOSITION OF SECURE COMPILERS). *If $\vdash \gamma_L^L : C_1$ and $\vdash \gamma_L^L : C_2$, then $\vdash \gamma_L^L \circ \gamma_L^L : C_1 \cap C_2$.*

Since the composition of secure compilers is again a secure compiler, the theorem generalises to a whole chain of n secure compilers.

4.1.1 Securing Optimisations. Notably, real-world compilation chains also perform a series of (sequential) passes whose main purpose is not necessarily to translate from one language to another, but to, e.g., optimise the code or enforce a certain property. Both examples can be seen in practice, e.g. as in the work of [Akritidis et al. 2009; Manjikian and Abdelrahman 1997; Nagarakatte et al. 2009, 2010; Wegman and Zadeck 1991] and many more. Consider the following two LLVM optimisation passes: CF, which rewrites constant expressions to the constant they evaluate to, and DCE, which removes dead code by rewriting conditional branches. The order in which CF and DCE are performed influences the final result of the compilation (see Figure 2). This *phase ordering*

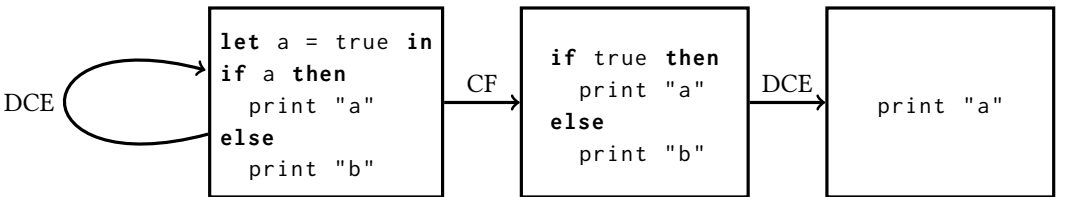



Fig. 2. Example program where the level of optimisations differ for one pass of applying CF and DCE in any order. Every edge is a compilation pass and the label on the edge states what the pass does, i.e., CF or DCE. The source code in the nodes is a glorified compiler intermediate representation and the code gets more optimised towards the right hand side of the figure.

problem is well-known in literature and a practical solution is to simply perform a fixpoint iteration

of the optimisation pipeline [Click and Cooper 1995]. Compiler engineers typically try to find an order of optimisations that yields well-optimised programs for either code size [Cooper et al. 1999] or performance [Kulkarni et al. 2006]. Corollary 4.3 justifies that any such order of compilation passes is valid with respect to security. So, given two compilation passes $\gamma_{1L}^L, \gamma_{2L}^L$, both robustly preserving class \mathbb{C}_1 or \mathbb{C}_2 , respectively, for any order of their composition the composed compiler robustly preserves the intersection of \mathbb{C}_1 and \mathbb{C}_2 .


COROLLARY 4.3 (SEQUENTIAL COMPOSITION OF SECURE COMPILERS). *If $\vdash \gamma_{1L}^L : \mathbb{C}_1$ and $\vdash \gamma_{2L}^L : \mathbb{C}_2$, then $\vdash \gamma_{1L}^L \circ \gamma_{2L}^L : \mathbb{C}_1 \cap \mathbb{C}_2$ and $\vdash \gamma_{2L}^L \circ \gamma_{1L}^L : \mathbb{C}_2 \cap \mathbb{C}_1$.* 

4.2 Secure Upper and Lower Composition

Besides sequential composition, there are two other compositions, namely an *upper*, i.e., a compiler that takes multiple inputs and yields one output, and a *lower* composition, i.e., a compiler that takes one input and yields multiple outputs. Define the upper composition γ_L^{L+L} as follows: Given a program p , its compiled counterpart is obtained by plugging p into γ_L^L if $p \in L$ or by plugging p into γ_L^L if $p \in L$.

Definition 4.4 (Upper Composition). $\gamma_L^{L+L} \stackrel{\text{def}}{=} \lambda p. \begin{cases} \text{if } p \in L, \text{ then } \gamma_L^L(p) \\ \text{if } p \in L, \text{ then } \gamma_L^L(p) \end{cases}$

Examples of this are present in industry: Consider the Java Virtual Machine bytecode **JVMBC**, which is a popular target for programming language designers due to its high performance and relevance in industry. Compilers for several programming languages have it as their target language, some popular instances are **Java** and **Kotlin**. Technically speaking, they both compile to class files and **Kotlin** objects are considered to be the same as **Java** objects at that point. Both languages can be used at the same time in one project [Google [n. d.]]. A compiler that accepts both **Java** and **Kotlin** code translating to the same target language or intermediate representation performs a kind of *upper* composition. Now, the following theorem tells us what happens if these are secure: Given γ_L^L robustly preserves \mathbb{C}_1 and γ_L^L robustly preserves \mathbb{C}_2 , it follows that their upper composition γ_L^{L+L} robustly preserves the intersection of classes \mathbb{C}_1 and \mathbb{C}_2 .

THEOREM 4.5 (UPPER COMPOSITION OF SECURE COMPILERS). *If $\vdash \gamma_L^L : \mathbb{C}_1$ and $\vdash \gamma_L^L : \mathbb{C}_2$, then $\vdash \gamma_L^{L+L} : \mathbb{C}_1 \cap \mathbb{C}_2$.* 

Dually, the *lower* composition is concerned about compilers that accept the same source but yield different target languages. Define the lower composition γ_{L+L}^L as follows: Given a program p , its compiled counterpart is obtained by plugging p into γ_L^L or by plugging p into γ_L^L , respectively, based on the internal decision.

Definition 4.6 (Lower Composition). $\gamma_{L+L}^L \stackrel{\text{def}}{=} \lambda p. L. \begin{cases} \text{if } L = L, \text{ then } \gamma_L^L(p) \\ \text{if } L = L, \text{ then } \gamma_L^L(p) \end{cases}$

Consider two compilers both accepting **LLVMIR** [Lattner and Adve 2004] and one of them emits **x86_64**, while the other emits **ARMv8**. It is intuitive that they are in some sense composed in the LLVM framework, but the decision of when to use one over the other is inherently *internal* to the formalisation effort of this kind of composition. For example, the user of this compiler provides an explicit flag that instructs to emit **x86_64** or the framework itself detects the target platform via heuristics, such as supported instructions.

The following theorem demonstrates what happens if the involved compilers are secure: Given γ_L^L robustly preserves \mathbb{C}_1 and γ_L^L robustly preserves \mathbb{C}_2 , it follows that their lower composition γ_{L+L}^L robustly preserves the intersection of classes \mathbb{C}_1 and \mathbb{C}_2 .

THEOREM 4.7 (LOWER COMPOSITION OF SECURE COMPILERS). *If $\vdash \gamma_L^L : \mathbb{C}_1$ and $\vdash \gamma_L^L : \mathbb{C}_2$, then $\vdash \gamma_{L+L}^L : \mathbb{C}_1 \cap \mathbb{C}_2$.*

5 CASE STUDY: LANGUAGE FORMALISATIONS

This section defines programming languages that the secure compilers defined in the next section will use. To this end, this section defines the languages L_{tms} , L , L_{ms} , and L_{sct} which share many common elements (presented in Section 5.1). L_{tms} is the only statically typed language and exhibits the property that all well-typed programs are TMS (Section 5.2). However, not all L_{tms} programs are SMS. That is, there are well-typed L_{tms} programs that perform an out-of-bounds access. Language L is untyped and does not provide any guarantees with regards to MS (Section 5.3). L_{ms} is exactly the same language as L , but this paper still distinguishes the two for sake of readability (Section 5.4). All three languages — so L_{tms} , L , and L_{ms} — assume CCT to hold.

Writing code attaining CCT should not be of the programmer's concerns [Cauligi et al. 2019]. Such consideration is also backed up by architecture providing a data (operand) independent timing mode, such as processors by Arm [Arm 2020, p. 543] and Intel [Intel 2023, p. 80]. In spirit of this, language L_{sct} allows violating CCT by emitting events on, e.g., branching and division, that contain secrets (Section 5.5), but provides a way to read and write to a *model-specific register* that enables a "CCT-mode".

5.1 Shared Language Definitions

(Expressions) $e ::= x \mid v \mid e_1 \oplus e_2 \mid x[e] \mid \text{let } x = \text{new } e_1 [e_2] \text{ in } e_3 \mid \text{delete } x \mid x[e_1] \leftarrow e_2$
 $\mid \langle e_1; e_2 \rangle \mid e.0 \mid e.1 \mid \text{let } x = e_1 \text{ in } e_2 \mid \text{return } e \mid \text{call } g \ e \mid \text{ifz } e \text{ then } e_1 \text{ else } e_2 \mid \text{abort}()$
 (Types) $\tau ::= \mathbb{N}_t \mid \tau_1 \times \tau_2$ (Functions) $F ::= fn \ \text{foo } x := e$ (Libraries) $\Xi ::= [\cdot] \mid F, \Xi$
 (Component Names) $\xi ::= [\cdot] \mid \text{foo}, \xi$ (Programs) $\langle \Xi_{\text{ctx}}; \Xi_{\text{comp}} \rangle$

Above is the shared syntax of all the programming languages of this paper. Variables are referred to as x, y, z, a, b, c, \dots while functions may be referred to as multi-character words, such as foo , as well as short-forms like f, g, h . All languages share the type \mathbb{N}_t representing natural numbers. Functions are constrained to take one argument and can only call other functions listed in libraries, which are just lists of functions. A program $\langle \Xi_{\text{ctx}}; \Xi_{\text{comp}} \rangle$ is indexed by two libraries that represent all context- and component-level functions, respectively. Lists of component-level names are referred to as ξ .

