

Relational Hoare Logic for Realistically Modelled Machine Code

Denis Mazzucato^{*1}[0000-0002-3613-2035], Abdalrhman
Mohamed^{*2}[0000-0003-1414-7073], Juneyoung Lee^{**3}[0000-0002-8152-9330], Clark
Barrett²[0000-0002-9522-3084], Jim Grundy³[0009-0006-5072-9520], John
Harrison³, and Corina S. Păsăreanu¹[0000-0002-5579-6961]



¹ Carnegie Mellon University
{dmazzuca, pcorina}@andrew.cmu.edu
² Stanford University
{abdal, barrettc}@stanford.edu
³ Amazon Web Services
{lebjuney, jmgruj, jargh}@amazon.com



Abstract. Many security- and performance-critical domains, such as cryptography, rely on low-level verification to minimize the trusted computing surface and allow code to be written directly in assembly. However, verifying assembly code against a realistic machine model is a challenging task. Furthermore, certain security properties—such as constant-time behavior—require relational reasoning that goes beyond traditional correctness by linking multiple execution traces within a single specification. Yet, relational verification has been extensively explored at a higher level of abstraction. In this work, we introduce a Hoare-style logic that provides low-level, expressive relational verification. We demonstrate our approach on the s2n-bignum library, proving both constant-time discipline and equivalence between optimized and verification-friendly routines. Formalized in HOL Light, our results confirm the real-world applicability of relational verification in large assembly codebases.

Keywords: Relational Verification · Machine Code · Mechanized Proofs

1 Introduction

Verification of low-level program properties is paramount for security-critical systems. This applies to microkernels [25], where processors and other hardware may have effects that are not captured by high-level abstractions, as well as cryptographic libraries [13, 6], which aim to minimize the trusted computing base and build-toolchain dependencies. Additionally, performance-critical code is often written directly in assembly to maximize performance.

Many challenges arise when verifying low-level code, as programs execute on machines with finite memory, bounded integers, unstructured control flow,

^{*} Denis Mazzucato and Abdalrhman Mohamed contributed equally to this work.

^{**} Juneyoung Lee is the corresponding author: lebjuney@amazon.com

and memory space that is shared between data and code. In contrast, high-level verification uses abstractions that simplify reasoning and hide hardware-specific details. For instance, low-level verification must consider that primitives, when storing data, need to have free memory space. Furthermore, the verification process becomes much harder when dealing with *relational properties* [16]. Relational properties link multiple execution traces together within a single property specification. They are necessary for critical applications, such as proving that a cryptographic routine runs in constant time with respect to secret data, or that two versions of a program are functionally equivalent.

In this work, we target the `s2n-bignum` library,⁴ a cryptographic library written in assembly for ARM and X86 architectures. It includes mathematical operations on large integers, such as modular multiplication, as well as more cryptographic-oriented operations, such as elliptic curve operations. As part of the AWS’s TLS/SSL implementation, these arithmetic routines are both performance- and security-critical. The library features both highly optimized routines that are hard to verify as well as verification-friendly variants that are easier to verify but slower in practice. By verifying the latter and proving that they are functionally equivalent to the former, we can ensure that the high-performance versions do not compromise correctness. We also aim to ensure that the high-performance routines execute in constant time, a property necessary to prevent timing side-channel attacks, which could compromise sensitive data.

A number of works previously studied Hoare-style logics for realistically modelled machine code [47, 41, 36, 29, 4, 14]. Hoare-style reasoning has also been pervasively studied for relational properties [9, 44, 12]. However, relational verification for low-level code remains underexplored, especially for machine code with realistic features such as finite memory and unstructured control flow. Ideally, a binary verification toolkit would use a robust Hoare-style logic that supports realistic machine code, can express relational properties, provides sound and complete proof rules, and retains key properties that users naturally expect. These include, for instance, commutativity, as well as the ability to weaken and strengthen pre- and postconditions and to unify contracts across different contexts. Such *natural properties* enable modular reasoning, support multiple proof strategies, and make the framework practical for real-world applications. To the best of our knowledge, no existing work has presented a relational Hoare logic for realistically modelled machine code that satisfies all these properties.

In this paper, we fill that gap by introducing a novel Hoare-style logic for low-level, relational verification. Our framework, fully formalized in the HOL Light theorem prover [18, 19], offers proof rules designed to meet users’ natural expectations. We demonstrate our approach via two major case studies: in the first, we show how our framework can be used to verify constant-time behavior of various routines; in the second, we show it can be used to prove the functional equivalence of two different implementations of the same routine (e.g., one optimized for speed and the other optimized for verifiability). These case studies

⁴ <https://github.com/awslabs/s2n-bignum>

are conducted on the s2n-bignum cryptographic library. Results show that our logic scales to large assembly programs and yields practical value.

While our primary application is the s2n-bignum library, the generality of our relational Hoare logic extends beyond cryptographic code. It supports low-level features including indirect branches and self-modifying code, even though cryptographic libraries such as s2n-bignum may not employ these features.

We summarize our contributions as follows:

- i) A novel relational Hoare logic tailored to realistically modelled machine code, formalized in HOL Light.
- ii) A first case study on constant-time behavior of cryptographic routines in the s2n-bignum library, including the copy and modular inversion routines.
- iii) A second case study involving equivalence proofs between optimized and verification-friendly implementations of s2n-bignum routines.

2 Related Work

Hoare-Style Reasoning for Realistically Modelled Machine Code. Verifying realistically modelled machine code is challenging due to unstructured control flow, which traditional Hoare logics [20] struggle to handle. While Affeldt [2] uses Hoare logic to verify low-level arithmetic routines, their work is limited to assembly fragments with structured control flow. Several approaches address unstructured control flow, such as the inductive assertion method [38, Section 2] used by Barthe et al. [7] and Lehner and Muller [28] to generate verification conditions, and the program logic by Tan et al. [46] based on continuation-passing style reasoning [8]. However, these methods often fail to specify pre- and postconditions with shared continuation labels.

Other notable efforts include a logic for total correctness of communicating unstructured programs [4, 21, 22], formalized in Isabelle/HOL, and one for reasoning about MIPS assembly in Coq [30]. However, these logics are compositional only for nonoverlapping fragments. Full compositionality is critical for modular reasoning, which, in turn, is needed for scalability.

Wang [47] proposes a logic for total correctness of unstructured programs with multi-exit postconditions, but it does not guarantee postconditions upon first encounter. Unstructured programs may, in fact, go through the last instruction, jump back, and then meet the postcondition later. While this might seem misleading, we argue that functional specifications are generally confined to function boundaries, where the final instruction is a return statement. This ensures the program cannot continue execution and revisit the postcondition later, effectively solving the issue of first-met postconditions.

Myreen and Gordon [37] introduce a logic for unstructured code applied in the verified CakeML compiler [27], leveraging decompilation into logic [36, 33, 34, 32, 35]. Despite its major impact in verifying seL4 compilation [45] and realistic executables [46], it lacks a conjunction rule, a property that is naturally expected from a Hoare-style logic in order to unify contracts over different postconditions. The lack of a conjunction rule significantly increases the proof burden.

