

SecurePtrs: Proving Secure Compilation with Data-Flow Back-Translation and Turn-Taking Simulation

Akram El-Korashy¹ Roberto Blanco² Jérémy Thibault² Adrien Durier² Deepak Garg¹ Cătălin Hrițcu²

¹Max Planck Institute for Software Systems (MPI-SWS) ²Max Planck Institute for Security and Privacy (MPI-SP)

Abstract—Proving secure compilation of partial programs typically requires back-translating an attack against the compiled program to an attack against the source program. To prove back-translation, one can syntactically translate the target attacker to a source one—i.e., syntax-directed back-translation—or show that the interaction traces of the target attacker can also be emitted by source attackers—i.e., trace-directed back-translation.

Syntax-directed back-translation is not suitable when the target attacker may use unstructured control flow that the source language cannot directly represent. Trace-directed back-translation works with such syntactic dissimilarity because only the external interactions of the target attacker have to be mimicked in the source, not its internal control flow. Revealing only external interactions is, however, inconvenient when sharing memory via unforgeable pointers, since information about shared pointers stashed in private memory is not present on the trace. This made prior proofs unnecessarily complex, since the generated attacker had to instead stash all reachable pointers.

In this work, we introduce more informative *data-flow traces*, combining the best of syntax- and trace-directed back-translation in a simpler technique that handles both syntactic dissimilarity and memory sharing well, and that is proved correct in Coq. Additionally, we develop a novel *turn-taking simulation* relation and use it to prove a recomposition lemma, which is key to reusing compiler correctness in such secure compilation proofs. We are the first to mechanize such a recomposition lemma in the presence of memory sharing.

We use these two innovations in a secure compilation proof for a code generation compiler pass between a source language with structured control flow and a target language with unstructured control flow, both with safe pointers and components.

1 Introduction

Compiler correctness, a.k.a. semantics preservation, is the current gold standard for formally verified compilers [22, 25, 27, 36]. However, compiler correctness alone is insufficient for reasoning about the security of compiled partial programs linked with arbitrary target contexts (e.g., components such as libraries) because compiler correctness shows that the compiled program simulates the source program, only under the *assumption* that the target context obeys all restrictions of the *source language* semantics, i.e., it does not perform any low-level attacks disallowed by the source language. This assumption is usually false in practice: compiled programs are routinely linked with arbitrary, unverified target-language code that may be buggy, compromised, or outright malicious in contravention of source semantics. In these cases, compiler correctness (even in compositional form [21, 31, 42, 43]), establishes no security guarantees for compiled partial programs.

This problem can be addressed by *secure compilation* [3, 34], by enforcing that any violation of a security property of a compiled program in some target context also appears for the source program in some source context. Formally, this requires proving the existence of a property-violating source context given a target-level violation and the corresponding violating target context. This proof step, often called *back-translation*, is crucial for establishing that a vulnerable compiled program only arises from a vulnerable source program, thus preserving the security of partial source programs even against adversarial target contexts. Although there is a long line of work on proving secure compilation for prototype compilation chains that differ in the specific security properties preserved and the way security is enforced [2, 3, 5, 6, 12, 14, 15, 17, 32, 34, 34, 35, 41, 45, 47], back-translation is a common, large element of such secure compilation proofs.

Back-translation is usually done in one of two different ways: *syntax-directed* or *trace-directed*. Syntax-directed back-translation defines a function from the violating target context (a piece of syntax) to a source context, basically treating back-translation as a target-to-source compiler. While this approach is easy to use in some situations [5, 6, 14, 15, 32, 34, 41, 45, 47], it has a significant limitation: it cannot be used if some constructs of the target language cannot be easily mimicked in the source language. For example, it is not well suited when the source language only has structured control flow, while the target language has unstructured control flow (goto or jump), as representing unstructured control flow in the source would require complex transformations or rely on heuristics that may not always work [29, 52].¹ Yet this kind of a difference between source and target languages is commonplace, e.g., when compiling any block-structured language to assembly.

In contrast, trace-directed back-translation works by defining a target *interaction trace semantics* that represents all the interactions between the compiled program and its context (e.g., cross-component calls and returns) and constructing the violating source context from the violating target trace instead of the target context [2, 17, 33, 35]. This has the advantage of not having to mimic the internal behavior of the target context in the source language. So in contrast to syntax-directed back-translation, this trace-directed method works well even when

¹An alternative to representing unstructured control flow in the source could be to write an emulator for the target language *in the source language*, but we think that proving such an emulator correct in a proof assistant would be a challenging undertaking.

some target language construct cannot be easily mimicked in the source language, as long as the construct’s effect does not cross linking (program-context) boundaries.

Although very powerful in principle, trace-directed back-translation is rather understudied for settings where the program and its context can *share private memory* by passing pointers or references to each other, something that is common in practical languages like Java, Rust, and ML. There is a good reason for this relative paucity of work: memory sharing is a source of interesting interaction between the program and its context, so allowing it makes the definition of traces [17, 24], the back-translation, and the proof of secure compilation significantly more complex. Moreover, as we explain below, memory sharing changes parts of the proof conceptually and needs fundamentally new techniques.

This is precisely the gap that this paper fills: it significantly advances proofs of secure compilation from a memory-safe source language with memory sharing to a target language that provides fine-grained memory protection. For this, we introduce two new proof techniques: (1) *data-flow back-translation*, a form of back-translation that is simpler and more mechanization-friendly than the closest prior work [17], and (2) *turn-taking simulation*, which we used to adapt another key lemma in the secure compilation proof to memory sharing. Next, we briefly explain the need for these new techniques.

Data-flow back-translation Consider a compiler from a memory-safe source language that prevents pointer forging (as in Java, Rust, or ML) to a target that provides fine-grained memory protection, such as a capability machine [48, 51] or a tagged-memory architecture [13, 16]. Suppose a compiled program has shared a pointer to its private memory with the co-linked target context in the past, and the context has stored this pointer somewhere in its (i.e., the context’s) private memory. Later in the execution, the context may use a *chain of memory dereferences within its private memory* to recover this shared pointer and write through it. Since this write changes *shared* memory, it must be recorded on the trace and must be mimicked by the source context constructed by back-translation. To mimic this write in the source language, the back-translated source context cannot forge a pointer. Instead, it must follow a similar chain of dereferences in the source to the one used by the target context. However, the chain of memory dereferences leading to this pointer is in the target context’s *private* memory and interaction traces omit these private dereferences by design!

Consequently, information needed to reconstruct *how* to access the shared pointer is missing from interaction traces, which led prior work [17, 33] to have the back-translation perform complex bookkeeping in order to reconstruct this missing information. For instance, the source context generated by El-Korashy et al. [17] had to fetch all reachable pointers every time it got control and stash them in its internal state. This required complex simulation invariants, on top of the usual invariants between the states of the target and source contexts. We think this informal stashing approach is

unnecessarily complex and would be difficult to mechanize in a proof assistant (see §3.1).

To back-translate, we instead first enrich the standard interaction traces with information about data-flows *within* the context. This considerably simplifies the back-translation definition by providing precisely the missing chain of private memory dereferences in the trace itself. The data-flow back-translation function then translates each such dereference (or in general each data-flow event) one by one to simple source expressions. For each data-flow event, we prove the correctness of its back-translation, i.e., that its corresponding predefined source expression keeps *source* memory related to the *target* memory. The proof relies on just setting up an invariant between the *target* memory that now appears in each data-flow event and the *source* memory in the state after executing the source expression obtained by back-translating the given event. Crucially, proofs about stashing all reachable pointers are not needed any more.

We see data-flow back-translation as a sweet spot between standard trace-directed back-translation, which abstracts away all internal behavior of the context, and syntax-directed back-translation, which mimics the internal behavior of the context in detail, but which cannot handle syntactic dissimilarity well.

Turn-taking simulation Turn-taking simulation is useful when one tries to *reuse* compiler correctness as a lemma in the secure compilation proof to separate concerns and avoid duplicating large amounts of work. Specifically, one defines a simulation relation between the run of the *source* program in the back-translated source context on one hand and the property-violating run of the *compiled* program in the target context on the other. Some of the source and target steps are executed by the source program and its compilation and are thus already related by compiler correctness. Having to reprove the simulation for these steps would be tantamount to duplicating an involved compiler correctness proof [25]. This duplication can be avoided by proving a so-called *recomposition* lemma in the target language, as proposed by Abate et al. [2]. Intuitively, recomposition says that if a program P_1 linked with a context C_1 , and a program P_2 linked with a context C_2 both emit the same trace, then one may *recompose*—link P_1 with C_2 —to obtain again the same trace.

The proof of recomposition is a ternary simulation between the runs of $P_1 \cup C_1$, $P_2 \cup C_2$, and the recomposed program $P_1 \cup C_2$. The question that becomes nuanced with memory sharing is how should the memory of the recomposed program be related to those of the given programs in this simulation. *Without memory sharing*, this is straightforward: at any point in the simulation, the projection of P_1 ’s memory in the recomposed run of $P_1 \cup C_2$ will equal the projection of P_1 ’s memory from the run $P_1 \cup C_1$ (and dually for C_2 ’s memory). With memory sharing, however, this simple relation does not work because C_2 may change parts of P_1 ’s shared memory in ways that C_1 does not. Specifically, while control is not in P_1 , the projections of P_1 ’s memories in the two runs mentioned above will not match.

This is where our turn-taking simulation comes in. We relate the memory of P_1 from the run of $P_1 \cup C_2$ to that from the run of $P_1 \cup C_1$ only while control is in P_1 . When control shifts to the contexts (C_2 or C_1), this relation is limited to P_1 's private memory (which is not shared with the context). The picture for C_2 's memory is exactly dual. Overall, the relation takes “turns”, alternating between two memory relations depending on where the control is. This non-trivial relation allows us to prove recomposition and therefore reuse a standard compiler correctness result even with memory sharing.

Concrete setting We illustrate our two new proof techniques by extending an existing mechanized secure compilation proof by Abate et al. [2] to cover dynamic memory sharing. The compilation pass we extend goes from an imperative source language with structured control flow (e.g., calls and returns, if-then-else) to an assembly-like target with unstructured jumps. Both languages had components and safe pointers—i.e., out of bound accesses are errors that stop execution.² In both languages, the program and the context had their own private memories, and pointers to these memories could not be shared with other components. The program and the context interacted only by calling each other's functions, and passing only primitive values via call arguments and return values.

We extend both languages by allowing their safe pointers to be passed to and dereferenced by other components, thus introducing dynamic memory sharing. We then prove that this extended compilation step is secure with respect to a criterion called “robust safety preservation” [3, 4, 33]. For this, we apply our two new techniques, data-flow back-translation and turn-taking simulation. Since the parts of the proof using these new techniques are fairly involved and non-trivial, we also fully mechanize them in the Coq proof assistant.

Summary of contributions:

- We introduce data-flow back-translation and turn-taking simulation, two new techniques for proving secure compilation from memory-safe source languages to target languages with fine-grained memory protection, when the languages support memory sharing and may be syntactically dissimilar.
- We apply these conceptual techniques to prove secure compilation for a code generation pass between a source language with structured control flow and a target with unstructured control flow. Both languages have safe pointers only, and in both memory is dynamically shared by passing safe pointers between components.
- We formalize this secure compilation proof in Coq, focusing on back-translation and recomposition, which illustrate our techniques and which we fully mechanized.