(Control Tags) $t ::= \text{ctx} \mid \text{comp}$ (Communication Tags) $c ::= ? \mid ! \mid \emptyset$ (Poison Tags) $\rho ::= \text{⊗} \mid \square$
 (Continuation Stacks) $\bar{K} ::= [\cdot] \mid (K; g), \bar{K}$ (Control Flow States) $\Psi ::= (\Xi; \bar{K})$
 (Locations) $l \in \mathbb{N}$ (Substitutions) $\gamma ::= [\cdot] \mid [v \text{ for } x], \gamma$
 (Stores) $\Delta ::= [\cdot] \mid x \mapsto (l; t; \rho; n), \Delta$ (Heaps) $H ::= [\cdot] \mid v, H$
 (Memory States) $\Phi ::= (H^{\text{ctx}}, H^{\text{comp}}; \Delta)$ (States) $\Omega ::= (\Psi; t; \Phi)$ (Runtime Terms) $r ::= \Omega \triangleright e$

States Ω are tuples⁴ containing a control flow state Ψ , control tags t , and a memory state Φ . A control flow state Ψ entails a library, which provides definitions for functions calls, as well as a stack of continuations \bar{K} . Elements of the stack \bar{K} are pairs of evaluation contexts and the name of a function associated to that evaluation context. Memory states Φ are tuples of two separate heaps H and a store Δ , which contains pointer metadata, such as the concrete memory location l , a control tag t indicating which heap the pointer points into, a poison tag ρ , as well as bounds information n . The bounds information is a mere proof artefact that has no semantic significance. The two separate heaps essentially model a sandbox, to prevent contexts from performing pointer arithmetic and reading from or writing to the data owned by a component. While this prevents the effects of out-of-bounds accesses across the context and component boundary, the goal of the design of the languages of this paper is to be able to express security violations. To this end, the poison tag ρ indicates whether a pointer has been freed \clubsuit or is still allocated \square , so that pointers can be used even after their deallocation without the semantics getting stuck. Runtime terms are simply expressions e paired with the operational state Ω .

$$\begin{aligned} (\text{Pre-Events}) \ a_b &::= \text{Alloc } l \ v \mid \text{Dealloc } l \mid \text{Get } l \ v \mid \text{Set } l \ v \ v' \mid \dots \\ (\text{Events}) \ a &::= \varepsilon \mid \frac{\!}{\!} \mid (a_b; t; \sigma) \end{aligned}$$

All languages use the same trace model, where events are either the empty event ε , the program crash event $\frac{\!}{\!}$, or a tuple consisting of a control tag t and a security tag σ . The former indicates whether the component compor the context ctx is to blame for emitting this event, the latter indicates the secrecy level of values of the emitted event, i.e., either \clubsuit or \spadesuit . As for pre-events, the memory-related ones are allocation ($\text{Alloc } l \ v$), deallocation ($\text{Dealloc } l$), reading from ($\text{Get } l \ v$) and writing to memory ($\text{Set } l \ v \ v'$). The following is an excerpt of the operational semantics handling some of the memory operations.

$$\begin{array}{c} \boxed{r \xrightarrow{a}_p r'} \text{ „}r \text{ does one primitive step to } r' \text{ emitting event } a\text{.”} \\ \begin{array}{c} (e - \text{get} - \varepsilon) \\ t = \Omega.t \quad \Omega.\Delta(x) = (l; t; \rho; m) \\ l + n \in \text{dom } \Omega.H^t(l + n) \end{array} \quad \begin{array}{c} (e - \text{set} - \notin) \\ t = \Omega.t \quad \Omega.\Delta(x) = (l; t; \rho; m) \\ l + n \notin \text{dom } \Omega.H^t \end{array} \\ \hline \Omega \triangleright x[n] \xrightarrow{(\text{Get } l \ n; t)}_p \Omega \triangleright H^t(l + n) \quad \Omega \triangleright x[n] \leftarrow v \xrightarrow{(\text{Set } l \ n; v; t)}_p \Omega \triangleright v \\ \begin{array}{c} (e - \text{new}) \\ \Omega \vdash z \text{ fresh} \quad \Omega \vdash l \text{ fresh} \quad H_1^t = \Omega.H^t \ll n \quad \Delta_1 = z \mapsto (l; \Omega.t; \square; n), \Omega.\Delta \end{array} \\ \hline \Omega \triangleright \text{new } x \ [n]e \xrightarrow{(\text{Alloc } l \ n; t)}_p \Omega \ [H^t \text{ for } H_1^t] \ [\Delta \text{ for } \Delta_1] \triangleright e \ [z \text{ for } x] \\ \begin{array}{c} (e - \text{dealloc}) \\ \Omega.\Delta(x) = (l; \Omega.t; \rho; n) \quad \Delta_1 = \Omega.\Delta(x \mapsto (l; \Omega.t; \clubsuit; n)) \end{array} \quad \begin{array}{c} (e - \text{abort}) \end{array} \\ \hline \Omega \triangleright \text{delete } x \xrightarrow{(\text{Dealloc } l; t)}_p \Omega \ [\Delta \text{ for } \Delta_1] \triangleright 0 \quad \Omega \triangleright \text{abort}() \xrightarrow{\frac{\!}{\!}}_p \frac{\!}{\!} \end{array}$$

To demonstrate the use of the poison tag ρ as metadata for pointers instead of removing them from the store Δ , consider Rules $e - \text{get} - \varepsilon$, $e - \text{set} - \notin$ and $e - \text{dealloc}$. In Rule $e - \text{dealloc}$, the premise does not care at all about the actual state of the poison tag ρ and just overwrites it, marking the location as freed \clubsuit . Besides that, the poison tag does not have any semantic meaning. For language L_{tms} , this tag is really just some semantic metadata that programmers have no access to. But, for the other languages, e.g., L , the poison tag is used to check pointer validity. Rule $e - \text{new}$ allocates enough space on the respective heap, either H^{ctx} or H^{comp} depending on the execution

⁴Throughout the paper, the substitution notation is also used to update entries in states Ω .

context, i.e., the value of $\Omega.t$, and adds the appropriate metadata associated to the pointer in Δ . Reading from Rule $e - \text{get} - \in$ and writing to memory Rule $e - \text{set} - \notin$ have two cases: Either the heap is large enough or not and, depending on that, either the actual value stored at that location is read from or written to, or some garbage data is returned. However, note that the execution does not get stuck in such cases, it performs a step, and emits an appropriate event. Also note that whether a pointer is poisoned or not is not reflected on the trace.

(Pre-Events) $a_b ::= \dots \mid \text{Call } c \text{ } g \text{ } v \mid \text{Ret } c \text{ } v \mid \text{Start} \mid \text{End } v$

A key difference in comparison with the specification trace model (Section 3.1) is that, as standard in secure compilation work [Abate et al. 2019; El-Korashy et al. 2021; Patrignani and Garg 2021], the traces have a call and return event that signals context switches, which are referred to as *interaction events*. The reason for these interaction events is technical: They are a proof artifact for reconstructing a source context from a potentially malicious target context, where during that translation, the insertion of some wrapper code right before context switching may be necessary to make the proof succeed. Hereby, a Call ? foo v and Return ? v signal that program execution transitions from context- to component-level. Contrary, Call ! foo v and Return ? v signal that program execution transitions from component- to context-level. For calls without this context switch, the environmental semantics attaches the \emptyset tag. In the following, $\neg\text{ctx} = \text{comp}$ and $\neg\text{comp} = \text{ctx}$.

$r \xrightarrow{a}_{\text{ctx}} r'$ „Contextual step from runtime-term r to r' emitting event a .”

$$\frac{\begin{array}{c} \Omega.\bar{K} = (K; \text{foo}), \bar{K}' \quad \xi \vdash \text{foo} : \Omega.t \vdash c \\ \Omega \triangleright K' [\text{return } v] \xrightarrow{(\text{Ret } c \text{ } v; \Omega.t)}_{\text{ctx}} \Omega [-\Omega.t \text{ for } t] [\bar{K}' \text{ for } \bar{K}] \triangleright K [v] \end{array}}{\begin{array}{c} \Omega \triangleright K [\text{call } \text{foo } v] \xrightarrow{(\text{Call } c \text{ } \text{foo } v; \Omega.t)}_{\text{ctx}} \Omega [-\Omega.t \text{ for } t] [(K; \text{foo}), \Omega.\bar{K} \text{ for } \bar{K}] \triangleright e [v \text{ for } x] \end{array}}$$

The environmental semantics is mostly straightforward. In Rules $e - \text{ret}$ and $e - \text{call} - \text{notsame}$, the judgement $\Omega.\xi \vdash \text{foo} : \Omega.t$ checks whether foo is a component-level name by looking it up in the list of component-level names ξ and emits the appropriate transfer tag, i.e., either ! or ?. Additional rules that are left out ensure that, e.g., when calling the main function, the event Start is emitted, which is a design choice this paper does for convenience when reasoning about call-chains. Note that the End v event is not emitted if the program crashes.

The top-level execution $\langle \Xi_{\text{ctx}}; \Xi_{\text{comp}} \rangle \xrightarrow{\bar{a}} r$ constructs an initial state Ω by linking Ξ_{ctx} and Ξ_{comp} and then starts execution by calling the main function. The trace \bar{a} emitted during that execution serves as abstraction of the behavior of the program enabling the use of Definitions 2.1 and 2.2.

5.2 L_{tms} : A Temporal but Not Spatial Memory Safe Language

L_{tms} uses the same syntax as presented earlier (Section 5) without extensions to the term level. But, L_{tms} is statically typed, where the type system is inspired by L^3 [Morrisett et al. 2005; Scherer et al. 2018]. The type system of L_{tms} exhibits the property that every well-typed L_{tms} program satisfies TMS (Theorem 5.1). The proof of this theorem relies on a projection $\text{Proj}^{L_{\text{tms}}}(\delta, a) = a$ from L_{tms} events to specification events a , because the properties defined earlier (Section 3.1) are defined in

the specification trace model. Hereby, the map $\delta(l) = l$ maps an L_{tms} location to a location l of the specification trace model.

$$\begin{array}{c}
 \boxed{\text{Proj}^{L_{\text{tms}}}(\delta, a) = a} \text{ „Project an } L_{\text{tms}} \text{ event } a \text{ to a specification event } a\text{.”} \\
 \frac{\text{(L}_{\text{tms}}\text{-filter-context)} \quad a_b \neq \downarrow}{\text{Proj}^{L_{\text{tms}}}(\delta, (a_b; \text{ctx}; \sigma) = \varepsilon)} \quad \frac{\text{(L}_{\text{tms}}\text{-filter-abort)}}{\text{Proj}^{L_{\text{tms}}}(\delta, \downarrow) = \downarrow} \quad \frac{\text{(L}_{\text{tms}}\text{-filter-start)}}{\text{Proj}^{L_{\text{tms}}}(\delta, (\text{Start}; \text{comp})) = \varepsilon} \\
 \frac{\text{(L}_{\text{tms}}\text{-filter-alloc)} \quad \delta(l) = l \quad n = n}{\text{Proj}^{L_{\text{tms}}}(\delta, (\text{Alloc } l \ n; \text{comp})) = (\text{Alloc } l \ n; \text{comp}; \blacksquare)}
 \end{array}$$

Most rules of the projection $\text{Proj}^{L_{\text{tms}}}(\delta, a)$ are left out since, for the most part, it does the expected, e.g., $\text{Proj}^{L_{\text{tms}}}(\delta, \text{Dealloc } l; \text{comp}; \sigma) = \text{Dealloc } \delta(l); \text{comp}; \sigma$. But, it also filters any action that a context does as well as the interaction events, since these are irrelevant for component-level TMS.