For instance, this assembly program increments `x0` by 1 until it reaches 3, then halts. Let $P = (x0 = 1)$, $Q = (x0 = 2)$, and $Q' = (x0 = 3)$. The program satisfies Q and Q' separately, on line 3, after two and three iterations, respectively, but Q and Q' cannot possibly hold simultaneously. In Section 5, we show how to handle such cases.

```

0  mov x0,xzr
1  loop:
2  add x0,x0,#1
3  cmp x0,#3
4  bne loop

```

Ray et al. [42] address conjunction rules and first-met postconditions by tracking execution steps, while Lundberg et al. [29] extend this to handle multi-exit locations, ensuring postconditions hold at the first encounter. However, their approaches assume deterministic semantics, incompatible with architectures like x86. Jensen et al. [23] addresses this gap by using separation logic [39, 43] but only for a subset of x86 code. Furthermore, EverCrypt [41] verifies cryptographic primitives using a Hoare-style logic through C and assembly code interoperability, but it does not support ARM architecture. Similar to our embedding of relational to unary Hoare triples, exploiting the event list, EverCrypt is able to prove constant-time. Fiat-Crypto [17] generates verified cryptographic code from high-level specifications with applications to big-number arithmetic, in scope similar to the s2n-bignum library. All these approaches lack a robust foundation for generic relational properties—not only the ones reducible to unary properties.

Relational Hoare Logics. Relational Hoare logics [38] extend unary Hoare triples to reason about multiple execution traces. Benton [9] gives a relational Hoare logic for two execution traces, later generalized by Blatter et al. [12] to any number of traces. While Benton also covers relational properties of low-level unstructured code [8], their logic relies on an idealized computational model. We propose a generic framework for manual proof of relational properties.

In credible compilation, Rinard [44] developed relational logics for pointer allocation, and Benton [10] proposes a sound-but-incomplete fully automatic tool for equivalence preservation of compiled programs with minor differences. They verify HHVM bytecode, which is not a high-level language but not as low as assembly; for instance, they do not handle physical registers. Instead, our approach sacrifices automation, gaining both soundness and completeness.

Barthe et al. [5] propose product program constructions for equivalence reasoning, later extended [11, 3] to support equivalence across multiple programs. Kang et al. [24] describe a relational logic for LLVM code which lacks support for indirect branches and self-modifying code. Pit-Claudiel et al. [40] further extend relational verification to low-level stack machines. Our logic handles relational properties but does not trade off any low level features of assembly code.

3 Running Example

We use the program `compare`, shown in Figure 1 (left), as a running example for the rest of the paper. It compares byte-by-byte the contents of a key buffer k and a data buffer x of length n . It takes as input the buffer length n and the memory addresses of the buffers k and x , provided via the registers `n`, `k`, and

<p>Program 1.1. compare</p> <pre style="border: 1px solid black; padding: 10px; margin: 0;"> 1 cbz n, eq 2 loop: 3 sub n, n, #1 4 ldr kn, [k, n, lsl #3] 5 ldr xn, [x, n, lsl #3] 6 cmp kn, xn 7 bne neq 8 cbnz n, loop 9 eq: 10 mov res, #1 11 ret 12 neq: 13 mov res, xzr 14 ret </pre>	<p>Program 1.2. cst-compare</p> <pre style="border: 1px solid black; padding: 10px; margin: 0;"> 1 mov diff, xzr 2 cbz n, end 3 loop: 4 sub n, n, #1 5 ldr kn, [k, n, lsl #3] 6 ldr xn, [x, n, lsl #3] 7 eor temp, kn, xn 8 orr diff, diff, temp 9 cbnz n, loop 10 end: 11 cmp diff, xzr 12 cset res, eq 13 ret </pre>
---	---

Fig. 1. Two programs to perform byte-by-byte comparison of buffers. The program `compare` (left) is not constant-time, while the program `compare-constant` (right) is.

x , respectively. Temporary values are stored in the registers `kn`, `xn`, `diff`, and `temp`, while the result is stored in `res`. The private data is the content of the key buffer. The program `compare` iterates backwards, comparing corresponding elements from both buffers. If a mismatch is detected, the program jumps to the label `neq` and sets `res` to 0. Otherwise, if all elements match, it reaches the label `eq` and sets `res` to 1. This behavior results in variable execution time depending on the buffer contents. An attacker can exploit this timing variation to deduce the position of mismatches and reconstruct the secret buffer k in linear time.

To address this issue, program `cst-compare` in Figure 1 (right) implements a constant-time comparison. It always iterates over the entire buffer length, regardless of mismatches, accumulating possible differences in `diff`. The program’s execution time is constant for any buffer content, ensuring no timing leaks and preventing attackers from inferring secret information. Furthermore, the two programs are functionally equivalent.

4 Unary Hoare Logic \mathcal{L}_1

This section provides background on an unary Hoare logic used in the verification framework described in [37, 29]. We refer to this logic as \mathcal{L}_1 . The definitions and theorems of \mathcal{L}_1 have been fully mechanized in HOL Light by previous researchers.

States. Let Σ denote a set of machine states, represented as the set of functions mapping observable resources \mathbb{L} (e.g., memory, registers, program counter) to their values. For example, in the ARM architecture, the resources \mathbb{L}_{ARM} include a 64-bit program counter `pc`, 32 general-purpose registers `regsi`, flags `flagsk`, and memory `memoryh`, indexed accordingly by i , k , and h . Similarly, the x86 architecture has resources like an instruction pointer (`rip`) and extended flags. To generalize across different architectures, the label `instr` refers to the address of the next instruction, where `instr = pc` for ARM and `instr = rip` for x86.

We use $s(l)$ for the value of resource $l \in \mathbb{L}$ in the state $s \in \Sigma$ and $s[l \mapsto v]$ for the updated state. Resource values depend on the architecture; e.g., for ARM, $s(\text{pc}) \in \text{int64}$, $s(\text{regs}_i) \in \text{int64}$, and $s(\text{memory}_h) \in \text{byte}$, where $\text{byte} \stackrel{\text{def}}{=} \{0, 1\}^8$ and $\text{int}n \stackrel{\text{def}}{=} \{0, 1\}^n$.

Properties. A property P is a subset of machine states Σ . A state s satisfies the property $P \subseteq \Sigma$ if $s \in P$. The execution of a single instruction is modelled as the small-step operational semantics $\tau \subseteq \Sigma \times \Sigma$, where $s \xrightarrow{\tau} s'$ describes the fetch-decode-execute cycle, updating state s to s' and advancing instr . The composition of two relations τ_1, τ_2 is defined as $\tau_1 \circ \tau_2 \stackrel{\text{def}}{=} \{(s, s'') \mid \exists s'. s \xrightarrow{\tau_1} s' \xrightarrow{\tau_2} s''\}$. The n -th composition of τ is τ^n . The decoding function $\text{DECODE}^\tau(s, i)$ maps bytes in the memory at address i either to an instruction or to \perp if undecodable. ARM instructions have a length of 4 bytes and are 4-byte aligned. x86 instructions have variable lengths. We write $\text{LENGTH}^\tau(C)$ for the number of bytes that the program C occupies in the memory without padding for alignment. Execution halts when the first undecodable instruction is encountered, denoted by $\text{END}^\tau(s, i) \stackrel{\text{def}}{\iff} s(\text{instr}) = i \wedge \text{DECODE}^\tau(s, i) = \perp$.

We use $\text{ALIGN}^\tau(s, i_0, C)$ in this paper to denote that the program C is stored in memory starting from the address i_0 where $s(\text{instr}) = i_0$ and i_0 satisfies the alignment constraint of a program if the architecture is ARM. The predicate $\text{ALIGN}^\tau(s, i_0, C)$ may appear as a conjunctive clause in P to describe the program of interest. The notation $\text{prog}(P)$ refers to the program C constrained by P .