Mechanized proof The Coq proof of secure compilation for the compilation pass outlined above is available as supplementary material uploaded with this submission. The size of

the back-translation and recomposition proof steps—which we fully mechanized and which constitute the biggest and most interesting parts of the proof—is 3k lines of specifications and 29k lines of proof. For comparison, in the Coq development without memory sharing on which we are building [2], these two steps were 2.7k lines in total.

As in Abate et al. [2], our mechanized secure compilation proof assumes only standard axioms (excluded middle, functional extensionality, etc.) and axioms about whole-program compiler correctness that are mostly standard and stated in the style of corresponding CompCert theorems (these are all documented in §5.3 and the included README.md). Compared to previous paper proofs of secure compilation with memory sharing [17, 33], all details of our proofs are mechanized with respect to the clear axioms mentioned above. We found that the use of a proof assistant was vital in getting all the invariants right and alleviating the human burden of checking our proof.

Outline The rest of the paper is organized as follows: In §2 we illustrate our secure compilation criterion and outline a previous proof [2] that did not support memory sharing. In §3 we explain the challenges of memory sharing, introduce data-flow back-translation and turn-taking simulation, and show how they fit into the existing proof outline. In §4 we show the source and target languages to which we apply these techniques. §5.1 and §5.2 provide details of applying data-flow back-translation and turn-taking simulation to our setting and §5.3 explains our assumptions. Finally, we discuss related work (§6), scope and limitations (§7), and future work (§8).

2 Background

We start with a motivating example explaining the broad setting we work with (§2.1), the formal secure compilation criterion we prove (called robust safety preservation; §2.2), and a proof strategy from prior work on which we build (§2.3).

2.1 Motivating Example and Setting

Broadly speaking, we are interested in the common scenario where a *part* of a program is written in a *memory-safe source language*, compiled to a target language and then linked against other target-language program parts, possibly untrusted or prone to be compromised, to finally obtain an executable target program. By “program part”, we mean a collection of components (modules), each of which contains a set of functions. These functions may call other functions, both within this part and those in other parts. We use the terms “program” and “program part” to refer to the program part we wrote and compiled, and “context” to refer to the remaining, co-linked program part that we didn't write.

As an example, consider the following source program part, a single component, `Main`, which implements a `main` function that calls two other functions `Net.init_network` and `Net.receive`, both implemented by a third-party networking library `Net` (not shown).

```
import component Net
component Main {
```

²While this is orthogonal to our current work on proof techniques, such safe pointers can be efficiently implemented using, for instance, hardware capabilities [48, 51] or programmable tagged architectures [13, 16].

```

static iobuffer[1024];
static user_balance_usd;

main () {
    Net.init_network(iobuffer);
    Net.receive();
}
}

```

Suppose that the source language is memory safe and that the program part above is compiled using a *correct* compiler to some lower-level language, then linked to a context that implements `Net.init_network` and `Net.receive`, and the resulting program is executed. Our goal is to ensure a safety property **nowrite**—that `Net.receive` never modifies the variable `user_balance_usd` (which is high integrity). Note that it is okay for `Net.receive` to modify the array `iobuffer`, whose pointer is passed as a parameter to the previous call to the `Net` library (to the function `Net.init_network`). The concern really is that a low-enough implementation of `Net.receive` may overflow the array `iobuffer` to overwrite `user_balance_usd`.

Broadly speaking, we can attain the invariant **nowrite** in at least two different ways, which we call **Setting 1** and **Setting 2**. In **Setting 1**, we compile to any target language, possibly memory-unsafe, but restrict the compilation of the program part above to be linked *only* to target-language contexts that were obtained by compiling program parts written in the same source language. Since the source language is safe, there is no way for any source function to cause a buffer overflow and a correct compiler will transfer this restriction to the target language so, in particular, the compilation of `Net.receive` cannot overwrite `user_balance_usd`, thus ensuring **nowrite**. This kind of restriction on linking—and the verification of compilers under such restrictions—has been studied extensively in *compositional compiler correctness* [21, 31, 42, 43].

In **Setting 2**, we compile to a target language with support for fine-grained memory protection, e.g., a capability machine [48, 51] or a tagged-memory architecture [13, 16], but allow target contexts to be arbitrary. The compiler uses the target language’s memory protection to defend against malicious attacks that do not necessarily adhere to source language’s memory-safety semantics. Now, **nowrite** does not follow from source memory safety and the correctness of the compiler. Instead, we must show that the compilation chain satisfies some additional security property. It is this second setting that interests us here and, more broadly, a large part of the literature on secure compilation.

2.2 Robust Safety Preservation (RSP^\sim)

The next question is what security criterion the compilation chain must satisfy to ensure that **nowrite** or, more generally, any property of interest, holds in **Setting 2**. The literature on secure compilation has proposed many such criteria (see Abate et al. [3], Patrignani et al. [34]). Here we describe and adopt one of the simplest criteria that ensures **nowrite**, namely, *robust safety preservation* or RSP^\sim [4].³

Definition 2.1 (Compilation chain has RSP^\sim [4]).

$$\begin{aligned}
 RSP^\sim &\stackrel{\text{def}}{=} \forall P \ C_t \ t. (C_t \cup P \downarrow) \rightsquigarrow^* t \\
 &\implies \exists C_s \ t'. (C_s \cup P) \rightsquigarrow^* t' \wedge t' \sim t
 \end{aligned}$$

This definition states the following: Consider any source program part⁴ P and its compilation $P \downarrow$. If $P \downarrow$ linked with some (arbitrarily chosen) target context C_t emits a finite trace prefix t , then there must exist a source context C_s that when linked to P is able to cause P to emit a related trace prefix t' .

To understand why this definition captures secure compilation in **Setting 2**, consider the case where t is a trace witnessing the violation of a safety property of interest. Then, if the compiler has RSP^\sim , there must be a source context which causes a similar violation entirely *in the source language*. In other words, an attack from some *target-level* context can only arise if the source program is vulnerable to a similar attack from some *source-level* context. In our particular example, since there clearly is no source context violating **nowrite**, RSP^\sim guarantees that no target context can violate it either.

A compilation chain attains RSP^\sim by enforcing source language abstractions against arbitrary target contexts. The specific source abstraction of interest to us here is memory safety. Our goal in this paper is to explain that *proving* RSP^\sim in the presence of memory sharing and source memory safety is difficult and to develop proof techniques for doing this. For simplicity, the concrete target language we use is memory safe, but our techniques benefit any compiler that targets a language with fine-grained memory protection. Our source and target languages differ significantly in their control flow constructs, which makes back-translation challenging.

The definition of RSP^\sim is indexed by a relation \sim between source and target traces. The concrete instantiation of this relation determines how safety properties transform from source to target [4]. In our setting, \sim is a bijective renaming relation on memory addresses, which we describe later (Definition 3.10).

2.3 A Proof Strategy for Robust Safety Preservation

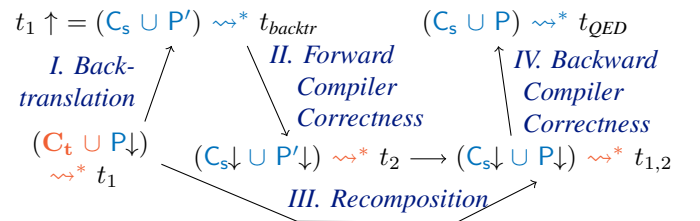


Fig. 2.1: Generic proof technique [2] for RSP^\sim . The traces t_1 , t_{backr} , t_2 , $t_{1,2}$, and t_{QED} are pairwise related by \sim .

³As a notational convention, we use different fonts and colors for **source language elements** and **target language elements**. Common elements are written in normal black font. We also use the symbol \downarrow for the compiler’s translation function.

⁴We use the notation uppercase P for a program, partial or whole, but only whole programs can execute. Whole-program execution is denoted $P \rightsquigarrow^* t$ or $P \xrightarrow{t} s$ where s is a state reached after emitting a trace prefix t .

RSP^\sim can be proved in various ways [2, 3, 33]. Here, we adapt a proof strategy by Abate et al. [2], since it *reuses* the proof of compiler correctness, thus avoiding duplication of work. Figure 2.1 summarizes the proof strategy.⁵

Overall, Abate et al. [2]’s proof of RSP^\sim consists of four steps, two of which are immediate from compiler correctness. RSP^\sim requires starting from $(C_t \cup P\downarrow) \rightsquigarrow^* t_1$ to demonstrate the existence of a C_s such that $(C_s \cup P) \rightsquigarrow^* t_{QED}$. The first proof step uses *back-translation* (Lemma 2.2) to show from $(C_t \cup P\downarrow) \rightsquigarrow^* t_1$ that there exist C_s and P' such that $(C_s \cup P') \rightsquigarrow^* t_{backtr}$ with $t_1 \sim t_{backtr}$. Note that the back-translation produces both a new context and a new program part, and that P' may be completely different from P . The second step directly uses a form of compiler correctness called forward compiler correctness (Assumption 2.3), to conclude that the compilation of this new source program, $(C_s \cup P')\downarrow = C_s\downarrow \cup P'\downarrow$, produces t_2 , related to t_1 . At this point, we have two target programs – $C_t \cup P\downarrow$ and $C_s\downarrow \cup P'\downarrow$ – that produce related traces t_1 and t_2 . The third step uses an innovative target-language lemma, *recomposition* (Lemma 2.4), to show that a third program $C_s\downarrow \cup P\downarrow$, which takes $P\downarrow$ from the first program and $C_s\downarrow$ from the second, also produces a related trace $t_{1,2}$. The final, fourth step uses another form of compiler correctness, called backward compiler correctness (Assumption 2.5), to conclude from this that the corresponding source, $C_s \cup P$ produces a related trace t_{QED} . This concludes the proof.

Lemma 2.2 (Whole-Program Back-translation [2]).

$$\forall P t. P \rightsquigarrow^* t \implies \exists P' t'. P \rightsquigarrow^* t' \wedge t' \sim t$$

Assumption 2.3 (Whole-Program Forward Compiler Correctness). $\forall P t. P \rightsquigarrow^* t \implies \exists t'. P\downarrow \rightsquigarrow^* t' \wedge t' \sim t$

Lemma 2.4 (Recomposition [2]).

$$\begin{aligned} &\forall P_1 C_1 P_2 C_2 t_1 t_2. \\ &(P_1 \cup C_1) \rightsquigarrow^* t_1 \implies (P_2 \cup C_2) \rightsquigarrow^* t_2 \implies \\ &t_1 \sim t_2 \implies \exists t_{1,2}. (P_1 \cup C_2) \rightsquigarrow^* t_{1,2} \wedge t_{1,2} \sim t_1 \end{aligned}$$

Assumption 2.5 (Whole-Program Backward Compiler Correctness). $\forall P t. P\downarrow \rightsquigarrow^* t \implies \exists t'. P \rightsquigarrow^* t' \wedge t' \sim t$

By following this proof strategy, Abate et al. [2] are able to reuse compiler correctness (Assumptions 2.3 and 2.5) and reduce the entire proof of RSP^\sim to two key lemmas: back-translation (Lemma 2.2) and recomposition (Lemma 2.4).