THEOREM 5.1 (L_{tms} -PROGRAMS ARE TMS). *For any $\Xi_{\text{comp}}, \vdash_R \Xi_{\text{comp}} : \text{tms}$* 

5.3 L: A Memory-Unsafe Language

(Expressions) $e ::= \dots \mid x \text{ is } \heartsuit \mid e \text{ has } \tau$

L extends the syntax presented earlier (Section 5.1) with dynamic typechecks $e \text{ has } \tau$ and a way to inspect poison tags $x \text{ is } \heartsuit$ in the metadata of pointers. For valid pointers (\square) bound to variable x , the check $x \text{ is } \heartsuit$ yields **1**. If the array bound to x was allocated, i.e., has been poisoned (\heartsuit), the check $x \text{ is } \heartsuit$ evaluates to **0**.

$$\frac{(e - x \text{ has } \mathbb{N}_t)}{\Omega \triangleright x \text{ has } \mathbb{N}_t \xrightarrow{\varepsilon}_p \Omega \triangleright \mathbf{1}} \quad \frac{(e - n \text{ has } \mathbb{N}_t)}{\Omega \triangleright n \text{ has } \mathbb{N}_t \xrightarrow{\varepsilon}_p \Omega \triangleright \mathbf{0}}$$

Dynamic typechecks $e \text{ has } \tau$ match on e and evaluate to **0** if the term is of type τ and **1** otherwise. The projection $\text{Proj}^L(\delta, \bar{a})$ is equal to $\text{Proj}^{L_{\text{tms}}}(\delta, \bar{a})$.

5.4 L_{ms} : Another Memory-Unsafe Language

To enhance readability, this paper uses L_{ms} , despite it being exactly equal to **L** (Section 5.3). The projection $\text{Proj}^{L_{\text{ms}}}(\delta, \bar{a})$ is also exactly equal to $\text{Proj}^L(\delta, \bar{a})$.

5.5 L_{sctt} : A Memory-Unsafe Language with a Data Independent Timing Mode

(Expressions) $e ::= n^\sigma \mid \dots \mid \text{let } x^\sigma = e_1 \text{ in } e_2 \mid \dots \mid \text{wrdoit } e \mid \text{rddoit } x \text{ in } e$
 (States) $\Omega ::= (\Psi; t; n; \heartsuit)$

L_{sctt} extends L_{ms} (Section 5.4) with a way to write to a model specific register that controls a data (operand) independent timing mode, a feature that is present in both Arm [Arm 2020, p. 543] and Intel [Intel 2023, p. 80] processors. To this end, states are extended with the value of the register, which is initially set to be not active. If the register is marked active, the intuition is that no secrets can appear on specification traces. If the register is marked inactive, secrets may appear on traces. For the other languages seen earlier, the mode is intuitively always-on, i.e., the mode of execution always uses data independent timing. The language also adds user-annotations to values and variables to know their secrecy σ , which is either high \blacksquare or low \blacktriangle . Security tags σ are on the usual secrecy lattice, where $\blacksquare \leq \sigma$ and $\sigma \leq \blacktriangle$.

(Pre-Events) $a_b ::= \dots \mid \widehat{\text{Get}}\ l\ v \mid \widehat{\text{Set}}\ l\ v\ v' \mid \text{Branch}\ n \mid \text{Binop}\ n$

To prevent secrets from leaking but still enable reasoning about memory safety, L_{sct} extends pre-events with $\widehat{\text{Get}}\ l\ v$ and $\widehat{\text{Set}}\ l\ v$. These indicate reads from and writes to memory without leaking secret information involved in the access. Moreover, the language extends pre-events with $\text{Branch}\ n$ and $\text{Binop}\ n$ that are emitted when evaluating a branch or certain binary expressions, such as division, respectively, whenever the data independent timing mode is inactive. The following rules demonstrate how this is handled semantically.

$$\begin{array}{c}
 \frac{(e - \oplus - \text{noleak}) \quad m \neq 0 \quad n_3 = n_1 \oplus n_2 \quad \sigma'' \leq \sigma \quad \sigma'' \leq \sigma'}{\Psi; t; m; \Phi \triangleright n_1^\sigma \oplus n_2^{\sigma'} \xrightarrow{e} \Psi; t; m; \Phi \triangleright n_3^{\sigma''}} \quad \frac{(e - \text{wrdoit}) \quad \Psi; t; m; \Phi \triangleright \text{wrdoit}\ n^\sigma \xrightarrow{e} \Psi; t; n; \Phi \triangleright n^\sigma}{(e - \text{ifz} - \text{true} - \text{leak})} \\
 \frac{\Psi; t; 0; \Phi \triangleright \text{ifz}\ 0^\sigma \text{ then } e_1 \text{ else } e_2 \xrightarrow{\text{Branch}\ 0; t; \sigma} \Psi; t; 0; \Phi \triangleright e_2}{}
 \end{array}$$

The evaluation steps are amended to propagate the security-tag annotations σ . When the data independent timing mode is active, pre-events $\text{Branch}\ n$ and $\text{Binop}\ n$ are emitted for conditionals and binary operations, respectively.

$\text{Proj}^{\text{L}_{\text{sct}}}(\delta, a) = a$ „Project an L_{sct} event a to a specification event a .”

$$\begin{array}{c}
 \frac{(L_{\text{sct}}\text{-filter-context}) \quad a_b \neq \text{z}}{\text{Proj}^{\text{L}_{\text{sct}}}(\delta, (a_b; \text{ctx}; \sigma)) = \varepsilon} \quad \frac{(L_{\text{sct}}\text{-filter-get}) \quad \delta(\mathbf{1}) = l \quad n = n}{\text{Proj}^{\text{L}_{\text{sct}}}(\delta, (\widehat{\text{Get}}\ l\ n; \text{comp}; \sigma)) = (\text{Get}\ l\ n; \text{comp}; \blacksquare)} \\
 \frac{(L_{\text{sct}}\text{-filter-get}) \quad \delta(\mathbf{1}) = l \quad n = n}{\text{Proj}^{\text{L}_{\text{sct}}}(\delta, (\widehat{\text{Set}}\ l\ n; \text{comp}; \sigma)) = (\text{Set}\ l\ n; \text{comp}; \sigma)}
 \end{array}$$

The projection to the specification trace model is mostly straightforward and similar to the others, e.g., Section 6.1. However, for events containing the pre-events $\widehat{\text{Get}}$ and $\widehat{\text{Set}}$, the projection always translates the security-tag σ to \blacksquare , regardless of its actual value, as seen in Rule $L_{\text{sct}}\text{-filter-get}$. The pre-events themselves still translate to just Get and Set, respectively. With this technical setup, the information whether a read or write happened on a secret value is not hidden by the semantics, e.g., by emitting ε , but when projecting to specification events. This allows flexibility: The trace can be checked to satisfy different properties, such as, in this case, TMS, SMS, sCCT, and their combined versions. Example 5.2 illustrates the differences of L_{sct} compared to the other languages.

Example 5.2 (L_{sct} with and without data independent timing). Consider again the context presented in Example 3.8, where everything is marked with a security tag of high \blacksquare . The following table shows parts of the execution trace, read from top to bottom, in the left column (*Active*) with and in the right column (*Inactive*) without data independent timing. The left side of the table (L_{sct}), i.e., the two columns on the left, describes the execution trace of the program, while the right side of the table (*Spec*), i.e., the two columns on the right, describes the respective projections $\text{Proj}^{\text{L}_{\text{sct}}}(\delta, \bar{a})$ to the specification trace model.

834

835

836

837

838

839

840

841

842

843

844

845

846

847

848

849

850

851

852

853

854

855

856

857

858

859

860

861

862

863

864

865

866

867

868

869

870

871

872

873

874

875

876

877

878

879

880

881

882

L_{sect}		$Spec$	
<i>Active</i>	<i>Inactive</i>	<i>Active</i>	<i>Inactive</i>
Alloc l_x 12; comp;	Alloc l_x 12; comp;	Alloc l_x 12; comp;	Alloc l_x 12; comp;
Alloc l_y 12; comp;	Alloc l_y 12; comp;	Alloc l_y 12; comp;	Alloc l_x 12; comp;
Get l_x 0; comp;	Get l_x 0; comp;	Use l_x 0; comp;	Use l_x 0; comp;
ϵ	Branch 0; comp;	ϵ	Branch 0; comp;
Get l_x 0; comp;	Get l_x 0; comp;	Use l_x 0; comp;	Use l_x 0; comp;
Set l_y 0 'H'; comp;	Set l_y 0 'H'; comp;	Use l_y 0; comp;	Use l_y 0; comp;
Get l_x 1; comp;	Get l_x 1; comp;	Use l_x 1; comp;	Use l_x 1; comp;
ϵ	Branch 0; comp;	ϵ	Branch 0; comp;
\vdots	\vdots	\vdots	\vdots
Get l_x 12; comp;	Get l_x 12; comp;	Use l_x 12; comp;	Use l_x 12; comp;
ϵ	Branch 1; comp;	ϵ	Branch 1; comp;
Dealloc l_y ; comp;	Dealloc l_y ; comp;	Dealloc l_y ; comp;	Dealloc l_y ; comp;
Get l_y 6; comp;	Get l_y 6; comp;	Use l_y 6; comp;	Use l_y 6; comp;

When the data independent timing mode is off, the execution yields events in similar fashion to before (Sections 5.2 to 5.4). But, if it is turned on, then the branching event **does not fire** anymore and both reading and writing to memory gets ultimately translated to a specification trace with no exposed secrets.