The eventually Property. Assume that a machine state $s \in \Sigma$ satisfies a precondition. A postcondition $Q \subseteq \Sigma$ must eventually hold after a finite number of steps from s . To represent such s , $\text{eventually}^\tau(Q)$ defines the set of states from which Q eventually holds along every possible path through τ .

Definition 1 (Eventually). *Given an operational semantics $\tau \subseteq \Sigma \times \Sigma$ and a property $Q \subseteq \Sigma$, the property $\text{eventually}^\tau(Q) \subseteq \Sigma$ is defined inductively as:*

$$\frac{s \in Q}{s \in \text{eventually}^\tau(Q)} \quad \frac{\exists s'. s \xrightarrow{\tau} s' \quad \forall s'. s \xrightarrow{\tau} s' \implies s' \in \text{eventually}^\tau(Q)}{s \in \text{eventually}^\tau(Q)}$$

The second inference rule expands $\text{eventually}^\tau(Q)$ if every next state s' is in $\text{eventually}^\tau(Q)$. This notion of *eventually* is essential for reasoning about nondeterministic operational semantics, such as in x86, where certain instructions exhibit nondeterministic behavior. For instance, the `mul` instruction⁶ nondeterministically sets the `SF` flag to either 0 or 1. A simplified small-step semantics for `mul` is as follows:

$$\frac{s(\text{instr}) = i \quad \text{DECODE}^{\tau_{\text{x86}}}(s, i) = \text{mul } r \quad r \in \text{int16} \quad s(r) = x \quad sf \in \{0, 1\}}{s \xrightarrow{\tau_{\text{x86}}} s \text{ [EAX} \mapsto s(\text{AX}) \cdot x, \text{ SF} \mapsto sf, \text{ instr} \mapsto i + \text{LENGTH}^\tau(\text{mul } r)]} \text{ MUL}}$$

⁵ We omit the τ symbol when the operational semantics is clear from the context.

⁶ <https://www.felixcloutier.com/x86/mul#flags-affected>

Example 1. Consider a program C consisting of the instructions `mul ax, sets dl,` and `imul edx, eax`. The program starts at instruction register i_0 , where `AX` (the least significant 16 bits of `EAX`) is multiplied by itself, setting `SF` to 0 or 1. In the second instruction, the least significant byte of `EDX`, referred to as `DL`, is set to `SF`. After then, the values of `EAX` and `EDX` are multiplied, and the truncated result up to 32 bits is stored in `EDX`. The program terminates at $i_0 + 9$ because each of the x86 instructions is 3 bytes, and `EDX` is equal to either 0 or x^2 . The postcondition can be expressed as: $\{s' \mid \text{END}^{\tau_{\text{x86}}}(s', i_0 + 9) \wedge (s'(\text{EDX}) = x^2 \vee s'(\text{EDX}) = 0)\}$. It holds under a precondition requiring `AX` to initially be equal to x and `EDX` to 0.

Unary Hoare Triple. Reasoning about machine code differs from reasoning about high-level languages in several ways. First, the machine code is not represented as a syntactic program but instead as a set of instructions in the memory space. Second, a machine code may modify anything during its execution, including itself and callee-save registers. To denote unwanted modifications after the program execution, a *frame condition* $F \subseteq \Sigma \times \Sigma$ bounds allowed changes of state components between the input and output states. We explain the formal definition of the predicate `ensures` which is the Hoare triple in \mathcal{L}_1 . The notation used in this paper follows the convention you may find in the `s2n-bignum` library.

Definition 2 (Ensures). *Given an operational semantics $\tau \subseteq \Sigma \times \Sigma$, a precondition $P \subseteq \Sigma$, a postcondition $Q \subseteq \Sigma$, and a frame condition $F \subseteq \Sigma \times \Sigma$, we define the predicate `ensures` $^\tau(P, Q, F)$ as follows:*

$$\begin{aligned} \text{ensures}^\tau(P, Q, F) &\stackrel{\text{def}}{\iff} \\ \forall s. s \in P &\implies s \in \text{eventually}^\tau(\{s' \mid s' \in Q \wedge (s, s') \in F\}) \end{aligned}$$

Example 2. Consider the program C from Example 1, starting from the precondition $\{s \mid \text{ALIGN}^{\tau_{\text{x86}}}(s, i_0, C) \wedge s(\text{AX}) = x \wedge s(\text{EDX}) = 0\}$, ensuring that the memory is aligned with C and `AX` is equal to x . By application of the operational semantics, we eventually satisfy the postcondition where C terminates with `EDX` equal to 0 or x^2 .

During execution, C may modify `EAX`, `EDX`, and the sign flag `SF`. We denote by `MAYCHANGE` : $\wp(\mathbb{L}) \rightarrow \wp(\Sigma \times \Sigma)$ the resources that the program may modify. Formally, `MAYCHANGE`(L) $\stackrel{\text{def}}{=} \{(s, s') \mid \forall l \in \mathbb{L}. l \notin L \implies s(l) = s'(l)\}$. Thus, the frame condition can be written as `MAYCHANGE`($\{\text{instr}, \text{EAX}, \text{EDX}, \text{SF}\}$). The correctness of C is captured by:

$$\text{ensures}^{\tau_{\text{x86}}} \left(\begin{array}{l} \{s \mid \text{ALIGN}^{\tau_{\text{x86}}}(s, i_0, C) \wedge s(\text{AX}) = x \wedge s(\text{EDX}) = 0\}, \\ \{s \mid \text{END}^{\tau_{\text{x86}}}(s, i_0 + 12) \wedge (s(\text{EDX}) = x^2 \vee s(\text{EDX}) = 0)\}, \\ \text{MAYCHANGE}(\{\text{instr}, \text{EAX}, \text{EDX}, \text{SF}\}) \end{array} \right)$$

Recall that the frame rule in separation logic [39, 43] states that if $\{P\} C \{Q\}$ holds, then for a disjoint memory region R , $\{P * R\} C \{Q * R\}$ also holds. Similarly, in our logic, if R is invariant under `MAYCHANGE`(L), i.e., $\forall s, s'. (s, s') \in \text{MAYCHANGE}(L) \implies (s \in R \iff s' \in R)$ then `ensures`($P, Q, \text{MAYCHANGE}(L)$)

implies $\text{ensures}(P \cap R, Q \cap R, \text{MAYCHANGE}(L))$. Therefore, \mathcal{L}_1 supports modular verification while preserving the simplicity of first-order predicates, enabling efficient proof automation.

The logic \mathcal{L}_1 is equipped with the usual derivation rules for reasoning about the program execution, cf. Appendix A. The logic core and tactics are implemented in 10k lines of HOL Light [18]. It is currently used to verify functional safety properties of the s2n-bignum library, comprising 615 arithmetic routines written in ARM and x86 assembly languages for P-256/384/521, x25519/ed25519 and RSA. A total of 1013 functional properties have been verified, amounting to 860k lines of proofs.

5 Program Logic \mathcal{L}_2 for Relational Verification

In this section, we first introduce a stronger variant of the **eventually** predicate. Then, we present the relational logic \mathcal{L}_2 as a natural extension of \mathcal{L}_1 . We show how to prove a unary Hoare triple from a relational one and vice versa. This last step is essential in demonstrating the robustness of our logic and allows proofs to transition between \mathcal{L}_1 and \mathcal{L}_2 . We highlight the main extensions that allow us to prove relational properties and leave a discussion about the details of the challenges in Appendix B.