However, Abate et al. execute this strategy for languages without any memory sharing between components. Their components—both source and target—interact *only through integers* passed as function call arguments and return values. As such, our earlier example cannot even be expressed in their setting. In the rest of this paper, we adapt their proof strategy for RSP^\sim to the setting where memory sharing is allowed. We show that memory sharing significantly complicates the proofs

⁵Abate et al. [2] instantiate the strategy mostly for \sim set to equality, while we use a nontrivial \sim everywhere, but this difference is less important here. We also removed everything they do about undefined behavior, which we do not consider in this work (see also §8).

of both back-translation and recomposition, and requires new proof techniques. However, before explaining these, we briefly show what traces actually look like.

Interaction traces A trace or, more precisely, an interaction trace, is a modeling and proof artifact that arises from an instrumented reduction semantics of a language, wherein certain steps are labeled with descriptors called *events*. The sequence of events along a reduction sequence forms a trace, denoted t . In prior work on secure compilation, only steps involving cross-component interactions or external communication (input-output) have been labeled with events. For example, in Abate et al.’s [2] setting without shared memory, cross-component interaction happens through calls and returns only, hence, their events are only cross-component calls and returns. We denote these events e_{no_shr} where the subscript no_shr stands for “no memory sharing”.

$$e_{no_shr} ::= \text{Call } c_{caller} c_{callee}.f(v) \mid \text{Ret } c_{prev} c_{next} v$$

The event $\text{Call } c_{caller} c_{callee}.f(v)$ represents a call from component c_{caller} to the function f of component c_{callee} with argument v . The dual event $\text{Ret } c_{prev} c_{next} v$ represents a return from component c_{prev} to component c_{next} with return value v . Along a trace, calls and returns are always well-bracketed (the semantics of both the source and target languages enforce this).

In our setting, memory shared between components is another medium of interaction, so reads and writes to it must be represented on interaction traces. However, our languages are sequential (only one component executes at a time), so writes to shared memory made by a component become visible to another component only when the writing component transfers control to the other component. As such, to capture interactions between components, it suffices to record the state of the shared memory only when control transfers from one component to another, i.e., at cross-component calls and returns. For this, we modify call and return events to also record the state of the memory shared up to the time of the event (the shared part of memory grows along an execution as more pointers are passed across components). The new events, denoted e , are defined below. The shared memory on each event, written Mem , is underlined for emphasis only. Technically, Mem is a just a partial map from locations l to values v , which themselves can be pointers to locations.

Definition 2.6 (Interaction-trace events w/ memory sharing).

$$e ::= \text{Call } \underline{Mem} c_{caller} c_{callee}.f(v) \mid \text{Ret } \underline{Mem} c_{prev} c_{next} v$$

Interaction traces serve two broad purposes. First, they are used to express safety properties of interest, such as the **nowrite** property in our earlier example. Second, as we explain in §3, interaction traces are essential to the proof of back-translation, Lemma 2.2. One of our key insights is that, with memory sharing, enriching interaction traces with selective information about data-flows *within* a component can simplify the proof of back-translation considerably.

3 Key Technical Ideas

We describe why the proofs of Lemmas 2.2 and 2.4 become substantially more difficult in the presence of memory sharing, and our new techniques—data-flow back-translations and turn-taking simulations—that offset some of the extra difficulty.

3.1 Data-Flow Back-translation

In proving back-translation (Lemma 2.2), we are given a target language whole program \mathbf{P} and an interaction trace t that it produces, and we have to construct a whole source program \mathbf{P} that produces a related interaction trace t' . For RSP^\sim , we can construct \mathbf{P} from either \mathbf{P} or t . Prior work has considered both approaches. Construction of \mathbf{P} from \mathbf{P} , which we call *syntax-directed back-translation*, typically works by simulating \mathbf{P} in the source language [5, 6, 14, 15, 32, 34, 41, 45, 47]. This is tractable when every construct of the target language can be simulated easily in the source. However, as explained earlier, this is not the case for many pairs of languages including our source and target languages (§4). The alternative then is to construct \mathbf{P} from the given target trace t [2, 17, 33, 35]. This alternative, which we call *trace-directed back-translation*, should be easier in principle, since the interaction trace only records cross-component interactions, so there is no need to simulate every language construct in the source; instead, only constructs that can influence cross-component interactions need to be simulated.

Indeed, trace-directed back-translation is fairly straightforward when there is no memory sharing [2, 35] or when memory references (pointers) can be constructed from primitive data like integers in the source language. However, with memory sharing and unforgeable memory references in the source—something that is common in safe source languages like Java, Rust, Go and ML—trace-directed back-translation is really difficult. To understand this, consider the following run of the *compiled version* of our example from §2.1.

Example 3.1. Suppose we want to back-translate the following four-event target interaction trace:

```

Call Mem cMain cNet.init_network(liobuffer)
:: Ret Mem cNet cMain 0
:: Call Mem cMain cNet.receive()
:: Ret Mem' cNet cMain 0

```

where $Mem = [l_{iobuffer} \mapsto 0, l_{iobuffer} + 1 \mapsto 0, \dots, l_{iobuffer} + 1023 \mapsto 0]$ and $Mem' = [l_{iobuffer} \mapsto 4, l_{iobuffer} + 1 \mapsto 4, \dots, l_{iobuffer} + 1023 \mapsto 4]$ ⁶

In this example run, the program first shares some memory (corresponding to `iobuffer`) by calling `Net.init_network` with the pointer `liobuffer`. This call does not modify the shared memory (the shared memory’s state is Mem both before and after the call). Later the program calls the function

`Net.receive` without any arguments, but this call changes the shared memory to Mem' . (Assuming that our compiler uses the target’s memory protection correctly, this could only have happened if the `Net` library stashed the pointer `liobuffer` during the first call and retrieved it during the second call.)

The question is how we can back-translate this interaction sequence into a source program, as required by Lemma 2.2. If pointers were forgeable in the source, this would be quite easy: `liobuffer`, being forgeable, could simply be hardcoded in the body of the simulating source function `Net.receive()`. However, in our memory-safe source language, the only option is to construct a source `Net.init_network` that stashes `liobuffer` for `Net.receive`’s use. Even though this “stashing” solution may seem straightforward, it is actually quite difficult because the back-translated context must fetch and stash (e.g., in an indexed data structure) *all* pointers that become accessible to it directly or indirectly by following shared pointers, since any of these pointers *may* be dereferenced later.

Prior work [17, 33] has used such a stashing data structure. They fetch pointers by a custom graph traversal (pointers are the edges and pointed locations are the nodes) whose output is a list of source commands. Each source command is responsible for traversing a path in memory and stashing the content of the destination location in private memory (in anticipation that this stashed content might be a pointer in which case it might be needed when back-translating a later interaction event). We found that proving the correctness of this construction in a proof assistant is difficult, even though the proofs seem easy on paper. For instance, even leaving aside the correctness of this traversal (which seems rather difficult), just proving its termination is nontrivial in a proof assistant.

Note that traversing the entire shared memory is a proactive, over-approximating strategy on the part of the generated source context, by which it mimics *all possible* stashing steps that the target context *could* have made. This complex strategy was needed in prior work because information about data flows within the target context is missing from standard interaction traces, which prior work relied on. If only a trace recorded precisely which memory paths were actually traversed, we could eliminate the complexity of the full traversal. This is exactly what our new data-flow back-translation idea supports.

The new idea: data-flow back-translation We enrich the interaction traces of the target language—only for the purposes of the back-translation proof—with information about *all* data-flows, even those *within* (the private state of) a single component. We call these enriched traces *data-flow traces*. From the target language’s reduction semantics, we can easily prove that every interaction trace as described above can be enriched to a data-flow trace (Lemma 3.4 below). And, given such a data-flow trace, we can easily back-translate to a simulating source program, since we know exactly how pointers flow. In the example above, the enriched trace would tell us exactly what `Net.init_network` did to stash `liobuffer` and how `Net.receive` retrieved it later. We can then mimic this in the constructed source program, without having to stash

⁶Technically, in our languages, function calls and returns and, hence, interaction traces carry *pointers* to locations, not locations themselves. However, in this section, we blur this distinction.

all reachable pointers in memory whenever passing control to the context (see Example 3.3 below).

Concretely, we define a new type of data-flow traces, denoted T , whose events, \mathcal{E} , extend those of interaction traces to capture all possible data flows in the target language. In the following, we show the events for our target language (§4), which is a memory-safe assembly-like language with registers and memory. The events `dfCall` and `dfRet` are just the `Call` and `Ret` events of interaction traces (Definition 2.6). The remaining events correspond to target language instructions that cause data flows: loading a constant to a register (`Const`), copying from a register to another (`Mov`), binary operations (`BinOp`), copying from a register to memory or vice-versa (`Store`, `Load`) and allocating a fresh location (`Alloc`). Importantly, in a data-flow trace, every event records the entire state—both shared state and state private to individual components. Accordingly, in the events below, Mem also includes locations that were not shared to other components, and Reg is the state of the register file.

Definition 3.2 (Events of data-flow traces).

$$\begin{aligned} \mathcal{E} ::= & \text{dfCall } Mem \text{ Reg } c_{caller} \ c_{callee} \cdot \text{proc}(v) \\ & | \text{dfRet } Mem \text{ Reg } c_{prev} \ c_{next} \ v \\ & | \text{Const } Mem \text{ Reg } c_{cur} \ v \ r_{dest} \\ & | \text{Mov } Mem \text{ Reg } c_{cur} \ r_{src} \ r_{dest} \\ & | \text{BinOp } Mem \text{ Reg } c_{cur} \ op \ r_{src1} \ r_{src2} \ r_{dest} \\ & | \text{Load } Mem \text{ Reg } c_{cur} \ r_{addr} \ r_{dest} \\ & | \text{Store } Mem \text{ Reg } c_{cur} \ r_{addr} \ r_{src} \\ & | \text{Alloc } Mem \text{ Reg } c_{cur} \ r_{ptr} \ r_{size} \end{aligned}$$

Example 3.3. Consider the following data-flow trace, which expands a part of Example 3.1’s interaction trace—the part that covers the call and return to `Net.init_network()` only. Here, l is a fixed, hardcodable location that can always be accessed by `Net`, r_{COM} is a special register used to pass arguments and return values, and Mem_l and Reg_l are some initial states of memory and registers, respectively.

$$\begin{aligned} & \text{dfCall } Mem_l \ (Reg_l[r_{COM} \mapsto l_{iobuffer}]) \\ & \quad c_{Main} \ c_{Net} \cdot \text{init_network}(l_{iobuffer}) \\ :: & \text{Const } Mem_l \ (Reg_l[r_{COM} \mapsto l_{iobuffer}, r_l \mapsto l]) \ c_{Net} \ l \ r_l \\ :: & \text{Store } (Mem_l[l \mapsto l_{iobuffer}]) \\ & \quad (Reg_l[r_{COM} \mapsto l_{iobuffer}, r_l \mapsto l]) \ c_{Net} \ r_l \ r_{COM} \\ :: & \text{dfRet } (Mem_l[l \mapsto l_{iobuffer}]) \\ & \quad (Reg_l[r_{COM} \mapsto l_{iobuffer}, r_l \mapsto 0]) \ c_{Net} \ c_{Main} \ l_{iobuffer} \end{aligned}$$

This data-flow trace shows clearly how `Net.init_network` stashed away $l_{iobuffer}$: It copied $l_{iobuffer}$ to its private memory location l . The rest of the data-flow trace (not shown) will also show precisely how `Net.receive()` later retrieved $l_{iobuffer}$. It is not difficult to construct a source program that mimics these data flows step-by-step, by using source memory locations to mimic the target’s register file and memory (see §5.1 for further details). The step-by-step

mimicking induces a step-for-step inductive invariant that we found much simpler to prove than the coarse-grained invariants from prior work [17] in which the back-translation input was just the non-informative trace of Example 3.1, and the lost target steps were compensated using the full graph traversal.