6 CASE STUDY: COMPOSING SECURE COMPILER PASSES AND OPTIMISATIONS

This section defines several secure compilers, each of which robustly preserves a different property of interest as depicted in Figure 3. The section demonstrates the power of the framework (Sections 3

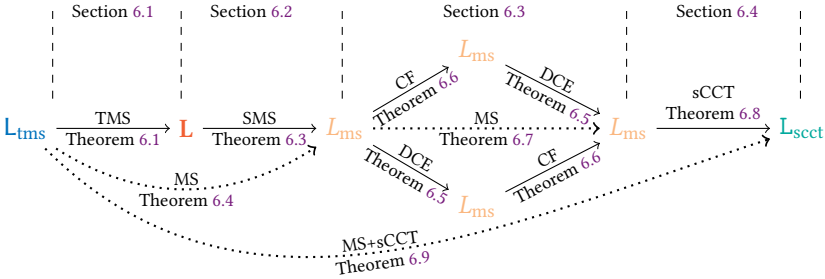


Fig. 3. Visualisation of the optimising compilation pipeline that attains a combination of MS and CCT. Vertices in the graph are the programming languages from earlier sections (Section 5). All edges are secure compilers, but dotted edges use the presented framework (Section 4) and strikethrough edges classic proof techniques. The dashed lines partition the graph into the sections where the respective theorems are presented.

and 4) by composing these compilers for a secure and optimising compilation chain that robustly preserves MS+sCCT. The first step in this chain is the compiler from L_{tms} to L that robustly preserves just TMS (Theorem 6.1). From here, an instrumentation from L to L_{ms} ensures that no out-of-bounds accesses can happen and, thus, programs at this point attain SMS (Theorem 6.3). Since these properties compose into MS, composing these passes yields a compiler that robustly

preserves MS (Theorem 6.4). At this stage, the section presents two optimising translations, namely CF and DCE, each of which robustly preserves MS (Theorems 6.5 and 6.6). These translations can be freely ordered in the compilation chain without compromising memory safety (Theorem 6.7). The last step of the chain ensures that code stays sCCT (Theorem 6.8) when lowered from L_{tms} to L_{sctt} . The final result is that the whole compilation chain robustly preserves MS+sCCT (Theorem 6.9).

6.1 Robust Temporal Memory Safety Preservation

This subsection defines a secure compiler from L_{tms} to L . To this end, the compiler needs to ensure that when execution switches from context to component, the type signatures are respected. It can do so by inserting dynamic typechecks prior to entering the body of a function belonging to the component.

$$\begin{aligned}
 \gamma_L^{\text{L}_{\text{tms}}}(\mathbf{x}) &= \mathbf{x} \\
 \gamma_L^{\text{L}_{\text{tms}}}(\mathbf{n}) &= \mathbf{n} \\
 \gamma_L^{\text{L}_{\text{tms}}}(\mathbf{e}_1 \oplus \mathbf{e}_2) &= \left[\gamma_L^{\text{L}_{\text{tms}}}(\mathbf{e}_1) \right] \oplus \left[\gamma_L^{\text{L}_{\text{tms}}}(\mathbf{e}_2) \right] \\
 \gamma_L^{\text{L}_{\text{tms}}}(\mathbf{x}[\mathbf{e}]) &= \mathbf{x} \left[\gamma_L^{\text{L}_{\text{tms}}}(\mathbf{e}) \right] \\
 \gamma_L^{\text{L}_{\text{tms}}}(\text{delete } \mathbf{x}) &= \text{delete} \left[\gamma_L^{\text{L}_{\text{tms}}}(\mathbf{x}) \right] \\
 \gamma_L^{\text{L}_{\text{tms}}}(\text{fn } \mathbf{g} \ \mathbf{x} : \mathbb{N}_t \rightarrow \tau_e := \mathbf{e}) &= \text{fn } \mathbf{g} \ \mathbf{x} := \text{ifz } \mathbf{x} \ \text{has } \mathbb{N}_t \ \text{then} \left[\gamma_L^{\text{L}_{\text{tms}}}(\mathbf{e}) \right] \ \text{else abort}()
 \end{aligned}$$

Since L has no static typechecks, it could happen that a bogus context Ξ_{ctx} invokes a callable object accepting a \mathbb{N}_t with (17:29). By inserting the check, the compiler ensures that execution does not proceed in such cases. The compiler does not insert other checks and proceeds as the identity function (which in this paper amounts to a simple re-colouring of L_{tms} to L expressions).

Compiling the `strncpy` function from Section 1 with $\gamma_L^{\text{L}_{\text{tms}}}$, the compiler would in this case ensure that the arguments that are evaluated in the compiled `strncpy` are valid.

THEOREM 6.1 (COMPILER $\gamma_L^{\text{L}_{\text{tms}}}$ IS SECURE WITH RESPECT TO TMS). $\vdash \gamma_L^{\text{L}_{\text{tms}}} : \text{tms}$

6.1.1 Proving Robust Safety Property Preservation. We illustrate the proof of Theorem 6.1 since the other secure compilation proofs of this paper follow the same approach. Unfolding the theorem statement yields the following assumptions: for any $\pi \in [\text{tms}]^5$, $\bar{\mathbf{a}}$, \mathbf{r} , and component Ξ_{comp} , we

have that $\vdash_R \Xi_{\text{comp}} : \pi$ and $\langle \Xi_{\text{ctx}}; \gamma_L^{\text{L}_{\text{tms}}}(\Xi_{\text{comp}}) \rangle \xRightarrow{\bar{\mathbf{a}}} \mathbf{r}$, where Ξ_{ctx} is arbitrary. The proof obligation is $\text{Proj}^L(\delta, \bar{\mathbf{a}}) \in \pi$, i.e., the specification trace associated to $\bar{\mathbf{a}}$ satisfies the property π . A way to show this is to relate trace $\bar{\mathbf{a}}$ to some L_{tms} trace $\bar{\mathbf{a}}$ (which already satisfies the property as per the assumptions). The assumptions already contain a target execution associated to this trace, so the task is to find an associated L_{tms} execution that yields $\bar{\mathbf{a}}$. The trace $\bar{\mathbf{a}}$ is split into different parts, as commonly done in secure compilation works [Abate et al. 2018; El-Korashy et al. 2021], where each part contains the events that either the context or the component does, but not both. Because of this, all such trace segments are „well-bracketed” in the sense that they start with either `Start`, `Call ! foo v`, or `Ret ! v` and end with either `End v`, `Ret ? v`, or `Call ? foo v`. In the following, the former is referred to as a context segment, since these executions happen in Ξ_{ctx} , and the latter is referred to as a component segment, since these executions happen in $\gamma_L^{\text{L}_{\text{tms}}}(\Xi_{\text{comp}})$. Figure 4 visualises this division for a program execution with one call from context to component and how

⁵ $[\cdot]$ lifts the property to a hyperproperty by applying the powerset operation [Clarkson and Schneider 2008].

the target execution is related to a source execution. In the figure, the green dashed lines encompass the component segments while the orange boxes contain the actual context switches from context to component or vice versa. From a technical perspective, as typically done in compilation proofs, the proof requires some setup to maintain a relation between Ω and $\bar{\Omega}$. Two cross-language relations make this precise: (i) $\dashv\!\!\dashv_{\delta}$ relates states that are involved in a context segment, allowing the target execution to perform internal calls, and (ii) \approx_{δ} relates states that are involved in a component segment, where both states need to agree exactly, i.e., the memory and the control flow states are required to contain the same information. The relations are indexed with δ , which is an injective mapping from L_{tms} locations l to L locations l . Note that the relations $\dashv\!\!\dashv_{\delta}$ and \approx_{δ} swap when context switching.




Fig. 4. Proof diagram for Theorem 6.1 depicting the general structure of robust preservation proofs. Nodes in the graph represent runtime states. Vertical lines indicate cross language relations, while horizontal ones are execution steps. The green dashed trapezoid encompasses the component segment, while the orange dotted rectangles entail the context switches. L traces are omitted for readability. L_{tms} trace segments \bar{a}_e and \bar{a}_r describe the events that happen at the boundaries, i.e., during a context switch. \bar{a}_p is the behavior of the component and the traces \bar{a}_1 and \bar{a}_2 describe the context.

So far, the paper explained how to relate an L_{tms} execution with a L execution. The next question is therefore how to build the corresponding L_{tms} execution. This is done using a standard secure compilation proof technique called trace-based backtranslation [Abate et al. 2019; El-Korashy et al. 2021; Patrignani and Garg 2021], which can be used to build a context Ξ_{ctx} that behaves similar to Ξ_{ctx} . For context segments of the trace \bar{a} it is also necessary to show that the execution behaves similarly, i.e., the context obtained from the backtranslation generates trace \bar{a} . For component segments of the trace, the relatedness of states and traces follows from a compiler correctness argument. These two arguments yield the source execution $\langle \Xi_{\text{ctx}}; \Xi_{\text{comp}} \rangle \xrightarrow{\bar{a}} r$.

The proof now works as follows. Given that $\text{Proj}^L(\delta, \bar{a}) = \text{Proj}^{L_{\text{tms}}}(\delta, \bar{a})$, the proof goal changes from $\text{Proj}^L(\delta, \bar{a}) \in \pi$ to $\text{Proj}^{L_{\text{tms}}}(\delta, \bar{a}) \in \pi$. This follows by specializing the robust satisfaction assumption $\vdash_R \Xi_{\text{comp}} : \pi$ to use the context Ξ_{ctx} , which is obtained from the backtranslation, and to use the source execution $\langle \Xi_{\text{ctx}}; \Xi_{\text{comp}} \rangle \xrightarrow{\bar{a}} r$.

6.2 Robust (Spatial) Memory Safety Preservation

981
982
983
984
985
986
987
988
989
990
991
992
993
994
995
996
997
998
999
1000
1001
1002
1003
1004
1005
1006
1007
1008
1009
1010
1011
1012
1013
1014
1015
1016
1017
1018
1019
1020
1021
1022
1023
1024
1025
1026
1027
1028
1029

$$\begin{aligned} \gamma_{L_{\text{ms}}}^L(\text{new } x [e_1] e_2) &= \text{let } x_{\text{SIZE}} = \gamma_{L_{\text{ms}}}^L(e_1) \text{ in new } x [x_{\text{SIZE}}] \gamma_{L_{\text{ms}}}^L(e_2) \\ \gamma_{L_{\text{ms}}}^L(x[e]) &= \text{let } x_n = \gamma_{L_{\text{ms}}}^L(e) \text{ in ifz } 0 \leq x_n < x_{\text{SIZE}} \text{ then } x[x_n] \text{ else abort}() \\ \gamma_{L_{\text{ms}}}^L(x[e_1] \leftarrow e_2) &= \text{let } x_n = \gamma_{L_{\text{ms}}}^L(e_1) \text{ in ifz } 0 \leq x_n < x_{\text{SIZE}} \text{ then } x[x_n] \leftarrow \gamma_{L_{\text{ms}}}^L(e_2) \text{ else abort}() \end{aligned}$$

The compiler $\gamma_{L_{\text{ms}}}^L$ only inserts bounds-checks whenever reading from or writing to memory in order to enforce SMS. For passing pointers, it has to pass them with their size information as well. To this end, the compiler introduces another, fresh identifier x_{SIZE} for each allocation that binds x to keep track of the allocation size.