5.1 Unary Hoare Triples with Number of Steps

Building on [42], we propose a stronger **eventually** operator that explicitly specifies the number of steps required to reach a given postcondition.

Definition 3 (Stronger Eventually). *Given an operational semantics $\tau \subseteq \Sigma \times \Sigma$ and a number of steps $n \in \mathbb{N}$, for any postcondition $Q \subseteq \Sigma$, we define:*

$$\text{eventually}_n^\tau(Q) \stackrel{\text{def}}{=} \left\{ s \in \Sigma \mid \begin{array}{l} \forall s'. s \xrightarrow{\tau^n} s' \implies s' \in Q \wedge \\ \forall s', l \in \mathbb{N}. l < n \wedge s \xrightarrow{\tau^l} s' \implies \exists s''. s' \xrightarrow{\tau} s'' \end{array} \right\}$$

That is, it defines the set of states such that for all states reachable in n steps, the postcondition Q must hold, and for all states reachable in less than n steps, there must exist a successor state.

There are two merits in specifying the number of steps n . First, it makes the conjunction rule sound. In low-level languages, a program execution that failed to satisfy the postcondition at **instr** may continue as long as it encounters decodable instructions, and then branch back prior to **instr** and eventually satisfy the postcondition. Therefore, writing multiple postconditions at **instr** that hold at different steps but not together would break the conjunction rule as shown in Section 2. Explicitly stating the exact number of steps to arrive at the postcondition as an additional constraint resolves such problem. Second, it retains soundness of the commutativity and composition of nested **eventually_n** operators, which are similarly important for proving natural properties of relational

Hoare triples. When a low-level program exhibits nondeterministic behavior, each trace may meet the postcondition after different numbers of steps.⁷ The definition of **eventually_n** is stronger than **eventually**, cf. Definition 1.

Lemma 1. $\forall Q \subseteq \Sigma, n \in \mathbb{N}. \text{eventually}_{\mathbf{n}}(Q) \subseteq \text{eventually}(Q)$

The stronger eventually operator supports the following properties:

Conjunction As we require postconditions to hold after exactly n steps, we can unify contracts stating different postconditions on the final states.

$$\frac{s \in \text{eventually}_{\mathbf{n}}(Q) \quad s \in \text{eventually}_{\mathbf{n}}(Q')}{s \in \text{eventually}_{\mathbf{n}}(Q \cap Q')} \text{CONJ}$$

Commutativity Nested eventually operators commute, implying that the order of the two programs specified by the relational property will not matter. Whenever $Q^\times \subseteq \Sigma \times \Sigma$ is eventually satisfied in n_0 and n_1 steps, for the first and second components of Q^\times , the inverse $\{(s_1, s_0) \mid (s_0, s_1) \in Q^\times\}$ is satisfied in n_1 and n_0 steps, respectively.

$$\frac{s_0 \in \text{eventually}_{\mathbf{n}_{n_0}}\left(\left\{s'_0 \mid s_1 \in \text{eventually}_{\mathbf{n}_{n_1}}\left(\dot{Q}_{\pi_0=s'_0}^\times\right)\right\}\right)}{s_1 \in \text{eventually}_{\mathbf{n}_{n_1}}\left(\left\{s'_1 \mid s_0 \in \text{eventually}_{\mathbf{n}_{n_0}}\left(\dot{Q}_{\pi_1=s'_1}^\times\right)\right\}\right)} \text{COMM}$$

Here, $\dot{Q}_{\pi_i=s_x}^\times \subseteq \Sigma$ contains all states that satisfy Q together with s_x in the i -th component, i.e., $\dot{Q}_{\pi_i=s_x}^\times \stackrel{\text{def}}{=} \{\pi_{1-i}(s, s') \mid \pi_i(s, s') = s_x \wedge (s, s') \in Q^\times\}$. The projection π_i retrieves the i -th component of a pair of states (zero indexed).

Projections are lifted to sets of states by $\pi_i(Q^\times) \stackrel{\text{def}}{=} \{\pi_i(s, s') \mid (s, s') \in Q^\times\}$.

Composition Two fragments reaching Q^\times and R^\times in n_0, n_1 and m_0, m_1 steps, respectively, can be composed to reach R^\times in $n_0 + m_0$ and $n_1 + m_1$ steps.

$$\frac{\forall s'_0, s'_1. (s'_0, s'_1) \in Q \implies s'_1 \in \text{eventually}_{\mathbf{n}_{m_0}}\left(\left\{s \mid s'_0 \in \text{eventually}_{\mathbf{n}_{m_1}}\left(\dot{R}_{\pi_0=s}^\times\right)\right\}\right)}{s_0 \in \text{eventually}_{\mathbf{n}_{n_0+m_0}}\left(\left\{s \mid s_1 \in \text{eventually}_{\mathbf{n}_{n_1+m_1}}\left(\dot{R}_{\pi_0=s}^\times\right)\right\}\right)} \text{COMP}$$

With the three properties of **eventually_n** (cf. CONJ, COMM, and COMP), we can define a unary Hoare triple that maintains the properties that users would naturally expect from a Hoare logic. To do so, we employ a step function $fn : \Sigma \rightarrow \mathbb{N}$ to make the number of steps dependent on a given state.

Definition 4 (Stronger Ensures). *Given an operational semantics $\tau \subseteq \Sigma \times \Sigma$, a precondition $P \subseteq \Sigma$, a postcondition $Q \subseteq \Sigma$, a frame condition $F \subseteq \Sigma \times \Sigma$, and a step function $fn : \Sigma \rightarrow \mathbb{N}$, a unary Hoare triple is a statement of the form **ensures_{fn}**(P, Q, F), where:*

$$\text{ensures}_{\mathbf{fn}}(P, Q, F) \stackrel{\text{def}}{\iff} \frac{\forall s. s \in P \implies s \in \text{eventually}_{\mathbf{fn}(s)}^\tau(\{s' \mid s' \in Q \wedge (s, s') \in F\})}{}$$

⁷ <https://github.com/aws-labs/s2n-bignum/blob/c747b1b66801e3975a8da502e18962838d3be945/common/relational2.ml#L86-L243>

Whenever precondition P holds for state s , postcondition Q holds for any state s' that is related by $fn(s)$ steps of the execution of the program $\mathbf{prog}(P)$, and $\mathbf{prog}(P)$ modifies only the memory locations specified by the frame condition F .

As a consequence of Lemma 1, the unary Hoare triple $\mathbf{ensures}_{fn}(P, Q, F)$ is stronger than $\mathbf{ensures}$, cf. Definition 2.

Theorem 1. $\forall P, Q, F, fn. \mathbf{ensures}_{fn}(P, Q, F) \implies \mathbf{ensures}(P, Q, F)$

The other direction of the implication is not always true; in fact, it holds only for deterministic programs. The reason is that the program may branch based on a nondeterministic choice, and the postcondition may hold in a different number of steps than the one specified in the Hoare triple.