Outline of data-flow back-translation proof Data-flow traces simplify the proof of back-translation (Lemma 2.2) by splitting it into two key lemmas: Enriching interaction traces to data-flow traces (Lemma 3.4) and back-translation of data-flow traces (Lemma 3.5), both of which are shown below and are much easier to prove than standard trace-directed backtranslation. Recall that T denotes a data-flow trace. $\text{remove_df}(T)$ denotes the interaction trace obtained by removing all internal data-flow events from T , i.e., by retaining only `Call` and `Return` events.

Lemma 3.4 (Enrichment).

$$\forall P \ t. \ P \rightsquigarrow^* t \implies \exists T. \ P \rightsquigarrow_{DF}^* T \wedge t = \text{remove_df}(T)$$

Proof. Immediate from the definition of the target-language semantics. \square

Lemma 3.5 (Data-flow back-translation).

$$\forall P \ T. \ P \rightsquigarrow_{DF}^* T \implies \exists P \ t. \ P \rightsquigarrow^* t \wedge t \sim \text{remove_df}(T)$$

Proof sketch. By constructing a P that simulates the data flows in T , thus keeping its state in lock-step with the state in T ’s events. See §5.1 for further details. \square

Composing these two lemmas yields Lemma 2.2.

3.2 Turn-Taking Simulation for Recomposition

Next, we turn to recomposition (Lemma 2.4). This lemma states that if two programs $P_1 \cup C_1$ and $P_2 \cup C_2$ produce two related interaction traces, then the program $P_1 \cup C_2$ can also produce an interaction trace related to both those traces. We refer to $P_1 \cup C_1$ and $P_2 \cup C_2$ as *base* programs, and to $P_1 \cup C_2$ as the *recomposed* program. We say that the partial programs P_1 and C_2 are *retained* by the recomposition, and that P_2 and C_1 are *discarded*. Traces in this section refer to the interaction traces of Definition 2.6. Data-flow traces are used only for back-translation, not for recomposition.

The proof of recomposition is a ternary simulation over executions of the three programs. For this, we need a ternary relation between a pair of states s_1 and s_2 of the base programs and a state $s_{1,2}$ of the recomposed program. The question is how we can relate the *memories* in s_1 and s_2 to that in $s_{1,2}$.

In the absence of memory sharing, as in Abate et al. [2], this is straightforward: We simply project P_1 ’s memory from s_1 , C_2 ’s memory from s_2 , put them together (take a disjoint union), and this yields the memory of $s_{1,2}$:

Definition 3.6 (Memory relation of Abate et al. [2]).⁷

$$\begin{aligned} \text{mem_rel}(s_1, s_2, s_{1,2}) & \stackrel{\text{def}}{=} \\ s_{1,2}.Mem & = (\text{proj}_{P_1}(s_1.Mem) \uplus \text{proj}_{C_2}(s_2.Mem)) \end{aligned}$$

⁷See Section 4 for the precise definition of `proj`.

However, with memory sharing, this definition no longer works, as illustrated by the following example.

Example 3.7. Consider the following three target-language components C_1 , C_2 and P_1 , represented in C-like syntax for simplicity. The fourth component P_2 is irrelevant for this explanation, hence not shown.

```
component  $C_1$  {
  int* ptr_to_P1 = malloc();
  void store(int* arg) {
    ptr_to_P1 = arg;
    int val_to_revert = *ptr_to_P1;
    *ptr_to_P1 = 42;
    ...
    *ptr_to_P1 = val_to_revert;
  }
}
```

```
component  $C_2$  {
  int* ptr_to_P1_or_P2 = malloc();
  void store(int* arg) {
    ptr_to_P1_or_P2 = arg;
  }
}
```

```
component  $P_1$  {
  int* priv_ptr = malloc();
  int* shared_ptr = malloc();
  void call_store() {
    store(shared_ptr);
  }
}
```

In the base program $P_1 \cup C_1$, P_1 shares `shared_ptr` with the function C_1 .`store()`. This function temporarily updates `shared_ptr` but reverts it to its original value before returning. Somewhat differently, in the recomposed program $P_1 \cup C_2$, C_2 .`store()` does *not* modify `shared_ptr` at all. Thus, even though the end-to-end interaction behavior of `store()` in both the programs is exactly the same, `shared_ptr` (which is actually in P_1 's memory) has been temporarily modified in C_1 .`store()` but not in C_2 .`store()`. Consequently, during the execution of the context's function `store()`, the memory relation of Definition 3.6 does not hold.

More abstractly, the problem here is that P_1 's *shared* memory in the recomposed program $P_1 \cup C_2$ can be related to that in the base program $P_1 \cup C_1$ *only while* control is in P_1 . When control is in C_2 , the contents of P_1 's shared memory can change unrelated to the base runs. This naturally leads to the following program counter-aware memory relation, where the relation \sim_{ren} captures location renaming and is formally defined later in this section.

Definition 3.8 (First attempt at our memory relation).

$$\text{mem_rel_pc}(s_1, s_2, s_{1,2}) \stackrel{\text{def}}{=}$$

if $s_{1,2}$ is executing in P_1 then:

$$\text{proj}_{P_1}(s_{1,2}.\text{Mem}) \sim_{\text{ren}} \text{proj}_{P_1}(s_1.\text{Mem})$$

else: (i.e., $s_{1,2}$ is executing in C_2)

$$\text{proj}_{C_2}(s_{1,2}.\text{Mem}) \sim_{\text{ren}} \text{proj}_{C_2}(s_2.\text{Mem})$$

Although this definition relates *shared* memory correctly, it is inadequate for P_1 's *private* memory—the memory P_1 has not shared with the context in the past, such as the pointer `priv_ptr` in Example 3.7. This private memory must remain related in the base program $P_1 \cup C_1$ and the recomposed program $P_1 \cup C_2$ independent of where the execution is. However, Definition 3.8 does not say this.

Accordingly, we revise our definition again. To determine which locations have been shared and which are still private, we rely on the interaction trace prefixes t_1 , t_2 and $t_{1,2}$ that are emitted before reaching the states s_1 , s_2 and $s_{1,2}$, respectively. For a memory *mem* and a trace *t*, we write `shared(mem, t)` for the projection of *mem* on addresses that are transitively shared on the trace *t* and `private(mem, t)` for the projection of *mem* on all the other addresses. With this, we can finally define a *turn-taking relation* `mem_rel_tt` that accurately describes the memory $s_{1,2}.\text{Mem}$ of the recomposed program in terms of the memories $s_1.\text{Mem}$ and $s_2.\text{Mem}$ of the two base programs:

Definition 3.9 (Turn-Taking Memory Relation).

$$\text{mem_rel_tt}(s_{1,2}, s_1, s_2, t_{1,2}, t_1, t_2) \stackrel{\text{def}}{=}$$

if $s_{1,2}$ is executing in P_1 then:

$$\text{mem_rel_exec}(P_1, t_1, t_{1,2}, s_1.\text{Mem}, s_{1,2}.\text{Mem}) \wedge \\ \text{mem_rel_not_exec}(C_2, t_2, t_{1,2}, s_2.\text{Mem}, s_{1,2}.\text{Mem})$$

else: (i.e., $s_{1,2}$ is executing in C_2)

$$\text{mem_rel_exec}(C_2, t_2, t_{1,2}, s_2.\text{Mem}, s_{1,2}.\text{Mem}) \wedge \\ \text{mem_rel_not_exec}(P_1, t_1, t_{1,2}, s_1.\text{Mem}, s_{1,2}.\text{Mem})$$

where

$$\text{mem_rel_exec}(\text{part}, t, t_{1,2}, \mathbf{m}_{\text{base}}, \mathbf{m}_{\text{recomp}}) \stackrel{\text{def}}{=} \\ \text{proj}_{\text{part}}(\mathbf{m}_{\text{recomp}}) \sim_{\text{ren}} \text{proj}_{\text{part}}(\mathbf{m}_{\text{base}}) \wedge \\ \text{shared}(\mathbf{m}_{\text{recomp}}, t_{1,2}) \sim_{\text{ren}} \text{shared}(\mathbf{m}_{\text{base}}, t)$$

and

$$\text{mem_rel_not_exec}(\text{part}, t, t_{1,2}, \mathbf{m}_{\text{base}}, \mathbf{m}_{\text{recomp}}) \stackrel{\text{def}}{=} \\ \text{proj}_{\text{part}}(\mathbf{m}_{\text{recomp}}) \cap \text{private}(\mathbf{m}_{\text{recomp}}, t_{1,2}) \\ \sim_{\text{ren}} \text{proj}_{\text{part}}(\mathbf{m}_{\text{base}}) \cap \text{private}(\mathbf{m}_{\text{base}}, t)$$

Intuitively, Definition 3.9 says the following about P_1 's memory: (a) While P_1 executes, P_1 's entire memory—both private and shared—is related in the runs of the base program $P_1 \cup C_1$ and the recomposed program $P_1 \cup C_2$. (b) While the contexts (C_1 and C_2) execute, only the private memory of P_1 in these two runs is related. For the context's memory, the dual relation holds. Figure 3.1 depicts this visually.

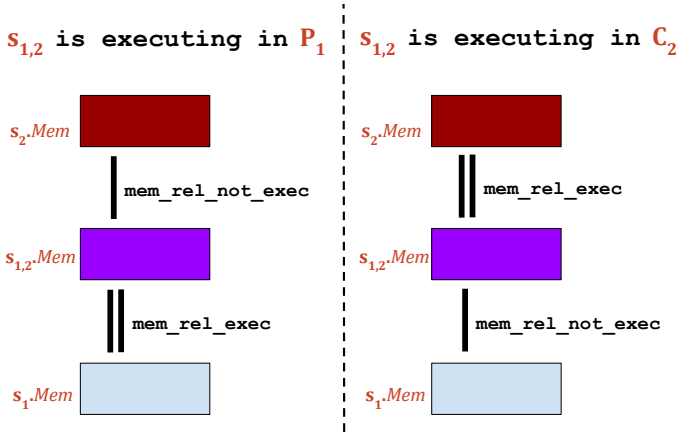


Fig. 3.1: The turn-taking memory relation, `mem_rel_tt`.

The memory relation \sim_{ren} . We now explain the memory relation \sim_{ren} that appears in the above definitions. This relation simply allows for a consistent renaming of memory locations up to a partial bijection. The need for this renaming arises because corresponding program parts may differ in the layouts of their private memories. In the example above, consider the case where the component P_2 , which we didn't show until now, is the same as P_1 , just without the private pointer `priv_ptr` and the corresponding `malloc`. In this case, the exact value of `shared_ptr` could differ across the base run $P_2 \cup C_2$ and the recomposed run $P_1 \cup C_2$. Formally, ren denotes a partial bijection that may depend on P_1, P_2, C_1 and C_2 , and \sim_{ren} is renaming of memories (both locations and their contents) up to ren .