Example 6.2 (Instrumented strncpy). Consider again `strncpy`, but instrumented for SMS:

```
void strncpy(size_t n, size_t dst_size, char *dst,
             size_t src_size, char *src) {
  for(size_t i = 0; i < src_size && src[i] != '\0' && i < n; ++i) {
    if(i < src_size && i < dst_size) {
      dst[i] = src[i];
    }
  }
}
```

When calling this in similar fashion to Example 3.10, the event $Use\ l_x\ 2; \text{comp}; \blacksquare$ would not be emitted during execution, since the bounds check prevents the condition `src[i] != '\0'` from executing.

THEOREM 6.3 (COMPILER $\gamma_{L_{\text{ms}}}^L$ IS SECURE WITH RESPECT TO SMS). $\vdash \gamma_{L_{\text{ms}}}^L : \text{sms}$

Theorem 6.4 states that the composition of $\gamma_L^{\text{L-tms}}$ and $\gamma_{L_{\text{ms}}}^L$ is secure with respect to MS and follows from Theorems 6.1 and 6.3 using Theorem 4.2.

THEOREM 6.4 (COMPILER $\gamma_L^{\text{L-tms}} \circ \gamma_{L_{\text{ms}}}^L$ IS SECURE WITH RESPECT TO MS). $\vdash \gamma_L^{\text{L-tms}} \circ \gamma_{L_{\text{ms}}}^L : \text{ms}$

PROOF. From Theorem 6.1 (Compiler $\gamma_L^{\text{L-tms}}$ is secure with respect to TMS) it follows that for any L_{tms} program p , it compiles to an L program \mathbf{p} that robustly satisfies TMS. Note that \mathbf{p} robustly satisfies TMS by the properties of the typesystem of L_{tms} . Then, Theorem 6.3 (Compiler $\gamma_{L_{\text{ms}}}^L$ is secure with respect to SMS) demonstrates that, assuming \mathbf{p} robustly satisfies SMS, the program \mathbf{p} compiles to an L_{ms} program p that also robustly satisfies SMS. From Theorem 6.4 (Compiler $\gamma_L^{\text{L-tms}} \circ \gamma_{L_{\text{ms}}}^L$ is secure with respect to MS) it follows that \mathbf{p} compiles to p that robustly satisfies MS, since MS is the intersection of TMS and SMS. \square

6.3 Optimising Compilers

$$\gamma_{DCE_{L_{ms}}}^{L_{ms}}(\text{ifz true then } e_1 \text{ else } e_2) = \gamma_{DCE_{L_{ms}}}^{L_{ms}}(e_1)$$

$$\gamma_{DCE_{L_{ms}}}^{L_{ms}}(\text{ifz false then } e_1 \text{ else } e_2) = \gamma_{DCE_{L_{ms}}}^{L_{ms}}(e_2)$$

$$\gamma_{DCE_{L_{ms}}}^{L_{ms}}(e_1 \oplus e_2) = \gamma_{DCE_{L_{ms}}}^{L_{ms}}(e_1) \oplus \gamma_{DCE_{L_{ms}}}^{L_{ms}}(e_2)$$

$$\gamma_{CF_{L_{ms}}}^{L_{ms}}(e) = \text{mix}(e, [\cdot])$$

$$\text{mix}(x, \bar{y}) = n \quad \text{if } [n \text{ for } x] \in \bar{y}$$

$$\text{mix}(x, \bar{y}) = x \quad \text{if } [n \text{ for } x] \notin \bar{y}$$

$$\text{mix}(n \oplus m, \bar{y}) = k \quad \text{if } n \oplus m = k$$

$$\text{mix}(\text{let } x=n \text{ in } e, \bar{y}) = \text{mix}(e, [x \text{ for } n], \bar{y})$$

$$\text{mix}(x[e], \bar{y}) = x[\text{mix}(e, \bar{y})]$$

$$\text{mix}(\text{let } x=e_1 \text{ in } e_2, \bar{y}) = \text{let } x=\text{mix}(e_1, \bar{y}) \text{ in } \text{mix}(e_2, \bar{y})$$

$$\text{mix}(\text{ifz } e_1 \text{ then } e_2 \text{ else } e_3, \bar{y}) = \text{ifz } \text{mix}(e_1, \bar{y}) \text{ then } \text{mix}(e_2, \bar{y}) \text{ else } \text{mix}(e_3, \bar{y})$$

The two optimising compiler passes from L_{ms} to L_{ms} perform DCE and CF, respectively. The DCE pass applies a naive rewrite rule on conditionals. For CF, the pass uses an auxiliary function `mix` that does the actual work. It rewrites constant binary operations, e.g., `17 - 1` to `16`, and replaces variables that are assigned to constants with their constant, e.g., `let x=7 in x` to `7`. Both passes are secure with respect to MS. The proof for either is relatively simple, because both DCE and CF do not change the way memory accesses happen. Moreover, since the input and output languages to these compilers are the same, attacker contexts do not have more power in the target language than in the source.

THEOREM 6.5 (COMPILER $\gamma_{DCE_{L_{ms}}}^{L_{ms}}$ IS SECURE WITH RESPECT TO MS). $\vdash \gamma_{DCE_{L_{ms}}}^{L_{ms}} : \text{ms}$

THEOREM 6.6 (COMPILER $\gamma_{CF_{L_{ms}}}^{L_{ms}}$ IS SECURE WITH RESPECT TO MS). $\vdash \gamma_{CF_{L_{ms}}}^{L_{ms}} : \text{ms}$

With both Theorems 6.5 and 6.6 it follows from Corollary 4.3 that the two passes can be interchanged arbitrarily:

THEOREM 6.7 (COMPILERS $\gamma_{CF_{L_{ms}}}^{L_{ms}} \circ \gamma_{DCE_{L_{ms}}}^{L_{ms}}$ AND $\gamma_{CF_{L_{ms}}}^{L_{ms}} \circ \gamma_{DCE_{L_{ms}}}^{L_{ms}}$ ARE SECURE WITH RESPECT TO MS). $\vdash \gamma_{CF_{L_{ms}}}^{L_{ms}} \circ \gamma_{DCE_{L_{ms}}}^{L_{ms}} : \text{ms}$ and $\vdash \gamma_{DCE_{L_{ms}}}^{L_{ms}} \circ \gamma_{CF_{L_{ms}}}^{L_{ms}} : \text{ms}$.

6.4 Robust Strict Cryptographic Constant Time Preservation

$$\gamma_{L_{sct}}^{L_{ms}}(\text{fn } g \text{ } x := e) = \text{fn } g \text{ } x := \text{wrdoit } 1; \gamma_{L_{sct}}^{L_{ms}}(e)$$

$$\gamma_{L_{sct}}^{L_{ms}}(\text{call } g \text{ } e) = \text{call } g \text{ } \gamma_{L_{sct}}^{L_{ms}}(e); \text{wrdoit } 1$$

$$\gamma_{L_{sct}}^{L_{ms}}(e_1 \oplus e_2) = \gamma_{L_{sct}}^{L_{ms}} e_1 \oplus \gamma_{L_{sct}}^{L_{ms}} e_2$$

Given the fact that L_{sct} provides a CCT-mode that can be turned on or off, the compiler inserts wrapper code for function bodies to ensure that execution in the component always happen in this CCT-mode. The context can overwrite the flag and exit the mode, but upon invoking a function

that is part of the component, the flag would be set again. Because of this, the compiler is secure with respect to sCCT, similarly proven as in Section 6.1.

THEOREM 6.8 (COMPILER $\gamma_{L_{\text{sct}}}^{L_{\text{ms}}}$ IS SECURE WITH RESPECT TO sCCT). $\vdash \gamma_{L_{\text{sct}}}^{L_{\text{ms}}} : \text{sct}$

6.5 Robust Preservation of Intersection of Memory Safety and Strict Cryptographic Constant Time

Let $\gamma_{L_{\text{sct}}}^{L_{\text{tms}}}$ be the compiler that is the composition of $\gamma_{L_{\text{tms}}}^{L_{\text{tms}}}$, $\gamma_{L_{\text{ms}}}^L$, $\gamma_{CF_{L_{\text{ms}}}}^{L_{\text{ms}}}$, $\gamma_{DCE_{L_{\text{ms}}}}^{L_{\text{ms}}}$, and $\gamma_{L_{\text{sct}}}^{L_{\text{ms}}}$, then the following theorem holds.

THEOREM 6.9 (COMPILER $\gamma_{L_{\text{sct}}}^{L_{\text{tms}}}$ IS SECURE WITH RESPECT TO sCCT). $\vdash \gamma_{L_{\text{sct}}}^{L_{\text{tms}}} : \text{ms} \cap \text{sct}$

PROOF. From Theorem 6.4 (Compiler $\gamma_{L_{\text{tms}}}^{L_{\text{tms}}} \circ \gamma_{L_{\text{ms}}}^L$ is secure with respect to MS), we have that any L_{tms} program p compiles into a L_{ms} program p that robustly satisfies MS. Then, from Theorem 6.7 (Compilers $\gamma_{CF_{L_{\text{ms}}}}^{L_{\text{ms}}} \circ \gamma_{DCE_{L_{\text{ms}}}}^{L_{\text{ms}}}$ and $\gamma_{CF_{L_{\text{ms}}}}^{L_{\text{ms}}} \circ \gamma_{DCE_{L_{\text{ms}}}}^{L_{\text{ms}}}$ are secure with respect to MS) we have that p gets optimised to a program p' that is also MS, where the order of optimisations does not matter for p' to be MS. Assuming p' robustly satisfies sCCT, by Theorem 6.8 (Compiler $\gamma_{L_{\text{sct}}}^{L_{\text{ms}}}$ is secure with respect to sCCT) it compiles to an L_{sct} program p that robustly satisfies sCCT as well. Finally, from Theorem 4.2 (Sequential Composition of Secure Compilers) it follows that, given p robustly satisfies sCCT and MS, p also robustly satisfies sCCT and MS. \square

7 RELATED WORK

This section discusses work on robust compilation (Section 7.1) and on other secure compilation criteria (Section 7.2). Since the case study of Sections 5 and 6 implements measures for preserving MS and CCT, this section then presents relevant related work as well (Sections 7.3 and 7.4).

7.1 Secure Compilation as Robust Preservation

The robust preservation of properties as a compiler-level criterion has been analyzed extensively [Abate et al. 2021a, 2019; Patrignani et al. 2019; Patrignani and Garg 2021] and thus we build on that framework. No existing work is concerned with composing robustly safe compilers. These works consider languages with different trace models and our technical setup can be adapted to that as long as security properties and their monitors are still defined on the same trace model. The work relating robust preservation with universal composability [Patrignani et al. 2022] is closest to what this paper presents. The authors demonstrate a similar compositionality theorem to what is presented here (Section 4) but use it in the context of protocols. They do not demonstrate the scalability of the approach. Moreover, they are missing the upper and lower compositions.