Theorem 2. *For any operational semantics τ , precondition P , postcondition Q , frame condition F , if τ is deterministic, then:*

$$\mathbf{ensures}^\tau(P, Q, F) \implies \exists fn. \mathbf{ensures}_{fn}^\tau(P, Q, F)$$

5.2 Relational Hoare Triples

We now define the relational Hoare triple $\mathbf{ensures}_{fn_0, fn_1}^2(P^\times, Q^\times, F^\times)$ that allows us to reason about the behavior of two programs. Whenever the precondition $P^\times \subseteq \Sigma \times \Sigma$ holds for a pair of states (s_0, s_1) , the postcondition $Q^\times \subseteq \Sigma \times \Sigma$ should eventually hold for any pair of states (s'_0, s'_1) that are related by respectively fn_0 and fn_1 steps of the execution of the two programs C_0 and C_1 . As for the logic \mathcal{L}_1 , the two programs are not explicitly given but instead are constrained in the memory space by P , i.e., $C_0 = \mathbf{prog}(\pi_0(P))$ and $C_1 = \mathbf{prog}(\pi_1(P))$. The frame condition $F^\times \subseteq (\Sigma \times \Sigma) \times (\Sigma \times \Sigma)$ specifies the memory locations that can be modified by the two programs. Formally:

Definition 5 (Relational Ensures). *Given operational semantics $\tau \subseteq \Sigma \times \Sigma$, a precondition $P^\times \subseteq \Sigma \times \Sigma$, a postcondition $Q^\times \subseteq \Sigma \times \Sigma$, and a frame condition $F^\times \subseteq (\Sigma \times \Sigma) \times (\Sigma \times \Sigma)$, two step functions $fn_0, fn_1 : \Sigma \rightarrow \mathbb{N}$, a relational Hoare triple is a statement of the form $\mathbf{ensures}_{fn_0, fn_1}^2(P^\times, Q^\times, F^\times)$, where:*

$$\begin{aligned} \mathbf{ensures}_{fn_0, fn_1}^2(P^\times, Q^\times, F^\times) &\stackrel{\text{def}}{\iff} \forall s_0, s_1. (s_0, s_1) \in P^\times \implies s_0 \in M_{s_0, s_1} \\ \text{where } M_{s_0, s_1} &\stackrel{\text{def}}{=} \mathbf{eventually}_{fn_0(s_0)}(\{s'_0 \mid s_1 \in N_{s_0, s_1, s'_0}\}) \\ \text{and } N_{s_0, s_1, s'_0} &\stackrel{\text{def}}{=} \mathbf{eventually}_{fn_1(s_1)}(\{s'_1 \mid (s'_0, s'_1) \in Q^\times \wedge ((s_0, s_1), (s'_0, s'_1)) \in F^\times\}) \end{aligned}$$

The definitions of M_{s_0, s_1} and N_{s_0, s_1, s'_0} nest $\mathbf{eventually}_{fn}$ requirements: M_{s_0, s_1} includes all the states where the program C_0 reaches a state s'_0 within $fn_0(s_0)$ steps, and N_{s_0, s_1, s'_0} includes all the states where the program C_1 reaches a state s'_1 where $(s'_0, s'_1) \in Q^\times$ and $((s_0, s_1), (s'_0, s'_1)) \in F^\times$ hold within $fn_1(s_1)$ steps.

As this definition is based on nested $\mathbf{eventually}_{fn}$ operators, thanks to its properties CONJ, COMM, and COMP, it follows that the relational Hoare triple $\mathbf{ensures}_{fn_0, fn_1}^2$ commutes, is compositional, and allows contract unification.

Lemma 2 (Commutativity). *Given precondition P^\times , postcondition Q^\times , frame condition F^\times , and step functions fn_0, fn_1 , the relational Hoare triple commutes:*

$$\mathbf{ensures}^2_{fn_0, fn_1}(P^\times, Q^\times, F^\times) \iff \mathbf{ensures}^2_{fn_1, fn_0}(P^S, Q^S, F^S)$$

where the swapped versions are defined as $X^S \stackrel{\text{def}}{=} \{(s_1, s_0) \mid (s_0, s_1) \in X\}$.

This symmetry above ensures that the relational logic is invariant to the program orders, allowing their roles to be interchanged without affecting the triple's validity.

Lemma 3 (Compositional). *Given three properties $P^\times, R^\times, Q^\times$, two frame conditions F_0^\times, F_1^\times , and four step numbers n_0, n_1, m_0, m_1 , it holds that two relational Hoare triples can be composed transitively:*

$$\begin{aligned} & \mathbf{ensures}^2_{\lambda s.n_0, \lambda s.m_0}(P^\times, R^\times, F_0^\times) \wedge \mathbf{ensures}^2_{\lambda s.n_1, \lambda s.m_1}(R^\times, Q^\times, F_1^\times) \\ \implies & \mathbf{ensures}^2_{\lambda s.n_0+n_1, \lambda s.m_0+m_1}(P^\times, Q^\times, F_0^\times \circ F_1^\times) \end{aligned}$$

Similarly, also the frame condition can be transitively composed. This is essential in Section 7 for the composition of program equivalences.

Lemma 4 (Compositional of Frame Conditions). *Given two preconditions P, P' , two postconditions Q, Q' , and three frame conditions F_0, F_1, F_2 , and three step functions fn_0, fn_1, fn_2 , it holds that two relational Hoare triples can be composed transitively with respect to the frame conditions:*

$$\frac{\mathbf{ensures}^2_{fn_0, fn_1}(P, Q, \{(s_0, s_1), (s'_0, s'_1) \mid (s_0, s'_0) \in F_0 \wedge (s_1, s'_1) \in F_1\}) \quad \mathbf{ensures}^2_{fn_1, fn_2}(P', Q', \{(s_0, s_1), (s'_0, s'_1) \mid (s_0, s'_0) \in F_1 \wedge (s_1, s'_1) \in F_2\})}{\mathbf{ensures}^2_{fn_0, fn_2}(P \circ P', Q \circ Q', \{(s_0, s_1), (s'_0, s'_1) \mid (s_0, s'_0) \in F_0 \wedge (s_1, s'_1) \in F_2\})}$$

Lemma 4 formalizes equivalence transitivity: when a program C_0 is equivalent to C_1 and C_1 is equivalent to C_2 , then C_0 is equivalent to C_2 . This Lemma is vital in the equivalence proofs because proving the correctness of each optimization step independently is easier than directly proving the equivalence of the original and optimized program.

Lemma 5 (Conjunction). *Given two preconditions P_0^\times, P_1^\times , two postconditions Q_0^\times, Q_1^\times , and a frame condition F^\times , two contracts can be unified with a conjunction:*

$$\begin{aligned} & \mathbf{ensures}^2_{fn_0, fn_1}(P_0^\times, Q_0^\times, F^\times) \wedge \mathbf{ensures}^2_{fn_0, fn_1}(P_1^\times, Q_1^\times, F^\times) \\ \implies & \mathbf{ensures}^2_{fn_0, fn_1}(P_0^\times \cap P_1^\times, Q_0^\times \cap Q_1^\times, F^\times) \end{aligned}$$

All these properties of our Hoare triples enable us to reason about the behavior of two programs, while maintaining the natural properties of a Hoare logic. Appendix C presents the additional properties of our program logic \mathcal{L}_2 , including the weakening and strengthening of pre-, post-, and frame conditions. Implemented in HOL Light, the core of the relational verification amounts to 1704 lines of code.