Proof of recomposition. In our Coq proof we effectively show that the turn-taking memory relation of Definition 3.9 is an invariant of the execution of any recomposed program of the target language of §4. Formally, this follows from two lemmas (Lemmas 5.4 and 5.5) that can be seen as *expected properties of the memory relation*. Using these lemmas, we are able to prove recomposition (Lemma 2.4). A key additional idea we use is *strengthening*, which we apply at cross-component calls and returns to strengthen `mem_rel_not_exec` into `mem_rel_exec`. The former relates only the private memory of a component, the latter relates private and shared memories of the same component. Strengthening follows from the assumption that the two base runs emit related interaction traces. §5.2 provides additional details.

3.3 Applying our ideas to an RSP^\sim proof

Figure 3.2 summarizes our overall proof technique for proving RSP^\sim with dynamic memory sharing. \uparrow denotes the data-flow back-translation function. Relative to Abate et al.'s [2] proof technique shown in Figure 2.1, the two key changes are that: (1) Step I (back-translation) has now been factored into two steps Ia and Ib to use data-flow traces. Steps Ia and Ib correspond to Lemma 3.4 and Lemma 3.5, respectively. (2) The proof of step III (recomposition) now relies on

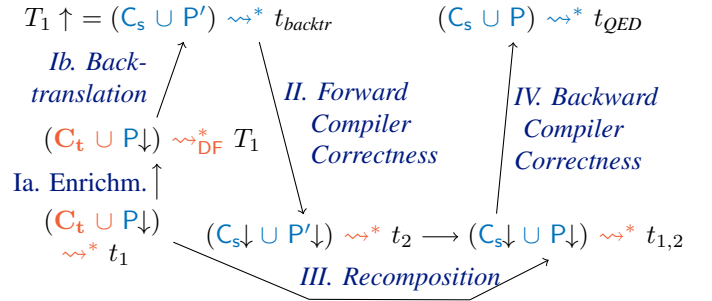


Fig. 3.2: Our proof technique for RSP^\sim with memory sharing. The interaction traces t_1 , `remove_df`(T_1), t_{backtr} , t_2 , $t_{1,2}$, and t_{QED} are pairwise related by the trace relation \sim .

turn-taking simulations. Steps II and IV, which simply reuse compiler correctness, remain unchanged.

Trace relation \sim We now also define the trace relation \sim , mentioned in §2 and §3. It says that two traces are related if corresponding events have the same kind (both call or both return, and between the same components), and there is a bijective renaming of locations ren such that the memories mentioned in corresponding events of the traces are related by \sim_{ren} (§3.2), and so are the arguments of calls and returns.

Definition 3.10 (Relation on interaction traces). For address renaming relations ren , suppose \sim_{ren} is the memory renaming relation described in §3.2.

$$t_1 \sim t_2 \stackrel{\text{def}}{=} \exists \text{ren}. \forall i. \begin{aligned} & t_1[i].\text{Mem} \sim_{\text{ren}} t_2[i].\text{Mem} \\ & \wedge \text{match_events}(t_1[i], t_2[i]) \\ & \wedge \text{valren}_{\text{ren}}(t_1[i].\text{arg}, t_2[i].\text{arg}) \end{aligned}$$

Here $t[i]$ denotes the i th event of trace t . Notation $t[i].\text{Mem}$ is the memory that appears in the event $t[i]$ (see the Definition 2.6 of events). `match_events`(e_1, e_2) says that the kind of events e_1 and e_2 (again see Definition 2.6 for the two possible kinds) and the component ids appearing on them (e.g., caller and callee) are the same. `valren_ren` is a value renaming relation that just lifts the address renaming relation ren to pointers. We give a precise definition of `valren_ren` in §4.

4 Concrete Languages and Compiler Pass

Next, we describe specific source and target languages—`SafeP` and `Mach`, respectively—and a specific compiler from the source to the target language. This specific setup is the testbed on which we have instantiated our new ideas from §3. In both languages, a program P consists of a set of named functions, a set of statically allocated data buffers and an interface. The interface divides the program into components (denoted c) and assigns every function to a component. It also defines which functions are imported and exported by each component.

Values, pointers and memory Both languages are memory safe and use the same memory model, which is adapted from CompCert's block-memory model [26]. A value v may

be an integer i , an (unforgeable) pointer, or a special error value `error` used to initialize memory.⁸ A *pointer* is a tuple $(perm, c, b, o)$ consisting of a permission $perm$ (used to distinguish code and data pointers), the identifier c of the allocating component, a unique block identifier b , and an integer offset o within the block. A location, which we denoted by l so far, is a triple of a component id, a block id, and an offset, (c, b, o) .

A memory maps locations to values. CompCert’s memory consists of an unbounded number of finite and isolated *blocks* of values. The memory in both our languages is similar, but is additionally partitioned by component ids. It can be seen as a collection $(c \mapsto cMem)$ of individual component memories $(cMem = b \mapsto (o \mapsto v))$. The projection operator that we used in Definitions 3.6 and 3.9 is formally defined as $\text{proj}_P(Mem) \stackrel{\text{def}}{=} [c \mapsto (Mem\ c) \mid c \in \text{component_ids}(P)]$, returning a sub-collection of the collection Mem containing just the component memories that correspond to the components of the program part P . Although each memory block is initially accessible to only the allocating component, memory sharing is allowed, so the *contents* (i.e., *values* v) of a component memory can be pointers to other component memories. In particular, the contents of a component memory in the collection $\text{proj}_P(Mem)$ can very well be pointers to a component memory that happens to *not* be in the collection.

Pointers can be incremented or decremented (pointer arithmetic), but this only changes the offset o . The block identifier b cannot be changed by any language operation. Additional metadata not shown here tracks the size of each allocated block. Any dereference of a data pointer with an offset beyond the allocated size or any call/jump to a code pointer with a non-zero offset causes the program to halt, which enforces memory safety. Code pointers can be shared between components, but a component cannot dereference code pointers to another component.⁹ Components interact only by calling exported functions of other components and by sharing memory.

Our languages are strongly inspired by those of Abate et al. [2] but, unlike them, we allow a component to pass pointers to other components. The receiving components can dereference these pointers, possibly after changing their offsets. However, a component cannot access a block without allocating it itself or receiving a location from it. Hence, our languages provide memory protection at block granularity.

Block ids are subject to renaming when relating two component implementations. Our memory and trace relations (Definitions 3.9 and 3.10) relate two implementations of a component even when the concrete block ids of pointers that they share with the outside world are different, as long as there exists a function¹⁰ that consistently renames the pointers shared by one implementation into those of the second. With such a block id renaming $\text{ren} : b \mapsto b$ in hand, one can

⁸Another possibility could have been to model `error` as a fixed default integer instead (like zero), so not necessarily a separate runtime type.

⁹This condition is not unrealistic and can be realized on, e.g., CHERI by implementing code pointers either as mere integer offsets or as sealed capabilities, but *not* as unsealed capabilities with execute permission.

¹⁰See the Coq file `Common/RenamingOption.v`

<code>exp ::= v</code>	values
<code> arg</code>	function argument
<code> local</code>	local static buffer
<code> exp₁ ⊗ exp₂</code>	binary operations
<code> exp₁; exp₂</code>	sequence
<code> if exp₁ then exp₂ else exp₃</code>	conditional
<code> alloc exp</code>	memory allocation
<code> !exp</code>	dereferencing
<code> exp₁ := exp₂</code>	assignment
<code> c.func(exp)</code>	function call
<code> *[exp₁](exp₂)</code>	call pointer
<code> &func</code>	function pointer
<code> exit</code>	terminate

Fig. 4.1: Syntax of source language expressions

define value renaming (which we introduced informally in Section 3.3) as follows:

- $i_1 = i_2 \implies \text{valren}_{\text{ren}}(i_1, i_2)$
- $\text{valren}_{\text{ren}}(\text{error}, \text{error})$
- $\text{ren}(b, b') \implies \text{valren}_{\text{ren}}(\text{(DATA, } c, b, o), \text{(DATA, } c, b', o))$
- $\text{valren}_{\text{ren}}(\text{(CODE, } c, b, o), \text{(CODE, } c, b, o))$

The only block-id-renaming relations we actually use in our proofs are the identity, and increment-by-1 (in Figure 5.1).

The operational semantics of both languages produce interaction traces of events from Definition 2.6, recording cross-component calls and returns. Calls and returns are necessarily well-bracketed in the semantics.

The two languages differ significantly in the constructs allowed within the bodies of functions, as we describe next.

The source language (SafeP) The body of a *SafeP* function is a single expression, `exp`, whose syntax is shown in Figure 4.1 and is inspired by the source language of Abate et al. [2]. The construct `arg` evaluates to the argument of the current function, which is a value (which may be a pointer). There are constructs for if-then-else, dereferencing a pointer (`!exp`), assigning value to a pointer (`exp1 := exp2`), calling a function `func` in component `c` with argument `exp` (`c.func(exp)`), calling a function pointer `exp1` (`*[exp1](exp2)`), and taking the address of a function (`&func`). Additionally, every component has access to a separate statically allocated memory block, whose pointer is returned by the construct `local`.

Importantly, the source language has only *structured control flow*: Calls and returns are well-bracketed by the semantics, the only explicit branching construct is if-then-else, and indirect function calls with non-zero offsets beyond function entry points are stopped by the semantics.

Function pointers exist in *SafeP* not only because they are a natural programming feature, but also to make specific steps of the back-translation convenient. Function pointers allow us, e.g., to easily mimic a store of the program counter to memory, an operation that a target-language program routinely

<code>instr ::= Const i -> r</code>	<code>Bnz r L</code>
<code>Mov r_s -> r_d</code>	<code>Jump r</code>
<code>BinOp r₁ ⊗ r₂ -> r_d</code>	<code>JumpFunPtr r</code>
<code>Label L</code>	<code>Jal L</code>
<code>PtrOffLabel L -> r_d</code>	<code>Call c func</code>
<code>Load *r_p -> r_d</code>	<code>Return</code>
<code>Store *r_p <- r_s</code>	<code>Nop</code>
<code>Alloc r₁ r₂</code>	<code>Halt</code>

Fig. 4.2: Instructions of the target language

performs. Without function pointers in the source, our cross-language value relation may have been more complex—a complexity that would propagate to the trace relation (Definition 3.10) and to the top-level theorem (Definition 2.1).

The target language (Mach) **Mach** is an assembly-like language inspired by RISC architectures, with two high-level features: the block-based memory model shared with **SafeP** and the component structure provided by interfaces. Its instructions are shown in Figure 4.2. Its state comprises a register file with a separate program counter and an abstract (protected) call stack for cross-component calls, which enforces well-bracketed cross-component **Calls** and **Returns**. A designated register `rCOM` is used for passing arguments and return values. At every cross-component call or return, all registers except `rCOM` are set to error.

Importantly, **Mach** has *unstructured control flow*: One may label statements (instruction **Label L**), jump to labeled statements (**Jump L**, **Bnz r L**), and call labeled statements (**Jal L**). Such unstructured jumps are confined to a single component (see the boxed premise of the rule for **Jump** that enforces this restriction), but may cross intra-component function boundaries. This makes it infeasible to syntactically back-translate **Mach** to **SafeP**.