7.2 Other Secure Compilation Criteria

While this paper focuses on the robust preservation framework [Abate et al. 2019], other secure compilation criteria exist. The survey on formal approaches to secure compilation [Patrignani et al. 2019] discusses a broad spectrum already, while this section presents a very high-level overview. Fully abstract compilation [Abadi 1999b] states that a compiler should preserve and reflect observational equivalence between source and target programs. It was shown [Abate et al. 2021b] that fully abstract compilers robustly preserve program properties that are either trivial or meaningless. As a mitigation for this, the authors presented a categorical approach based on maps of distributive laws [Watanabe 2002], which they call many maps of distributive laws. Maps of distributive laws have been investigated before as a possible secure compilation criterion [Tsampas et al. 2020]. Other approaches are extensions of the compiler correctness criterion as discussed in other work [Patterson and Ahmed 2019] or the introduction of opaque observations [Vu et al.

2021] to reconcile compiler optimisations with security. Note that this work also presents secure compilers that are optimising, but contrary to the other [Vu et al. 2021], provides a formal account of these in the robust preservation framework.

7.3 Memory Safety Mechanisms

Different mechanisms for enforcing memory safety exist that also consider the secure compilation domain, i.e., have an active attacker model. For example, the „pointers as capabilities” principle represents pointers as machine-level capabilities [El-Korashy et al. 2021], which behave in a similar fashion to capabilities by means of linear typing [Morrisett et al. 2005]. The approach of this paper also uses linear typing, but differs from L^3 [Morrisett et al. 2005] in the way that functions are not first-class. Moreover, this paper considers an active attacker, while the work on L^3 only discusses whole programs and, thus, has no active attacker model. The instrumentation to ensure memory safety that this paper presents is inspired by Softbounds [Nagarakatte et al. 2009]. That work inserts bounds-checks in front of pointer-dereferences and, for this to work, inserts meta-data information on pointer creation. Softbounds also works in a more advanced setting with structured fields accesses and also introduces a table-lookup for pointers that are stored in memory. This paper only considers arrays of primitive data, i.e., there are no pointers to pointers or structures. Several other approaches to memory-safety exist in literature, specifically as compiler instrumentations [Akritidis et al. 2009; Dhumbumroong and Piromsopa 2020; Jung et al. 2021; Nam et al. 2019; Shankaranarayana et al. 2023; Younan et al. 2010; Zhou et al. 2023], hardware-extensions [Chen et al. 2023; Kim et al. 2023; Kwon et al. 2013; Saileshwar et al. 2022], or programming language extensions [Benoit and Jacobs 2019; Elliott et al. 2018, 2015; Jim et al. 2002; Li et al. 2022; Weis et al. 2019; West and Wong 2005]. What differentiates this work from them is that this work uses known, compiler-based approaches to ensure memory-safety as a means to investigate secure compiler compositions. This paper does not provide efficient memory-safety, but serves as a theoretical foundation for the secure compilation domain.

To extend the languages in this paper with a less restricted form of pointer arithmetic, the region coloring memory safety monitor presented in earlier work [Michael et al. 2023] can be used. The work presenting this monitor provides an approach for the robust preservation of memory safety compiling from C to WASM. However, they do not discuss composition of secure compilers but rather investigate an instance of a secure compiler.

7.4 Cryptographic Constant Time Mechanisms

The approach to preserving cryptographic constant time in this paper is high-level, where a programming language exposes a way to switch the semantics to a data (operand) independent timing mode. Since identifiers in L_{scct} are annotated with a secrecy tag, this approach is similar to others with information flow control. For example, Vale [Bond et al. 2017] uses Dafny to ensure constant-time assembly code, while Jasmin [Almeida et al. 2017] makes use of the Coq proof assistant to reject non-constant-time programs. CT-Wasm [Watt et al. 2019] enforces constant-timeness by means of a type system. Different to the approach of this paper, these approaches necessitate that the programmer writes CCT code. An approach to allow programmers to write more high-level code is CryptOpt [Kuepper et al. 2023], which generates efficient target-code by means of a randomised search. This paper abstracts over concrete mitigation strategies and simply assumes that there is a flag to switch to a cryptographic-constant time execution mode. This can be realised by employing the FaCT [Cauligi et al. 2019] compiler, which translates common non-constant time code patterns to be constant-time, and the data (object) independent timing execution mode of modern processors.

8 CONCLUSION

This paper tackled the problem of understanding what kind of security properties does a secure compiler preserve, when said compiler is the combination of compiler passes that preserve possibly different security properties. For this, this paper first formalised security properties of interest and their composition. Then, it proved that composing secure compilers that preserve certain properties results in a secure compiler that preserves the composition of these properties. Finally, this paper defines a multi-pass compiler and proves that it preserves MS+sCCT. Crucially, this paper derives the security of the multi-pass compiler from the composition of the security properties preserved by its individual passes, which include security-preserving as well as optimisation passes.

REFERENCES

- Martín Abadi. 1999a. *Protection in Programming-Language Translations*. Springer Berlin Heidelberg, Berlin, Heidelberg, 19–34. https://doi.org/10.1007/3-540-48749-2_2
- Martín Abadi. 1999b. *Protection in Programming-Language Translations*. Springer Berlin Heidelberg, Berlin, Heidelberg, 19–34. https://doi.org/10.1007/3-540-48749-2_2
- Carmine Abate, Arthur Azevedo de Amorim, Roberto Blanco, Ana Nora Evans, Guglielmo Fachini, Catalin Hritcu, Théo Laurent, Benjamin C. Pierce, Marco Stronati, and Andrew Tolmach. 2018. When Good Components Go Bad: Formally Secure Compilation Despite Dynamic Compromise. In *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security (Toronto, Canada) (CCS '18)*. Association for Computing Machinery, New York, NY, USA, 1351–1368. <https://doi.org/10.1145/3243734.3243745>
- Carmine Abate, Roberto Blanco, Ștefan Ciobăcă, Adrien Durier, Deepak Garg, Cătălin Hrițcu, Marco Patrignani, Éric Tanter, and Jérémy Thibault. 2021a. An Extended Account of Trace-Relating Compiler Correctness and Secure Compilation. *ACM Trans. Program. Lang. Syst.* 43, 4, Article 14 (nov 2021), 48 pages. <https://doi.org/10.1145/3460860>
- Carmine Abate, Roberto Blanco, Deepak Garg, Catalin Hritcu, Marco Patrignani, and Jérémy Thibault. 2019. Journey Beyond Full Abstraction: Exploring Robust Property Preservation for Secure Compilation. In *2019 IEEE 32nd Computer Security Foundations Symposium (CSF)*. 256–25615. <https://doi.org/10.1109/CSF.2019.00025>
- Carmine Abate, Matteo Busi, and Stelios Tsampas. 2021b. Fully Abstract and Robust Compilation. In *Programming Languages and Systems*, Hakjoo Oh (Ed.). Springer International Publishing, Cham, 83–101.
- Amal Ahmed and Matthias Blume. 2011. An Equivalence-Preserving CPS Translation via Multi-Language Semantics. In *Proceedings of the 16th ACM SIGPLAN International Conference on Functional Programming (Tokyo, Japan) (ICFP '11)*. Association for Computing Machinery, New York, NY, USA, 431–444. <https://doi.org/10.1145/2034773.2034830>
- Amal Ahmed, Deepak Garg, Catalin Hritcu, and Frank Piessens. 2018. Secure Compilation (Dagstuhl Seminar 18201). *Dagstuhl Reports* 8, 5 (2018), 1–30. <https://doi.org/10.4230/DagRep.8.5.1>
- Periklis Akrkitidis, Manuel Costa, Miguel Castro, and Steven Hand. 2009. Baggy Bounds Checking: An Efficient and Backwards-Compatible Defense against out-of-Bounds Errors. In *Proceedings of the 18th Conference on USENIX Security Symposium (Montreal, Canada) (SSYM'09)*. USENIX Association, USA, 51–66.
- José Bacelar Almeida, Manuel Barbosa, Gilles Barthe, Arthur Blot, Benjamin Grégoire, Vincent Laporte, Tiago Oliveira, Hugo Pacheco, Benedikt Schmidt, and Pierre-Yves Strub. 2017. Jasmin: High-Assurance and High-Speed Cryptography. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security (Dallas, Texas, USA) (CCS '17)*. Association for Computing Machinery, New York, NY, USA, 1807–1823. <https://doi.org/10.1145/3133956.3134078>
- Arm. 2020. *Arm®A-profile Architecture Registers*. Accessed: 2023-06-09.
- Arm. 2022. *Morello for A-profile Architecture*. Accessed: 2023-06-10.
- Arthur Azevedo de Amorim, Cătălin Hrițcu, and Benjamin C. Pierce. 2018. The Meaning of Memory Safety. In *Principles of Security and Trust*, Lujo Bauer and Ralf Küsters (Eds.). Springer International Publishing, Cham, 79–105.
- Michael Backes, Cătălin Hrițcu, and Matteo Maffei. 2014. Union, intersection and refinement types and reasoning about type disjointness for secure protocol implementations. *Journal of Computer Security* 22, 2 (2014), 301–353. <https://doi.org/10.3233/jcs-130493>
- Gilles Barthe, Benjamin Grégoire, and Vincent Laporte. 2018. Secure Compilation of Side-Channel Countermeasures: The Case of Cryptographic “Constant-Time”. In *2018 IEEE 31st Computer Security Foundations Symposium (CSF)*. 328–343. <https://doi.org/10.1109/CSF.2018.00031>
- Jesper Bengtson, Karthikeyan Bhargavan, Cédric Fournet, Andrew D. Gordon, and Sergio Maffei. 2011. Refinement Types for Secure Implementations. *ACM Trans. Program. Lang. Syst.* 33, 2, Article 8 (feb 2011), 45 pages. <https://doi.org/10.1145/1890028.1890031>
- Tuur Benoit and Bart Jacobs. 2019. Uniqueness Types for Efficient and Verifiable Aliasing-Free Embedded Systems Programming. In *International Conference on Integrated Formal Methods*.