5.3 Connection with Unary Hoare Triples

We compare the relational Hoare triple **ensures2** with the unary counterpart **ensuresn**, demonstrating two key transformations: (1) deriving relational Hoare triples from two unary ones, and (2) extracting a unary Hoare triple from a *hybrid* relational one. These transformations serve a dual purpose. First, deriving a relational triple from unary ones enables reasoning about the behavior of two programs by analyzing each independently:

Theorem 3. *Given two sets of pre-, post-, and frame conditions P, P', Q, Q', F, F' , and two step functions fn_0, fn_1 , it holds that:*

$$\begin{aligned} & \mathbf{ensuresn}_{fn_0}(P, Q, F) \wedge \mathbf{ensuresn}_{fn_1}(P', Q', F') \\ & \implies \mathbf{ensures2}_{fn_0, fn_1}(P \times P', Q \times Q', F \times F') \end{aligned}$$

Second, extracting a unary triple from a hybrid relational one allows results obtained in the unary logic to be seamlessly promoted to the relational framework. A hybrid relational triple is a relational triple where the pre-, post-, and frame conditions relate to unary pre-, post-, and frame conditions, respectively. The goal is to be able to extract a unary Hoare triple from a relational one; hence: (i) the relational precondition should always have a satisfying pair (s_0, s_1) when s_1 satisfies the unary precondition; (ii) if a pair (s_0, s_1) satisfies the relational postcondition, then s_1 should satisfy the unary postcondition; and (iii) the frame condition should be satisfied for the product relation whenever the second component satisfies the frame condition of the unary relation.

Definition 6 (Hybrid Relational Ensures). *Given the pre-, post-, and frame conditions for the product relation $P^\times, Q^\times, F^\times$, and unary pre-, post-, and frame conditions P, Q, F , and two step functions $fn_0, fn_1 : \Sigma \rightarrow \mathbb{N}$, a hybrid relational Hoare triple, written $\mathbf{hensures2}_{fn_0, fn_1}(P^\times, Q^\times, F^\times \mid P, Q, F)$, holds if:*

$$\begin{aligned} & \mathbf{ensures2}_{fn_0, fn_1}(P^\times, Q^\times, F^\times) \\ & \wedge \forall s_1. s_1 \in P \implies \exists s_0. (s_0, s_1) \in P^\times & (i) \\ & \wedge \forall s_0, s_1. (s_0, s_1) \in Q^\times \implies s_1 \in Q & (ii) \\ & \wedge \exists F'. \forall s_0, s_1, s'_0, s'_1. \left(\begin{array}{l} ((s'_0, s'_1), (s_0, s_1)) \in F^\times \iff \\ (s'_0, s'_1) \in F' \wedge (s_0, s_1) \in F \end{array} \right) & (iii) \end{aligned}$$

Employing the hybrid relational triple **hensures2** (with the prefix *h* denoting “hybrid”) simplifies the verification process and makes the logic more robust. For instance, it enables translating correctness proofs for one program to another, equivalent program without having to reprove them, saving time and effort. The next result shows that a hybrid relational Hoare triple can be transformed into a unary Hoare triple.

Theorem 4. $\mathbf{hensures2}_{fn_0, fn_1}(P^\times, Q^\times, F^\times \mid P, Q, F) \implies \mathbf{ensuresn}_{fn_1}(P, Q, F)$.

6 Constant-Time Behavior

In this section, we show how our relational logic \mathcal{L}_2 can be applied to reason about constant-time behavior. As is customary in security analysis, we discriminate between public and private input data by partitioning the state labels into two disjoint sets, i.e., $\mathbb{L} = \mathbb{L}_{\text{pub}} \cup \mathbb{L}_{\text{pri}}$ and $\mathbb{L}_{\text{pub}} \cap \mathbb{L}_{\text{pri}} = \emptyset$. Public and private data induce equivalence relations on states, i.e., \simeq_{pub} and \simeq_{pri} respectively. The public data is accessible to the attacker, while the private data is kept secret. A program is constant-time if, for the same public input data, any two executions terminate with the same number of clock cycles. As a result, private data does not influence the execution time of the program.

Constant-Time via Events Accumulation. It is not practical to specify constant-time behavior by ensuring that the number of steps is equivalent in executions with the same public data, as it is highly dependent on the underlying hardware. Due to microarchitectural effects, such as memory access patterns or branch prediction, the number of clock cycles can vary significantly between executions. Instead, we can safely reason about constant-time behavior by employing a stronger notion of timing security: a program is *constant-time* if—for the same public input data—any two executions of the program induce *the same trace of microarchitectural events*.

To observe these events, we extend the state space Σ with an **events** component in $\mathbb{L}_e \stackrel{\text{def}}{=} \mathbb{L} \cup \{\mathbf{events}\}$. The **events** component records an ordered list of events, such as memory accesses (**load** x, n or **store** x, n ; where x is the accessed address and n the operation size in bytes), and branch jumps (**branch** x, y ; where x and y are the current and the destination program counter, respectively). Any other variable-time instruction, such as division or floating point operations can also be included in the event trace. In our case studies, these operations are intentionally left unresolved by the operational semantics, and thus not included in the event trace. These events are public data, i.e., $\mathbf{events} \in \mathbb{L}_{\text{pub}}$. The extended state space is Σ_e with operational semantics τ^e . For instance, loading a memory address x into a 16 bit register r collects a load event of 2 bytes:

$$\frac{\begin{array}{l} s(\mathbf{instr}) = i \quad \text{DECODE}(s, i) = \mathbf{load } r, x \\ s(\mathbf{memory}_x) = v \quad s(\mathbf{events}) = e \quad \text{LENGTH}(r) = 16 \end{array}}{s \xrightarrow{\tau^e} s[r \mapsto v, \mathbf{events} \mapsto (e \ ++ \ \mathbf{load } x, 2), \mathbf{instr} \mapsto i + \text{LENGTH}^r(\mathbf{load } r, x)]} \text{LOAD}$$

Therefore, we are now able to specify constant-time behavior by ensuring that the list of microarchitectural events is the same in both executions. Our approach can be easily extended to include other side-channels, such as power consumption.

While we do not include opcode-level information in our events, instruction opcodes can influence the number of cycles (e.g., the **cbz** and **b.ne** instructions in ARM). This relies on an assumption that a program is public information, and therefore the events do not need to carry opcode information. This assumption can be broken if a program runs assembly instructions that are separately stored in a private input buffer. We prove that such things do not happen individually.

Definition 7 (Constant-Time via Event Accumulation). Let $\tau^e \subseteq \Sigma_e \times \Sigma_e$ be an operational semantics that collects the microarchitectural events, $P \subseteq \Sigma_e$ be a precondition, $Q \subseteq \Sigma_e$ a postcondition, $F \subseteq \Sigma_e \times \Sigma_e$ a frame condition, and $fn_0, fn_1 : \Sigma_e \rightarrow \mathbb{N}$ two step functions. The program $\text{prog}(P)$ is constant-time with respect to private data \mathbb{L}_{pri} if it holds that:

$$\text{ensures}_{fn_0, fn_1}^{\tau^e} \left(\begin{array}{l} \{(s_0, s_1) \in P \times P \mid s_0(\mathbb{L}_{\text{pub}}) = s_1(\mathbb{L}_{\text{pub}})\}, \\ \{(s_0, s_1) \in Q \times Q \mid s_0(\text{events}) = s_1(\text{events})\}, \\ F \times F \end{array} \right)$$

Note that, by constraining the public data to be equal in the precondition, we also require that states share the same event trace before executing the program.