JUMP

$$\frac{\text{fetch}(\mathbf{E}, pc) = \mathbf{Jump} \ r \quad \boxed{pc' = \text{reg}[r]} \quad \boxed{is_code_pointer(pc')} \quad \boxed{comp(pc) = comp(pc')}}{\mathbf{E} \vdash (\sigma, mem, reg, pc) \Downarrow (\sigma, mem, reg, pc')}$$

In addition to the interaction trace semantics (Definition 2.6) like **SafeP**'s, **Mach** also enriches the trace with data-flow events (Definition 3.2) as explained in §3.1 and illustrated by the boxed premise of the **Store** rule:

STORE

$$\frac{\text{fetch}(\mathbf{E}, pc) = \mathbf{Store} \ *r_p \ <- \ r_s \quad ptr = \text{reg}[r_p] \quad v = \text{reg}[r_s] \quad mem' = mem[ptr \mapsto v] \quad \boxed{\alpha = \mathbf{Store} \ mem' \ reg \ comp(pc) \ r_p \ r_s}}{\mathbf{E} \vdash (\sigma, mem, reg, pc) \xrightarrow{[\alpha]} (\sigma, mem', reg, pc + 1)}$$

Compiler from SafeP to Mach Our compiler from **SafeP** to **Mach** is single-pass and quite simple. It implements **SafeP**'s structured control flow with labels and direct jumps, intra-component calls using jump-and-link (**Jal**),

and function pointer calls using indirect jumps (**Jump** and **JumpFunPtr**).

Even this simple compilation chain, where security is mostly enforced by the target language semantics, suffices to bring out the difficulties in proving secure compilation in the presence of memory sharing. Using the ideas developed in §3, we have proved that this compilation chain provides RSP^\sim security.

Theorem 4.1. *Our **SafeP** to **Mach** compiler is RSP^\sim (i.e., it satisfies Definition 2.1).*

5 Some Details of the Coq Proof

5.1 Data-Flow Back-Translation of Mach

We provide some details of how we back-translate **Mach**'s data-flow traces to **SafeP**, i.e., how we prove Lemma 3.5. The back-translation function, written \uparrow , takes as input a data-flow trace T and outputs a **SafeP** whole program P that produces the (standard) trace $\text{remove_df}(T)$ in **SafeP**.¹¹ As for Abate et al. [2], each component in P maintains an *event counter* to keep track of which trace event the component is currently mimicking. This counter, as well as a small amount of other metadata used by the back-translation, is stored inside the statically allocated buffers of each component of P , which are accessed using the `local` construct.

Control flow of the result of the back-translation The outermost structure and control flow of the result of our data-flow back-translation is very similar to that of Abate et al. [2]'s interaction-trace-directed back-translation. Every procedure has a main loop (implemented using a tail-recursive call) that emits, one after the other, the events this procedure's component is responsible for emitting. In the “loop body”, the event counter mentioned above is checked using a switch statement to determine the event whose turn it is to be emitted.

Mimicking register operations A technical difficulty in the back-translation is that, unlike **Mach**, **SafeP** does not have registers. In order to mimic data-flow events involving registers, P simulates these registers and operations on them within the static buffer of the active component. For instance, a `Mov Mem Reg ccur rsrc rdest` event (which copies a value from register r_{src} to register r_{dest}) is simulated by the expression $(\text{local} + \text{OFFSET}(r_{dest}) := !(\text{local} + \text{OFFSET}(r_{src}))$, where $\text{OFFSET}(r)$ is statically expanded to the offset corresponding to register r in the simulated register file.

Mimicking memory operations Because like in **CompCert** the source and target memory models coincide, we are able to back-translate memory events quite easily. That is, a `Store` event is back-translated using assignment ($:=$) and a `Load` event is back-translated using dereferencing ($!$). Since the static buffer (whose block number is 0 in our semantics) is already used by the back-translation to store metadata and simulated

¹¹ \uparrow also takes as input the *interface* of the given target-language program to be able to mimic the same interface in the source program, but we elide the details as they are not very insightful, and largely similar to those in Abate et al. [2]. P can then be split into a context C_S and a program part P' by slicing it along this interface.

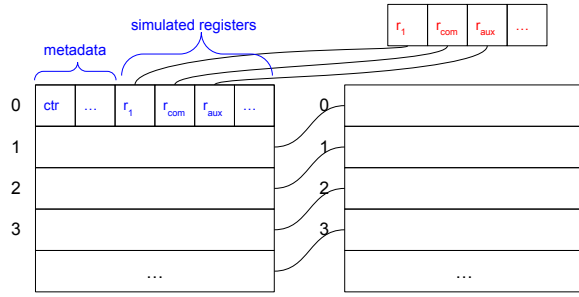


Fig. 5.1: Memory layout of a back-translated component (left) compared to a target component (right) witnessing the increment-by-1 block-id-renaming relation

registers, the back-translated program’s memory shifts by one block relative to the memory in the target: for each component, block b in the target corresponds to block $b+1$ in the source. The memory layout of the back-translated program ($P=T^\uparrow$) relative to the given **Mach** program P is shown in Figure 5.1. P maintains the invariant that, after simulating an event in T , P ’s memory and its current component’s simulated registers are synchronized with the target memory Mem and the target register file Reg mentioned in the simulated event (This is part of a mimicking_state invariant—see Lemma 5.1 below).

Mimicking calls and returns **Mach**’s semantics enforce a calling convention: calls and returns store the argument or return value in r_{COM} , and set all other registers to error. Therefore, calls and returns in P need extra administrative steps to mimic this convention. For example, mimicking a call event requires two administrative steps: (1) In the caller, dereference the content of the location simulating r_{COM} to get the argument and pass it to the function. (2) In the callee, assign the function argument arg to the location simulating r_{COM} , and set all other registers to error. Similar administrative steps are needed for mimicking a return event.

Proof of back-translation To prove back-translation (Lemma 3.5), we use a simulation lemma that ensures a relation `mimicking_state` holds between the state of P and the prefix mimicked so far. Intuitively, `mimicking_state` $T_{pref} T_{suff} s$ means that s is the state reached after mimicking all the data-flow events in T_{pref} , and that the starting state of the remaining trace T_{suff} matches s .

Lemma 5.1 (Trace-prefix mimicking).

$$\begin{aligned} \forall P T T_{pref} T_{suff}. P \rightsquigarrow_{DF}^* T \implies T = T_{pref} ++ T_{suff} \implies \\ \exists s t'_{pref}. T \uparrow \xrightarrow{t'_{pref}}^* s \wedge t'_{pref} \sim \text{remove_df}(T_{pref}) \\ \wedge \text{mimicking_state } T_{pref} T_{suff} s \end{aligned}$$

Because Lemma 5.1 ensures the relation `mimicking_state` holds for every prefix, it effectively states that the memory of the back-translation is in lock-step with the Mem and Reg appearing in each data-flow event \mathcal{E} from T . `mimicking_state`

is also strong enough to ensure that the trace relation holds between the projection of the prefix mimicked so far $\text{remove_df}(T_{pref})$ and the corresponding prefix t'_{pref} that the back-translation emits.

The fully mechanized Coq proof of Lemma 3.5 is in `Source/DefinabilityEnd.v`, which in turn uses `Source/Definability.v` and `Source/NoLeak.v`.¹²

5.2 Proof of Recomposition for **Mach**

We use the turn-taking memory relation from §3.2 to prove recomposition (Lemma 2.4). To do that, we prove that Definition 3.9 of `mem_rel_tt` is an invariant. Definition 3.9 is part of a bigger invariant `state_rel_tt` on execution states that we elide here for space reasons. The Coq proof of Lemma 2.4 is, however, available in `Intermediate/RecompositionRel.v`, which in turn uses all of `RecompositionRelCommon.v`, `RecompositionRelOptionSim.v`, `RecompositionRelLockStepSim.v` and `RecompositionRelStrengthening.v`.¹³

As explained at the end of Section 3.2, a key requirement of the recomposition proof is a strengthening lemma that recovers a stronger invariant, `state_rel_border`, which holds at states that emit interaction events. We show the memory part of `state_rel_border`:

Definition 5.2 (Memory Relation At Interaction Events).

$$\begin{aligned} \text{mem_rel_border}(s_{1,2}, s_1, s_2, t_{1,2}, t_1, t_2) \stackrel{\text{def}}{=} \\ \text{mem_rel_exec}(P_1, t_1, t_{1,2}, s_1.Mem, s_{1,2}.Mem) \wedge \\ \text{mem_rel_exec}(C_2, t_2, t_{1,2}, s_2.Mem, s_{1,2}.Mem) \end{aligned}$$

where `mem_rel_exec` is exactly as in Definition 3.9.

Among other things, `mem_rel_border` ensures that the shared memories of the three states (of the recomposed program and the two base programs) are all in sync. We are able to instantiate this strong invariant *only* at interaction events, because at these points we can use the assumption that the traces of the two base programs are related (last assumption of Lemma 5.3), which implies that the shared memories of the base programs are related. This assumption can be combined with `mem_rel_tt` (which holds universally for every triple of corresponding states) to obtain `mem_rel_border`.

Lemma 5.3 (Strengthening at interaction events).

$$\begin{aligned} \forall s_{1,2} s_1 s_2 t_{1,2} t_1 t_2 s'_1 s'_2 e_1 e_2. \\ \text{state_rel_tt}(s_{1,2}, s_1, s_2, t_{1,2}, t_1, t_2) \implies \\ s_1 \xrightarrow{[e_1]} s'_1 \implies \\ s_2 \xrightarrow{[e_2]} s'_2 \implies \\ t_1 ++ [e_1] \sim t_2 ++ [e_2] \implies \\ \exists s'_{1,2} e_{1,2}. s_{1,2} \xrightarrow{[e_{1,2}]} s'_{1,2} \wedge \\ \text{state_rel_border}(s'_{1,2}, s'_1, s'_2, \end{aligned}$$

¹²For a total of 1.3k lines of specification and 14.3k lines of proof.

¹³For a total of 830 lines of specification and 12.6k lines of proof.

$$t_{1,2} ++ [e_{1,2}], t_1 ++ [e_1], t_2 ++ [e_2])$$

The relation `state_rel_tt` is a turn-taking simulation invariant. It ensures that the memory relation `mem_rel_tt` holds of the memories of the three related states. Similarly, the stronger state relation `state_rel_border` ensures that the memory relation `mem_rel_border` holds of the memories of the three related states.