- 1226 BlockSec. 2021. CVE-2021-3564. Available from MITRE, CVE-ID CVE-2021-2564.. <http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2021-2564>
- 1227 Barry Bond, Chris Hawblitzel, Manos Kapritsos, K. Rustan M. Leino, Jacob R. Lorch, Bryan Parno, Ashay Rane, Srinath Setty, and Laure Thompson. 2017. Vale: Verifying High-Performance Cryptographic Assembly Code. In *26th USENIX Security Symposium (USENIX Security 17)*. USENIX Association, Vancouver, BC, 917–934. <https://www.usenix.org/conference/usenixsecurity17/technical-sessions/presentation/bond>
- 1228 William J. Bowman and Amal Ahmed. 2015. Noninterference for Free. *SIGPLAN Not.* 50, 9 (aug 2015), 101–113. <https://doi.org/10.1145/2858949.2784733>
- 1229 Ran Canetti, Yevgeniy Dodis, Rafael Pass, and Shabsi Walfish. 2006. Universally Composable Security with Global Setup. Cryptology ePrint Archive, Paper 2006/432. <https://eprint.iacr.org/2006/432> <https://eprint.iacr.org/2006/432>.
- 1230 Sunjay Cauligi, Gary Soeller, Brian Johannesmeyer, Fraser Brown, Riad S. Wahby, John Renner, Benjamin Grégoire, Gilles Barthe, Ranjit Jhala, and Deian Stefan. 2019. FaCT: A DSL for Timing-Sensitive Computation. In *Proceedings of the 40th ACM SIGPLAN Conference on Programming Language Design and Implementation (Phoenix, AZ, USA) (PLDI 2019)*. Association for Computing Machinery, New York, NY, USA, 174–189. <https://doi.org/10.1145/3314221.3314605>
- 1231 Dongwei Chen, Dong Tong, Chun Yang, Jiangfang Yi, and Xu Cheng. 2023. FlexPointer: Fast Address Translation Based on Range TLB and Tagged Pointers. *ACM Trans. Archit. Code Optim.* 20, 2, Article 30 (mar 2023), 24 pages. <https://doi.org/10.1145/3579854>
- 1232 Michael R. Clarkson and Fred B. Schneider. 2008. Hyperproperties. In *2008 21st IEEE Computer Security Foundations Symposium*. 51–65. <https://doi.org/10.1109/CSF.2008.7>
- 1233 Cliff Click and Keith D. Cooper. 1995. Combining Analyses, Combining Optimizations. *ACM Trans. Program. Lang. Syst.* 17, 2 (mar 1995), 181–196. <https://doi.org/10.1145/201059.201061>
- 1234 Keith D. Cooper, Philip J. Schielke, and Devika Subramanian. 1999. Optimizing for Reduced Code Space Using Genetic Algorithms. In *Proceedings of the ACM SIGPLAN 1999 Workshop on Languages, Compilers, and Tools for Embedded Systems (Atlanta, Georgia, USA) (LCTES '99)*. Association for Computing Machinery, New York, NY, USA, 1–9. <https://doi.org/10.1145/314403.314414>
- 1235 Dominique Devriese, Marco Patrignani, and Frank Piessens. 2017a. Parametricity versus the Universal Type. *Proc. ACM Program. Lang.* 2, POPL, Article 38 (dec 2017), 23 pages. <https://doi.org/10.1145/3158126>
- 1236 Dominique Devriese, Marco Patrignani, Frank Piessens, and Steven Keuchel. 2017b. Modular, Fully-abstract Compilation by Approximate Back-translation. *Logical Methods in Computer Science* Volume 13, Issue 4 (Oct. 2017). [https://doi.org/10.23638/LMCS-13\(4:2\)2017](https://doi.org/10.23638/LMCS-13(4:2)2017)
- 1237 Smith Dhumbumroong and Kerk Pirosopa. 2020. BoundWarden: Thread-enforced spatial memory safety through compile-time transformations. *Science of Computer Programming* 198 (2020), 102519. <https://doi.org/10.1016/j.scico.2020.102519>
- 1238 Akram El-Korashy, Stelios Tsampas, Marco Patrignani, Dominique Devriese, Deepak Garg, and Frank Piessens. 2021. CapablePtrs: Securely Compiling Partial Programs Using the Pointers-as-Capabilities Principle. In *2021 IEEE 34th Computer Security Foundations Symposium (CSF)*. 1–16. <https://doi.org/10.1109/CSF51468.2021.00036>
- 1239 Archibald Samuel Elliott, Andrew Ruef, Michael Hicks, and David Tarditi. 2018. Checked C: Making C Safe by Extension. In *2018 IEEE Cybersecurity Development (SecDev)*. 53–60. <https://doi.org/10.1109/SecDev.2018.00015>
- 1240 Trevor Elliott, Lee Pike, Simon Winwood, Patrick C. Hickey, James Bielman, Jamey Sharp, Eric L. Seidel, and John Launchbury. 2015. Guilt free ivory. *Proceedings of the 2015 ACM SIGPLAN Symposium on Haskell (2015)*.
- 1241 Xaver Fabian, Marco Patrignani, and Marco Guarnieri. 2022. Automatic Detection of Speculative Execution Combinations. In *Proceedings of the 29th ACM Conference on Computer and Communications Security (CCS 2022)*. ACM.
- 1242 Cédric Fournet, Andrew D. Gordon, and Sergio Maffei. 2007. A Type Discipline for Authorization Policies. *ACM Trans. Program. Lang. Syst.* 29, 5 (aug 2007), 25–es. <https://doi.org/10.1145/1275497.1275500>
- 1243 Google. [n. d.]. Android Studio Webpage. <https://developer.android.com/>. Accessed: 2023-05-30.
- 1244 Google. 2008. V8 Javascript Engine. <https://v8.dev/blog/10-years>. Accessed: 2023-05-30.
- 1245 Andrew D. Gordon and Alan Jeffrey. 2003. Authenticity by Typing for Security Protocols. *J. Comput. Secur.* 11, 4 (jul 2003), 451–519.
- 1246 Intel. 2023. *Intel®64 and IA-32 Architectures Software Developer Manual*. Accessed: 2023-06-09.
- 1247 Trevor Jim, J. Gregory Morrisett, Dan Grossman, Michael W. Hicks, James Cheney, and Yanling Wang. 2002. Cyclone: A Safe Dialect of C. In *USENIX Annual Technical Conference, General Track*.
- 1248 Tina Jung, Fabian Ritter, and Sebastian Hack. 2021. PICO: A Presburger In-Bounds Check Optimization for Compiler-Based Memory Safety Instrumentations. *ACM Trans. Archit. Code Optim.* 18, 4, Article 45 (jul 2021), 27 pages. <https://doi.org/10.1145/3460434>
- 1249 Andrew Kennedy. 2006. Securing the .NET programming model. *Theoretical Computer Science* 364, 3 (2006), 311–317. <https://doi.org/10.1016/j.tcs.2006.08.014> Applied Semantics.
- 1250 Sungkeun Kim, Farabi Mahmud, Jiayi Huang, Pritam Majumder, Chia che Tsai, Abdullah Muzahid, and Eun Jung Kim. 2023. WHISTLE: CPU Abstractions for Hardware and Software Memory Safety Invariants. *IEEE Trans. Comput.* 72 (2023),