Example 3. The program `cst-compare` in Figure 1 (right) is constant-time with respect to the microarchitectural events of Definition 7. Indeed, `cst-compare` first branches on the length n of the buffers if $n = 0$ at Line 2; otherwise, it compares the buffers byte-by-byte. Assuming registers of 32 bits, each iteration collects two 4-bytes load events: one for each buffer at Lines 5 and 6. Then, it branches to start the next iteration at Line 9 until the end of the buffers, no matter what the comparison result is. Hence, for any public input value, `cst-compare` induces the same event trace.

In contrast, the program `compare` in Figure 1 (left) is not constant-time since the event trace may be different for two executions. Consider the following counterexample, where $n = 1$ and the buffers are $k = 10$ and $x = 20$. In memory, the two executions contain $s_0(\text{memory}_{10}) = 0$ and $s_0(\text{memory}_{20}) = 0$; and $s_1(\text{memory}_{10}) = 0$ and $s_1(\text{memory}_{20}) = 1$ respectively. The two traces differ at the first mismatch, as the loop in the second execution is terminated early. For brevity, the following event traces are simplified omitting the address of branch instructions with the evaluation of the condition:

$$\begin{aligned} s_0(\text{events}) &= [\text{branch FALSE}, \text{load } 10, 4, \text{load } 20, 4, \text{branch FALSE}, \text{branch FALSE}] \\ s_1(\text{events}) &= [\text{branch FALSE}, \text{load } 10, 4, \text{load } 20, 4, \text{branch TRUE}] \end{aligned}$$

Constant-Time via Unary to Relational Embedding. We can employ unary Hoare logic to prove constant-time behavior by showing that private data does not influence the event trace generated during program execution. In other words, it is sufficient to provide a witness trace that depends only on public data.

Definition 8 (Constant-Time via Unary to Relational Embedding). Let τ^e be an operational semantics that collects the microarchitectural events, P be a precondition, Q a postcondition, and F a frame condition. The program $\text{prog}(P)$ is constant-time with respect to private data \mathbb{L}_{pri} if there exists a function $f : \Sigma_e(\mathbb{L}_{\text{pub}}) \rightarrow \mathbb{E}$ such that:

$$\forall v_{\text{pub}}, e_0. \text{ensures}_{fn}^{\tau^e} \left(\begin{array}{l} \{s \in P \mid v_{\text{pub}} = s(\mathbb{L}_{\text{pub}}) \wedge e_0 = s(\text{events})\}, \\ \{s \in Q \mid s(\text{events}) = e_0 \text{ ++ } f(v_{\text{pub}})\}, \\ F \end{array} \right)$$

where $\Sigma_e(\mathbb{L}_{\text{pub}})$ is the partial projection of states Σ_e on public data \mathbb{L}_{pub} , and $\#$ is the list concatenation.

This approach eliminates the need to run the symbolic simulation tactic twice, but requires providing an explicit witness for the event trace function f . Since this approach proves a statement about a single program execution, the proof structure is very similar to the correctness proof. Therefore, we can merge the two proofs for correctness and constant-time behavior into a single one; thus eliminating the computational effort of checking each proof separately and greatly reducing the overhead of writing and maintaining them. We can retrieve the relational definition by instantiating Theorem 3 with two instances of the same `ensuresn` proof, renamed accordingly.

Example 4. Using list comprehension, for a given public input v_{pub} , the witness f for the program `cst-compare` in Figure 1 is defined as:

$$[\text{branch } (v_{\text{pub}}(\mathbf{n}) = 0)] \# \left[\begin{array}{l} \text{load } (v_{\text{pub}}(\mathbf{x}) + v_{\text{pub}}(\mathbf{n}) - 1 - i), 4 \\ \text{load } (v_{\text{pub}}(\mathbf{y}) + v_{\text{pub}}(\mathbf{n}) - 1 - i), 4 \\ \text{branch } (i < v_{\text{pub}}(\mathbf{n})) \end{array} \middle| i \in [0, v_{\text{pub}}(\mathbf{n})] \right]$$

where \mathbf{n} , \mathbf{x} , and \mathbf{y} are public data and therefore accessible in v_{pub} .

Note that, routines in the `s2n-bignum` library can be proven constant-time by instantiating either Definition 7 or Definition 8, the two are equivalent.

7 Equivalence Checking

In this section, we demonstrate the application of our relational Hoare logic framework to equivalence checking between performance and verification-friendly implementations of the same routine in the `s2n-bignum` library.

Equivalence between Two Programs. Two programs are considered functionally equivalent if they produce the same output states starting from equivalent input states. When dealing with assembly-level programs, we must carefully define what it means for two states to be “equal”. For instance, two equal input states should not require the exact same code in memory; otherwise, only identical programs could be compared. Similarly, because the calling convention allows callee-save registers to hold different values, the value of these registers should not be constrained.

On the output side, certain registers or memory regions may differ if they are not designated as outputs. For example, eliminating dead stores to the stack frame is a valid optimization because the stack frame is not used after function returned. Two equivalent output states must allow those parts of memory to contain different data.

As a consequence, the equivalence checking takes as a parameter the equivalence relations $\simeq_{\text{in}} \subseteq \Sigma \times \Sigma$ and $\simeq_{\text{out}} \subseteq \Sigma \times \Sigma$ that define when input and output states are considered equivalent. This relation has to be defined manually for each pair of programs to be compared.

Example 5. Consider the two programs `compare` and `cst-compare` in Figure 1. Assuming a proof of correctness for `compare` already exists, our goal is to prove that the secure constant-time version is functionally equivalent to the original program, without needing to reprove the correctness of `cst-compare` from scratch. To do so, we define the input equivalence \simeq_{in} , relating the program counter, input registers, and relevant part of the memory as follows:

$$\simeq_{\text{in}} = \text{MAYCHANGE} \left(\mathbb{L} \setminus \left(\{\text{instr}, \mathbf{n}, \mathbf{x}, \mathbf{y}\} \cup \{\text{memory}_i \mid i \in [x, x+n) \vee i \in [y, y+n)\} \right) \right)$$

Note the use of the `MAYCHANGE` operator to define a relation that allows two states to differ in all labels but the ones specified. For output equivalence \simeq_{out} , we relate only the output register, i.e., $\simeq_{\text{out}} = \text{MAYCHANGE}(\mathbb{L} \setminus \{\text{res}\})$.