The exact definition of the relation `state_rel_tt` is in `RecompositionRelCommon.v`. We show here two key lemmas:

Lemma 5.4 (Option simulation w.r.t. *non-executing part*).

$$\begin{aligned} & \forall s_{1,2} s_1 s_2 t_{1,2} t_1 t_2 s'_1. \\ & s_{1,2} \text{ is executing in } C_2 \text{ (i.e., not in } P_1) \implies \\ & \text{state_rel_tt}(s_{1,2}, s_1, s_2, t_{1,2}, t_1, t_2) \implies \\ & s_1 \xrightarrow{\square}^* s'_1 \implies \\ & \text{state_rel_tt}(s_{1,2}, s'_1, s_2, t_{1,2}, t_1, t_2) \end{aligned}$$

The last assumption ($s_1 \xrightarrow{\square}^* s'_1$) of the option simulation (Lemma 5.4) says that state s_1 of the base program $P_1 \cup C_1$ takes some non-interaction steps. This base program contributes just P_1 to the recomposed program ($P_1 \cup C_2$), and we know by assumption “ $s_{1,2}$ is executing in C_2 ” that the recomposed state $s_{1,2}$ is *not* executing in P_1 . The invariant `state_rel_tt` ensures that $s_{1,2}$ executes in P_1 whenever s_1 executes in P_1 . Thus, the steps that s_1 has made must be taken by the discarded part C_1 , not the retained part P_1 . As shown in Example 3.7, we know that steps taken by C_1 can cause a mismatch between the memory of the recomposed program and the memory of the base program $P_1 \cup C_1$. The option simulation lemma ensures that this mismatch is tolerated by the `state_rel_tt` invariant.

Lemma 5.5 (Lock-step simulation w.r.t. *executing part*).

$$\begin{aligned} & \forall s_{1,2} s_1 s_2 t_{1,2} t_1 t_2 s'_1. \\ & s_{1,2} \text{ is executing in } C_2 \text{ (i.e., not in } P_1) \implies \\ & \text{state_rel_tt}(s_{1,2}, s_1, s_2, t_{1,2}, t_1, t_2) \implies \\ & s_2 \xrightarrow{\square} s'_2 \implies \\ & \exists s'_{1,2}. s_{1,2} \xrightarrow{\square} s'_{1,2} \wedge \\ & \text{state_rel_tt}(s'_{1,2}, s_1, s'_2, t_{1,2}, t_1, t_2) \end{aligned}$$

Lock-step simulation (Lemma 5.5) ensures that the invariant `state_rel_tt` is strong enough to keep every non-interaction step of a retained part in sync between the recomposed program and the corresponding base program.

Although both Lemmas 5.4 and 5.5 hold only for the scenario when “ $s_{1,2}$ is executing in C_2 (i.e., *not* in P_1)”, they are still general enough because we can apply symmetry lemmas to our invariant `state_rel_tt` to reduce the other scenario “ $s_{1,2}$ is executing in P_1 ” to the former scenario—thus avoiding lots of duplicate proof. The symmetry

lemmas are proved in `RecompositionRelCommon.v`. Here is the main symmetry lemma we use:

Lemma 5.6 (Symmetry of `state_rel_tt`).

$$\begin{aligned} & \forall s_{1,2} s_1 s_2 t_{1,2} t_1 t_2. \\ & \text{state_rel_tt}(s_{1,2}, s_1, s_2, t_{1,2}, t_1, t_2) \implies \\ & \text{state_rel_tt}(s_{1,2}, s_2, s_1, t_{1,2}, t_2, t_1) \end{aligned}$$

Intuitively, the two situations that Lemma 5.6 asserts are symmetric are those where the **main** function of the recomposed program $P_1 \cup C_2$ is (a) implemented by P_1 and (b) by C_2 . The symmetry lemmas consequently allow us to apply our simulation lemmas to *both* of these cases even though these simulation lemmas are proved for just one of the cases.

In `RecompositionRel.v`, the reader can find the top-level proof of recomposition (Lemma 2.4) that uses these symmetry lemmas in addition to strengthening (Lemma 5.3), option simulation (Lemma 5.4), and lock-step simulation (Lemma 5.5).

To summarize, the new idea of turn-taking simulations helped us complete the recomposition proof with memory sharing, which is fully mechanized in Coq.

5.3 Axioms

Our three novel proof steps (data-flow back-translation, Lemma 3.5, recomposition using the turn-taking simulation, Lemma 2.4, and enrichment, Lemma 3.4) are fully mechanized and rely only on standard logical axioms: proof irrelevance, functional extensionality, and classical excluded middle.

The proof of Theorem 4.1 relies (in addition to the fully mechanized lemmas above) intuitively on assumptions 2.3 and 2.5 about (separate) compilation of *whole programs* from **SafeP** to **Mach**. These kind of assumptions are axiomatized in our Coq development in a similar way to that of Abate et al. [2]. In detail: (a) we have four axioms stating that the result of compilation is syntactically well-formed if the source program is well-formed; (b) one separate compilation axiom stating that compilation and linking commute; (c) two axioms stating the existence of a forward and a backward simulation for whole-program compilation; and (d) one axiom ensuring that our compiler preserves the privacy of the local buffer. We expect this last axiom to hold because our compiler pass does not merge memory blocks.¹⁴ To prove it, we expect one can use fine-grained simulation invariants very similar to the ones one would use for a compiler correctness proof. The precise statements of these axioms are given in our `README.md` file.

One key motivation for building on the strategy of Abate et al. [2] (§2.3) is to benefit from separation of concerns between secure compilation concern and whole-program compiler correctness concern). *Axiomatizing* whole-program compiler correctness, however, is only reasonable for the purposes of a methodology-oriented case study like ours, but is not reasonable if the goal were to provide **RSP**-style *assurance* for a real system. In that case though, our methodology will enable *reuse* of the compiler correctness proof.

¹⁴If the compiler did merge blocks, then satisfying the axiom would require ensuring that it never merges a private block with a shared one.

6 Related Work

Memory relations similar to turn-taking simulations (`mem_rel_tt`) El-Korashy et al. [17] and Stewart et al. [43] use memory relations that are similar to `mem_rel_tt` in that the shared memories of two related executions may mismatch and the memory relation guarantees that the context does not modify the *private* memory of the compiled program. However, there are notable differences. First, their relations are *binary*—between two runs that differ in one component—unlike ours, which is *ternary*. This allows their relations to be strengthened whenever the compiled program is executing, while our relation can be strengthened (Definition 5.2) only for single steps right after interaction events. Second, the applications are quite different. Stewart et al. [43]’s relation is used in a non-security proof about compositional compiler correctness where guarantees come from assumptions about the target context (Setting 1, Section 2.1), while our guarantees come from memory protection features of **Mach** (Setting 2). El-Korashy et al. [17]’s memory relation is used to establish a different security criterion, full abstraction [1, 34].

Reuse of standard compiler correctness for secure compilation We are aware of only two works that reuse compiler correctness lemmas in a secure compilation proof. Abate et al. [2], which we directly build on, have goals similar to ours, but without memory sharing, which is really the focus of our paper. El-Korashy et al. [17] support memory sharing and proof reuse using a different proof technique they call *TrICL*. As explained in the paragraph on memory relations above, their memory relation (which is part of *TrICL*) is technically very different from our turn-taking simulations. Additionally, unlike our technique, their proof is not mechanically verified and, as explained in Section 3.1, mechanizing their proof is very difficult due to their use of complex bookkeeping.

Other kinds of informative traces Using inspiration from fully abstract trace semantics [24], Patrignani and Garg [33] perform back-translation (with shared memory) for a compiler pass using traces that record the *whole* memory but still only emit it at just interaction events. Although more informative than traces that record only shared memory at interaction events [17, 24], these traces still do not eliminate the need for bookkeeping, unlike our data-flow traces that selectively expose *non-interaction events* to simplify back-translation.

Handling memory sharing as message passing Patrignani et al. [35] describe a completely different secure compilation of object-oriented programs with memory sharing: Their compiled code implements shared memory in a trusted third party (realized as a hardware-protected module), and all reads and writes become explicit RPCs to this third party. Under the hood, the third party relies on dynamic sealing to hide memory addresses [28]. This effectively reduces memory sharing to message passing and elides most of the complications in proofs with true memory sharing, but also results in extremely inefficient code that requires heavyweight calls at every read-

/write to shared memory, thus largely defeating the purpose of sharing memory in the first place.

It would be interesting to study whether enforcing encapsulation while also allowing more direct memory sharing is feasible, and if so, whether the same challenges we faced still arise, and hence whether our proof techniques still apply.

Secure linking To ensure safe interaction with low-level code, typed assembly languages and multi-language semantics have been used by Patterson et al. [37]. Their technique restricts the low-level language not with runtime enforcement of memory isolation like in some architectures [16, 44, 48, 49, 50, 51] and in our **Mach** model, but instead with a static type system. The type system and the accompanying logical relation allow reasoning about the equivalence of a “mixed-language” setting, which is similar to our **Setting 2** but sometimes requires exposing low-level abstractions to high-level code. The secure compilation approach we follow has a chance of avoiding that. For example, by avoiding the need for directly reasoning in **Setting 2**, our secure compilation result beneficially hides from the programmer the fact that a **Mach** function can jump to non-entry points of other functions in the same component.

Robust safety preservation Robust safety preservation (RSP), the secure compilation criterion we use, was first described by Abate et al. [3], Patrignani and Garg [33] and Abate et al. [2]. However, this initial work uses a trivial relation (equality) between source and target traces. With a general relation, as in our setting, RSP was first examined by Abate et al. [4]. RSP further traces lineage to the robust verification of safety properties of a given program (not a given compiler), which is often called “robust satisfaction of safety properties” [23]. Robust satisfaction is a well-developed concept, used in model checking [19], type systems [18], and program logics [20, 38, 46].

Secure compilation of information-flow-like properties A long line of work [7, 8, 9, 10, 11, 30, 39, 40] develops proof techniques and verified compilers to ensure that information flow properties like non-interference, the constant-time policy, or side-channel resistance are preserved by compilation. These techniques, however, are all concerned with whole programs, unlike our work, which starts with the premise that partial programs will interact with untrusted code.

7 Scope and Limitations

We emphasize again that the key benefit of our data-flow back-translation lies in simplifying secure compilation proofs when memory is shared via pointers and the source and target languages are syntactically dissimilar. However, our technique is useful only if the target language has fine-grained memory protection, which is available in some capability architectures like **CHERI** [51] or tagged architectures [13], but not in mainstream architectures such as **x86**. Nonetheless, this is not a limitation of our *proof* technique, but rather a fundamental *enforcement* problem. Even leaving aside the

proof, we believe it is not known how to efficiently compile a memory-safe source language with fine-grained memory sharing to an architecture without support for fine-grained memory protection in a way that maintains security against arbitrary target-level contexts.

While the presentation in the paper kept the renaming relation abstract, our RSP theorem in Coq is stated only for a concrete subclass of renaming relations, which was sufficient for our particular back-translation function and our particular compiler pass. Our compiler satisfies compiler correctness for the identity renaming, meaning that it does not rename pointers. To simplify our formalization, we exploit this fact by only considering renamings that are constant shifts. We leave for future work the generalization of this subclass. Such a generalization would be needed for applying our proof technique to a more interesting compiler that needs a more complex renaming relation [25, 26]. For instance, instead of storing the whole stack in a single block, the compiler could implement the stack by allocating a new block for each stack frame. In this case the renaming relation needed for compiler correctness would relate blocks in a more subtle way than the simple identity or increment-by-1 relation. While the proof of back-translation would not be affected by such a change, with a generalized renaming relation, we will have to think more explicitly about the properties needed for the recomposition proof and the top-level proof to go through. We expect that consistency of the renaming is one such property, but there may be other properties on which we relied implicitly for our special subclass.