- 1275 811–825.
- 1276 Paul C. Kocher. 1996. Timing Attacks on Implementations of Diffie-Hellman, RSA, DSS, and Other Systems. In *Advances in*
1277 *Cryptology — CRYPTO '96*, Neal Koblitz (Ed.). Springer Berlin Heidelberg, Berlin, Heidelberg, 104–113.
- 1278 Joel Kuepper, Andres Erbsen, Jason Gross, Owen Conoly, Chuyue Sun, Samuel Tian, David Wu, Adam Chlipala, Chitchanok
1279 Chuengsatiansup, Daniel Genkin, Markus Wagner, and Yuval Yarom. 2023. CryptOpt: Verified Compilation with
1280 Randomized Program Search for Cryptographic Primitives. *Proc. ACM Program. Lang.* 7, PLDI, Article 158 (jun 2023),
25 pages. <https://doi.org/10.1145/3591272>
- 1281 P.A. Kulkarni, D.B. Whalley, G.S. Tyson, and J.W. Davidson. 2006. Exhaustive optimization phase order space exploration. In
1282 *International Symposium on Code Generation and Optimization (CGO'06)*. 13 pp.–318. <https://doi.org/10.1109/CGO.2006.15>
- 1283 Albert Kwon, Udit Dhawan, Jonathan M. Smith, Thomas F. Knight, and Andre DeHon. 2013. Low-Fat Pointers: Compact
1284 Encoding and Efficient Gate-Level Implementation of Fat Pointers for Spatial Safety and Capability-Based Security. In
1285 *Proceedings of the 2013 ACM SIGSAC Conference on Computer & Communications Security (Berlin, Germany) (CCS '13)*.
Association for Computing Machinery, New York, NY, USA, 721–732. <https://doi.org/10.1145/2508859.2516713>
- 1286 Chris Lattner and Vikram Adve. 2004. LLVM: A Compilation Framework for Lifelong Program Analysis and Transformation.
1287 San Jose, CA, USA, 75–88.
- 1288 Michael LeMay, Joydeep Rakshit, Sergej Deutsch, David M. Durham, Santosh Ghosh, Anant Nori, Jayesh Gaur, Andrew
1289 Weiler, Salmin Sultana, Karanvir Grewal, and Sreenivas Subramoney. 2021. Cryptographic Capability Computing. In
1290 *MICRO-54: 54th Annual IEEE/ACM International Symposium on Microarchitecture (Virtual Event, Greece) (MICRO '21)*.
Association for Computing Machinery, New York, NY, USA, 253–267. <https://doi.org/10.1145/3466752.3480076>
- 1291 Liyi Li, Yiyun Liu, Deena Postol, Leonidas Lampropoulos, David Van Horn, and Michael Hicks. 2022. A Formal Model of
1292 Checked C. In *2022 IEEE 35th Computer Security Foundations Symposium (CSF)*. 49–63. [https://doi.org/10.1109/CSF54842.](https://doi.org/10.1109/CSF54842.2022.9919657)
1293 [2022.9919657](https://doi.org/10.1109/CSF54842.2022.9919657)
- 1294 Sergio Maffeis, Martin Abadi, Cédric Fournet, and Andy Gordon. 2008. Code-Carrying Authorization. In *13th European*
1295 *Symposium on Research in Computer Security, Málaga, Spain, October 6-8, 2008. Proceedings* (13th european symposium
1296 on research in computer security, Málaga, Spain, October 6-8, 2008. proceedings ed.), Vol. 5283. Springer Berlin Heidelberg,
563–579. <https://www.microsoft.com/en-us/research/publication/code-carrying-authorization/>
- 1297 N. Manjikian and T.S. Abdelrahman. 1997. Fusion of loops for parallelism and locality. *IEEE Transactions on Parallel and*
1298 *Distributed Systems* 8, 2 (1997), 193–209. <https://doi.org/10.1109/71.577265>
- 1299 D. McCullough. 1988. Noninterference and the Composability of Security Properties. In *2012 IEEE Symposium on Security*
1300 *and Privacy*. IEEE Computer Society, Los Alamitos, CA, USA, 177. <https://doi.org/10.1109/SECPRI.1988.8110>
- 1301 Alexandra E. Michael, Anitha Gollamudi, Jay Bosamiya, Evan Johnson, Aidan Denlinger, Craig Disselkoen, Conrad Watt,
1302 Bryan Parno, Marco Patrignani, Marco Vassena, and Deian Stefan. 2023. MSWasm: Soundly Enforcing Memory-Safe
1303 Execution of Unsafe Code. *Proc. ACM Program. Lang.* 7, POPL, Article 15 (jan 2023), 30 pages. [https://doi.org/10.1145/](https://doi.org/10.1145/3571208)
1304 [3571208](https://doi.org/10.1145/3571208)
- 1305 Microsoft. 2010a. CVE-2010-2557. Available from MITRE, CVE-ID CVE-2010-2557.. [http://cve.mitre.org/cgi-bin/cvename.](http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2010-2557)
1306 [cgi?name=CVE-2010-2557](http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2010-2557)
- 1307 Microsoft. 2010b. CVE-2011-0035. Available from MITRE, CVE-ID CVE-2011-0035.. [http://cve.mitre.org/cgi-bin/cvename.](http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2011-0035)
1308 [cgi?name=CVE-2011-0035](http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2011-0035)
- 1309 Microsoft. 2010c. CVE-2011-0036. Available from MITRE, CVE-ID CVE-2011-0036.. [http://cve.mitre.org/cgi-bin/cvename.](http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2011-0036)
1310 [cgi?name=CVE-2011-0036](http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2011-0036)
- 1311 Microsoft. 2015. CVE-2015-1770. Available from MITRE, CVE-ID CVE-2015-1770.. [http://cve.mitre.org/cgi-bin/cvename.](http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2015-1770)
1312 [cgi?name=CVE-2015-1770](http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2015-1770)
- 1313 Greg Morrisett, Amal Ahmed, and Matthew Fluet. 2005. L3: A Linear Language with Locations. In *Typed Lambda Calculi*
1314 *and Applications*, Paweł Urzyczyn (Ed.). Springer Berlin Heidelberg, Berlin, Heidelberg, 293–307.
- 1315 Santosh Nagarakatte, Jianzhou Zhao, Milo M.K. Martin, and Steve Zdancewic. 2009. SoftBound: Highly Compatible and
1316 Complete Spatial Memory Safety for c. *SIGPLAN Not.* 44, 6 (jun 2009), 245–258. <https://doi.org/10.1145/1543135.1542504>
- 1317 Santosh Nagarakatte, Jianzhou Zhao, Milo M.K. Martin, and Steve Zdancewic. 2010. CETS: Compiler Enforced Temporal
1318 Safety for C. *SIGPLAN Not.* 45, 8 (jun 2010), 31–40. <https://doi.org/10.1145/1837855.1806657>
- 1319 Myoung Jin Nam, Periklis Akritidis, and David J Greaves. 2019. FRAMER: A Tagged-Pointer Capability System with Memory
1320 Safety Applications. In *Proceedings of the 35th Annual Computer Security Applications Conference (San Juan, Puerto Rico,*
1321 *USA) (ACSAC '19)*. Association for Computing Machinery, New York, NY, USA, 612–626. [https://doi.org/10.1145/3359789.](https://doi.org/10.1145/3359789.3359799)
1322 [3359799](https://doi.org/10.1145/3359789.3359799)
- 1323 George C. Necula, Jeremy Condit, Matthew Harren, Scott McPeak, and Westley Weimer. 2005. CCured: Type-Safe Retrofitting
1324 of Legacy Software. *ACM Trans. Program. Lang. Syst.* 27, 3 (may 2005), 477–526. <https://doi.org/10.1145/1065887.1065892>
- 1325 Marco Patrignani, Amal Ahmed, and Dave Clarke. 2019. Formal Approaches to Secure Compilation: A Survey of Fully Abstract
1326 Compilation and Related Work. *ACM Comput. Surv.* 51, 6, Article 125 (feb 2019), 36 pages. <https://doi.org/10.1145/3280984>

- 1324 Marco Patrignani and Deepak Garg. 2021. Robustly Safe Compilation, an Efficient Form of Secure Compilation. *ACM Trans.*
1325 *Program. Lang. Syst.* 43, 1, Article 1 (feb 2021), 41 pages. <https://doi.org/10.1145/3436809>
- 1326 Marco Patrignani, Robert Künnemann, and Riad S. Wahby. 2022. Universal Composability is Robust Compilation.
1327 arXiv:1910.08634 [cs.PL]
- 1328 Daniel Patterson and Amal Ahmed. 2019. The next 700 Compiler Correctness Theorems (Functional Pearl). *Proc. ACM*
1329 *Program. Lang.* 3, ICFP, Article 85 (jul 2019), 29 pages. <https://doi.org/10.1145/3341689>
- 1330 Daniel Patterson and Amal J. Ahmed. 2017. Linking Types for Multi-Language Software: Have Your Cake and Eat It Too.
1331 *ArXiv abs/1711.04559* (2017).
- 1332 Gururaj Saileshwar, Rick Boivie, Tong Chen, Benjamin Segal, and Alper Buyuktosunoglu. 2022. HeapCheck: Low-Cost
1333 Hardware Support for Memory Safety. *ACM Trans. Archit. Code Optim.* 19, 1, Article 10 (jan 2022), 24 pages. <https://doi.org/10.1145/3495152>
- 1334 Michael Sammler, Deepak Garg, Derek Dreyer, and Tadeusz Litak. 2019. The High-Level Benefits of Low-Level Sandboxing.
1335 *Proc. ACM Program. Lang.* 4, POPL, Article 32 (dec 2019), 32 pages. <https://doi.org/10.1145/3371100>
- 1336 Gabriel Scherer, Max New, Nick Rioux, and Amal Ahmed. 2018. FabULous Interoperability for ML and a Linear Language.
1337 In *International Conference on Foundations of Software Science and Computation Structures (FoSSaCS) (FabOpen image in*
1338 *new windowous Interoperability for ML and a Linear Language, Vol. LNCS - Lecture Notes in Computer Science)*, Christel
1339 Baier and Ugo Dal Lago (Eds.). Springer, Thessaloniki, Greece. https://doi.org/10.1007/978-3-319-89366-2_8
- 1340 Amogha Udupa Shankaranarayana, Gopal Raveendra Soori, Michael Ferdman, and Dongyoon Lee. 2023. TAILCHECK: A
1341 Lightweight Heap Overflow Detection Mechanism with Page Protection and Tagged Pointers.
- 1342 David Swasey, Deepak Garg, and Derek Dreyer. 2017. Robust and Compositional Verification of Object Capability Patterns.
1343 *Proc. ACM Program. Lang.* 1, OOPSLA, Article 89 (oct 2017), 26 pages. <https://doi.org/10.1145/3133913>
- 1344 Stelios Tsampas, Andreas Nuyts, Dominique Devriese, and Frank Piessens. 2020. A categorical approach to secure compilation.
1345 *ArXiv abs/2004.03557* (2020).
- 1346 Thomas Van Strydonck, Frank Piessens, and Dominique Devriese. 2019. Linear Capabilities for Fully Abstract Compilation
1347 of Separation-Logic-Verified Code. *Proc. ACM Program. Lang.* 3, ICFP, Article 84 (jul 2019), 29 pages. [https://doi.org/10.](https://doi.org/10.1145/3341688)
1348 [1145/3341688](https://doi.org/10.1145/3341688)
- 1349 VMWare. 2023. CVE-2023-20892. Available from MITRE, CVE-ID CVE-2023-20892.. [http://cve.mitre.org/cgi-bin/cvename.](http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2023-20892)
1350 [cgi?name=CVE-2023-20892](http://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2023-20892)
- 1351 Son Tuan Vu, Albert Cohen, Arnaud De Grandmaison, Christophe Guillon, and Karine Heydemann. 2021. Reconciling
1352 optimization with secure compilation. *Proceedings of the ACM on Programming Languages* 5 (2021), 1 – 30.
- 1353 Hiroshi Watanabe. 2002. Well-behaved Translations between Structural Operational Semantics. In *International Workshop*
1354 *on Coalgebraic Methods in Computer Science*.
- 1355 Conrad Watt, John Renner, Natalie Popescu, Sunjay Cauligi, and Deian Stefan. 2019. CT-Wasm: Type-Driven Secure
1356 Cryptography for the Web Ecosystem. *Proc. ACM Program. Lang.* 3, POPL, Article 77 (jan 2019), 29 pages. <https://doi.org/10.1145/3290390>
- 1357 Mark N. Wegman and F. Kenneth Zadeck. 1991. Constant Propagation with Conditional Branches. *ACM Trans. Program.*
1358 *Lang. Syst.* 13, 2 (apr 1991), 181–210. <https://doi.org/10.1145/103135.103136>
- 1359 Torben Weis, Marian Waltereit, and Maximilian Uphoff. 2019. Fyr: a memory-safe and thread-safe systems programming
1360 language. *Proceedings of the 34th ACM/SIGAPP Symposium on Applied Computing* (2019).
- 1361 Richard West and Gary T. Wong. 2005. Cuckoo: a Language for Implementing Memory- and Thread-safe System Services.
1362 In *International Conference on Programming Languages and Compilers*.
- 1363 Jonathan Woodruff, Robert N.M. Watson, David Chisnall, Simon W. Moore, Jonathan Anderson, Brooks Davis, Ben Laurie,
1364 Peter G. Neumann, Robert Norton, and Michael Roe. 2014. The CHERI Capability Model: Revisiting RISC in an Age of
1365 Risk. In *Proceeding of the 41st Annual International Symposium on Computer Architecture* (Minneapolis, Minnesota, USA)
1366 (*ISCA '14*). IEEE Press, 457–468.
- 1367 Yves Younan, Pieter Philippaerts, Lorenzo Cavallaro, R. Sekar, Frank Piessens, and Wouter Joosen. 2010. PAriCheck: An
1368 Efficient Pointer Arithmetic Checker for C Programs. In *Proceedings of the 5th ACM Symposium on Information, Computer*
1369 *and Communications Security* (Beijing, China) (*ASIACCS '10*). Association for Computing Machinery, New York, NY,
1370 USA, 145–156. <https://doi.org/10.1145/1755688.1755707>
- 1371 Jie Zhou, John Criswell, and Michael Hicks. 2023. Fat Pointers for Temporal Memory Safety of C. *Proc. ACM Program. Lang.*
1372 7, OOPSLA1, Article 86 (apr 2023), 32 pages. <https://doi.org/10.1145/3586038>