8 Future Work

In the future, we would like to apply our proof techniques to more realistic compilers and also to lower-level compiler passes that implement enforcement mechanisms, for instance based on capability machines [48, 51] or programmable tagged architectures [13, 16]. We also think our techniques can be extended to stronger secure compilation criteria, building on work by Abate et al. [3], who illustrate that the robust preservation of a large class of *relational* safety properties can be proved by trace-directed back-translation. This is stronger than both RSP^{\sim} and a full abstraction variant, but their back-translation technique does not yet cope with mutable state.

The languages we studied are both dynamically typed. It would be interesting to study how our proof techniques apply to secure compilers from *statically-typed* source languages too.

Another line for extending our work is to study a more realistic calling convention involving a single stack for data and control (our target language uses just a control stack and passes arguments only in registers). We expect data-flow back-translation to still be applicable, but to build such a secure compiler one will need to specify the interface of the low-level context and dynamically enforce that the low-level context's use of the stack adheres to its interface.

Finally, allowing undefined behavior, as done by Abate et al. [2], should be compatible with our techniques, as long as the *cross-component* memory operations of the source language

are compiled safely, not left completely undefined—e.g., out-of-bounds accesses to pointers shared by other components.

References

- [1] M. Abadi. Protection in programming-language translations. *ICALP*. Springer, 1998.
- [2] C. Abate, A. Azevedo de Amorim, R. Blanco, A. N. Evans, G. Fachini, C. Hritcu, T. Laurent, B. C. Pierce, M. Stronati, J. Thibault, and A. Tolmach. [When good components go bad: Formally secure compilation despite dynamic compromise](#). arXiv:1802.00588 v5 (previous version appeared at CCS'18), 2019.
- [3] C. Abate, R. Blanco, D. Garg, C. Hritcu, M. Patrignani, and J. Thibault. Journey beyond full abstraction: Exploring robust property preservation for secure compilation. *CSF*, 2019.
- [4] C. Abate, R. Blanco, Ș. Ciobâcă, A. Durier, D. Garg, C. Hritcu, M. Patrignani, É. Tanter, and J. Thibault. Trace-relating compiler correctness and secure compilation. *ESOP*. 2020.
- [5] A. Ahmed and M. Blume. [Typed closure conversion preserves observational equivalence](#). *SIGPLAN Not.*'08.
- [6] A. Ahmed and M. Blume. [An equivalence-preserving CPS translation via multi-language semantics](#). *SIGPLAN Not.*, 46(9), 2011.
- [7] G. Barthe, A. Basu, and T. Rezk. Security types preserving compilation. *VMCAI*. 2004.
- [8] G. Barthe, B. Grégoire, and V. Laporte. Secure compilation of side-channel countermeasures: the case of cryptographic “constant-time”. *CSF*. IEEE, 2018.
- [9] G. Barthe, S. Blazy, B. Grégoire, R. Hutin, V. Laporte, D. Pichardie, and A. Trieu. [Formal verification of a constant-time preserving C compiler](#). *PACMPL*, 4 (POPL), 2019.
- [10] G. Barthe, S. Blazy, R. Hutin, and D. Pichardie. [Secure compilation of constant-resource programs](#). *CSF*. 2021.
- [11] F. Besson, A. Dang, and T. Jensen. Securing compilation against memory probing. *PLAS*. ACM, 2018.
- [12] M. Busi, J. Noorman, J. Van Bulck, L. Galletta, P. Degano, J. T. Mühlberg, and F. Piessens. Provably secure isolation for interruptible enclaved execution on small microprocessors. *CSF*. IEEE, 2020.
- [13] A. A. De Amorim, M. Dénes, N. Giannarakis, C. Hritcu, B. C. Pierce, A. Spector-Zabusky, and A. Tolmach. Micro-policies: Formally verified, tag-based security monitors. *S&P*. IEEE, 2015.
- [14] D. Devriese, M. Patrignani, and F. Piessens. [Fully-abstract compilation by approximate back-translation](#). *POPL*, 2016.
- [15] D. Devriese, M. Patrignani, F. Piessens, and S. Keuchel. [Modular, fully-abstract compilation by approximate back-translation](#). *LMCS*, 13(4), 2017.
- [16] U. Dhawan, C. Hritcu, R. Rubin, N. Vasilakis, S. Chiricescu, J. M. Smith, T. F. Knight, Jr., B. C. Pierce, and A. DeHon. [Architectural support for software-](#)

- defined metadata processing. *SIGARCH Comput. Archit. News*, 43(1), 2015.
- [17] A. El-Korashy, S. Tsampas, M. Patrignani, D. Devriese, D. Garg, and F. Piessens. [CapablePtrs: Securely compiling partial programs using the pointers-as-capabilities principle](#). *CSF*, 2021.
- [18] A. D. Gordon and A. Jeffrey. Authenticity by typing for security protocols 1. *JCS*, 11(4), 2003.
- [19] O. Grumberg and D. E. Long. Model checking and modular verification. *CONCUR*. Springer, 1991.
- [20] L. Jia, S. Sen, D. Garg, and A. Datta. [A logic of programs with interface-confined code](#). *CSF*, 2015.
- [21] J. Kang, Y. Kim, C.-K. Hur, D. Dreyer, and V. Vafeiadis. [Lightweight verification of separate compilation](#). *SIGPLAN Not.*, 51(1), 2016.
- [22] R. Kumar, M. O. Myreen, M. Norrish, and S. Owens. [CakeML: a verified implementation of ML](#). *POPL*. 2014.
- [23] O. Kupferman and M. Y. Vardi. Robust satisfaction. *CONCUR*. 1999.
- [24] J. Laird. [A fully abstract trace semantics for general references](#), pages 667–679. Springer Verlag, 2007.
- [25] X. Leroy. [A formally verified compiler back-end](#). *Journal of Automated Reasoning*, 43(4), 2009.
- [26] X. Leroy and S. Blazy. [Formal verification of a C-like memory model and its uses for verifying program transformations](#). *J. Autom. Reason.*, 41(1), 2008.
- [27] R. Milner and R. Weyhrauch. Proving compiler correctness in a mechanized logic. *Machine Intelligence*, 1972.
- [28] J. H. Morris, Jr. [Protection in programming languages](#). *Commun. ACM*, 16(1), 1973.
- [29] M. O. Myreen, M. J. C. Gordon, and K. Slind. [Machine-code verification for multiple architectures: An application of decompilation into logic](#). *FMCAD*. 2008.
- [30] K. S. Namjoshi and L. M. Tabajara. Witnessing secure compilation. *VMCAI*. 2020.
- [31] G. Neis, C.-K. Hur, J.-O. Kaiser, C. McLaughlin, D. Dreyer, and V. Vafeiadis. [Pilsner: A compositionally verified compiler for a higher-order imperative language](#). *ICFP*. 2015.
- [32] M. S. New, W. J. Bowman, and A. Ahmed. [Fully abstract compilation via universal embedding](#). *ICFP*. 2016.
- [33] M. Patrignani and D. Garg. [Robustly safe compilation, an efficient form of secure compilation](#). *TOPLAS'21*.
- [34] M. Patrignani, A. Ahmed, and D. Clarke. [Formal approaches to secure compilation: A survey of fully abstract compilation and related work](#). *ACM Comput. Surv.'19*.
- [35] M. Patrignani, P. Agten, R. Strackx, B. Jacobs, D. Clarke, and F. Piessens. Secure compilation to protected module architectures. *TOPLAS*, 37(2), 2015.
- [36] D. Patterson and A. Ahmed. [The next 700 compiler correctness theorems \(functional pearl\)](#). *PACMPL*, 3 (ICFP), 2019.
- [37] D. Patterson, J. Perconti, C. Dimoulas, and A. Ahmed. [Funtal: Reasonably mixing a functional language with assembly](#). *SIGPLAN Not.*, 52(6), 2017.
- [38] M. Sammler, D. Garg, D. Dreyer, and T. Litak. [The high-level benefits of low-level sandboxing](#). *POPL'19*.
- [39] R. Sison and T. Murray. [Verifying That a Compiler Preserves Concurrent Value-Dependent Information-Flow Security](#). *ITP*. 2019.
- [40] R. Sison and T. Murray. [Verified secure compilation for mixed-sensitivity concurrent programs](#). *JFP*, 2021.
- [41] L. Skorstengaard, D. Devriese, and L. Birkedal. [StkTokens: Enforcing well-bracketed control flow and stack encapsulation using linear capabilities](#). *PACMPL*, 3 (POPL), 2019.
- [42] Y. Song, M. Cho, D. Kim, Y. Kim, J. Kang, and C.-K. Hur. [CompCertM: CompCert with C-assembly linking and lightweight modular verification](#). *PACMPL*, 4 (POPL), 2019.
- [43] G. Stewart, L. Beringer, S. Cuellar, and A. W. Appel. [Compositional compcert](#). *POPL*. 2015.
- [44] R. Strackx, J. Noorman, I. Verbauwhe, B. Preneel, and F. Piessens. Protected software module architectures. pages 241–251. Springer, 2013.
- [45] T. V. Strydomck, F. Piessens, and D. Devriese. [Linear capabilities for fully abstract compilation of separation-logic-verified code](#). *PACMLP*, 3(ICFP), 2019.
- [46] D. Swasey, D. Garg, and D. Dreyer. Robust and compositional verification of object capability patterns. *Proc. ACM Program. Lang.*, 1(OOPSLA), 2017.
- [47] S. Tsampas, D. Devriese, and F. Piessens. [Temporal safety for stack allocated memory on capability machines](#). *CSF*, 2019.
- [48] R. N. Watson, J. Woodruff, P. G. Neumann, S. W. Moore, J. Anderson, D. Chisnall, N. Dave, B. Davis, K. Gudka, B. Laurie, et al. [CHERI: A hybrid capability-system architecture for scalable software compartmentalization](#). *S&P. IEEE*, 2015.
- [49] R. N. M. Watson, R. M. Norton, J. Woodruff, S. W. Moore, P. G. Neumann, J. Anderson, D. Chisnall, B. Davis, B. Laurie, M. Roe, N. H. Dave, K. Gudka, A. Joannou, A. T. Markettos, E. Maste, S. J. Murdoch, C. Rothwell, S. D. Son, and M. Vadera. [Fast protection-domain crossing in the CHERI capability-system architecture](#). *IEEE Micro*, 36(5), 2016.
- [50] R. N. M. Watson, P. G. Neumann, J. Woodruff, M. Roe, H. Almatary, J. Anderson, J. Baldwin, D. Chisnall, B. Davis, N. W. Filardo, A. Joannou, B. Laurie, A. T. Markettos, S. W. Moore, S. J. Murdoch, K. Nienhuis, R. Norton, A. Richardson, P. Rugg, P. Sewell, S. Son, and H. Xia. [Capability Hardware Enhanced RISC Instructions: CHERI Instruction-Set Architecture \(Version 7\)](#). Technical Report UCAM-CL-TR-927, University of Cambridge, Computer Laboratory, 2019.
- [51] J. Woodruff, R. N. Watson, D. Chisnall, S. W. Moore, J. Anderson, B. Davis, B. Laurie, P. G. Neumann, R. Norton, and M. Roe. [The CHERI Capability Model: Revisiting RISC in an Age of Risk](#). *SIGARCH Comput. Archit. News*, 42(3), 2014.
- [52] F. Zhang and E. D'Hollander. [Using hammock graphs to structure programs](#). *IEEE Trans. Softw. Eng.*, 2004.