Definition 9 (Equivalence). *Let $P_0, P_1 \subseteq \Sigma$ be two preconditions, $Q_0, Q_1 \subseteq \Sigma$ two postconditions, $F_0, F_1 \subseteq \Sigma \times \Sigma$ two frame conditions, and $fn_0, fn_1 : \Sigma \rightarrow \mathbb{N}$ two step functions. Given the input and output equivalences $\simeq_{\text{in}}, \simeq_{\text{out}} \subseteq \Sigma \times \Sigma$, the programs `prog`(P_0) and `prog`(P_1) are equivalent if it holds that:*

$$\text{ensures}^2_{fn_0, fn_1} \left(\begin{array}{l} \{(s_0, s_1) \in P_0 \times P_1 \mid s_0 \simeq_{\text{in}} s_1\}, \\ \{(s_0, s_1) \in Q_0 \times Q_1 \mid s_0 \simeq_{\text{out}} s_1\}, \\ F_0 \times F_1 \end{array} \right)$$

Example 6. We can prove that the constant-time program `cst-compare` is equivalent to the original program in `compare` by applying Definition 9 with the input and output equivalences defined in Example 5. Along the lines of the pre- and postconditions defined in Section 6, we define:

$$\begin{aligned} P' &= \left\{ s \mid \begin{array}{l} s(\mathbf{n}) = n \wedge s(\mathbf{x}) = x \wedge s(\mathbf{y}) = y \wedge \\ \forall i \leq n. s(\text{memory}_{x+i}) = \mathbf{x}_i \wedge s(\text{memory}_{y+i}) = \mathbf{y}_i \end{array} \right\}, \\ P_0 &= \{s \in P' \mid \text{ALIGN}^\tau(s, i_0, \text{cst-compare})\}, P_1 = \{s \in P' \mid \text{ALIGN}^\tau(s, i_0, \text{compare})\}, \\ Q_0 &= \{s \mid \text{END}^\tau(s, \text{LENGTH}^\tau(\text{cst-compare}))\}, Q_1 = \{s \mid \text{END}^\tau(s, \text{LENGTH}^\tau(\text{compare}))\}, \\ F_0 &= \text{MAYCHANGE}(\{\text{instr}, \mathbf{n}, \mathbf{xn}, \mathbf{yn}\}), \\ F_1 &= \text{MAYCHANGE}(\{\text{instr}, \mathbf{n}, \mathbf{xn}, \mathbf{yn}, \text{diff}, \text{temp}\}), \\ fn_0(s) &= \text{LARGESTPREFIX}_n(s, x, y), \text{ and } fn_1(s) = s(\mathbf{n}), \end{aligned}$$

where $\text{LARGESTPREFIX}_n(s, x, y)$ is the length of the largest prefix among the two given memory addresses x and y of length n . In conclusion, Definition 9 provides the specification for the equivalence proof between the two programs.

Composition of Program Equivalences. We slightly abuse notation and define `eqensures` as a shorthand for the equivalence of two programs C_0, C_1 with \simeq_{in} in the precondition starting from pc_0, pc_1 , eventually reaching \simeq_{out} in the postcondition at $\text{pc}'_0, \text{pc}'_1$: `eqensures` $_{\text{pc}_0, \text{pc}'_0, \text{pc}_1, \text{pc}'_1}(C_0, C_1, \simeq_{\text{in}}, \simeq_{\text{out}})$. Notably, Lemma 3 proves that the sequential composition of two equivalence proofs is sound if

$\forall s, s'. s \simeq_{\text{out}} s' \implies s \simeq_{\text{in}} s'$. Formally, the *sequential composition* of two equivalences is defined as follows:

$$\frac{\text{eqensures}_{\text{pc}_0, \text{pc}'_0, \text{pc}_1, \text{pc}'_1}(C_0, C_1, \simeq_{\text{in}}, \simeq_{\text{out}}) \quad \text{eqensures}_{\text{pc}'_0, \text{pc}''_0, \text{pc}'_1, \text{pc}''_1}(C_0, C_1, \simeq'_{\text{in}}, \simeq'_{\text{out}})}{\text{eqensures}_{\text{pc}_0, \text{pc}'_0, \text{pc}_1, \text{pc}'_1}(C_0, C_1, \simeq_{\text{in}}, \simeq_{\text{out}})}$$

Lemma 4 instead proves the soundness of the transitive composition of two equivalences, only if the result input and output equivalences preserve the existence of an intermediate state, i.e., $s \simeq_{\text{in}} s' \iff \exists s''. (s \simeq_{\text{in}} s'' \wedge s'' \simeq_{\text{in}} s')$ and $s \simeq_{\text{out}} s' \iff \exists s''. (s \simeq_{\text{out}} s'' \wedge s'' \simeq_{\text{out}} s')$. Formally, the *transitive composition* of two equivalences is defined as follows:

$$\frac{\text{eqensures}_{\text{pc}_0, \text{pc}'_0, \text{pc}_1, \text{pc}'_1}(C_0, C_1, \simeq_{\text{in}}, \simeq_{\text{out}}) \quad \text{eqensures}_{\text{pc}_1, \text{pc}'_1, \text{pc}_2, \text{pc}'_2}(C_1, C_2, \simeq_{\text{in}}, \simeq_{\text{out}})}{\text{eqensures}_{\text{pc}_0, \text{pc}'_0, \text{pc}_2, \text{pc}'_2}(C_0, C_2, \simeq'_{\text{in}}, \simeq'_{\text{out}})}$$

Combining Equivalence and Correctness Proofs. In the following, we show how to reuse a correctness proof of an original program to obtain a correctness proof of an optimized program through program equivalence. Indeed, given the functional correctness of the original program in the form of an **ensures_n** proof, we can apply it to the optimized program by proving the equivalence of the two via the relational Hoare triple **ensures₂**. The correctness proof of the optimized program is given in the form of a hybrid relational Hoare triple **hensures₂**, presented in Definition 6.

Theorem 5 (Transfer of Correctness through Equality).

$$\begin{aligned} & \text{ensures}_{fn_0}^{\tau}(P, Q, F) \wedge \text{ensures}_{fn_0, fn_1}^{2\tau}(P^{\times}, Q^{\times}, F^{\times}) \\ \implies & \text{hensures}_{fn_0, fn_1}^{2\tau} \left(\begin{array}{l} \{(s_0, s_1) \in P^{\times} \mid s_0 \in P\}, \\ \{(s_0, s_1) \in Q^{\times} \mid s_0 \in Q\}, \\ \{(s_0, s_1), (s'_0, s'_1) \in F^{\times} \mid (s_0, s'_0) \in F\} \end{array} \middle| \begin{array}{l} P, \\ Q, \\ F \end{array} \right) \end{aligned}$$

Let $P_0 \subseteq \Sigma$ be the precondition, $Q_0 \subseteq \Sigma$ the postcondition, $F_0 \subseteq \Sigma \times \Sigma$ the frame condition, and $fn_0 : \Sigma \rightarrow \mathbb{N}$ the step function. We state functional correctness as: **ensures_n** $_{fn_0}(\{s \in P_0 \mid s(\text{pc}) = x_0\}, \{s \in Q_0 \mid s(\text{pc}) = x_{\omega}\}, F_0)$. Afterwards, from Definition 9, given the two input-output equivalences \simeq_{in} and \simeq_{out} , the equivalence between two programs is achieved by proving:

$$\text{ensures}_{fn_0, fn_1}^2 \left(\begin{array}{l} \{(s_0, s_1) \in P_0 \times P_1 \mid s_0 \simeq_{\text{in}} s_1\}, \\ \{(s_0, s_1) \in Q_0 \times Q_1 \mid s_0 \simeq_{\text{out}} s_1\}, \\ F_0 \times F_1 \end{array} \right)$$

where P_1, Q_1, F_1 are the pre-, post-, and frame conditions of the second program, respectively. Theorem 5 transfers the correctness and equivalence proofs to the following hybrid relational Hoare triple:

$$\text{hensures}_{fn_0, fn_1}^2 \left(\begin{array}{l} \{(s_0, s_1) \in P_0 \times P_1 \mid s_0 \simeq_{\text{in}} s_1 \wedge s_0(\text{pc}) = x_0\}, \\ \{(s_0, s_1) \in Q_0 \times Q_1 \mid s_0 \simeq_{\text{out}} s_1 \wedge s_0(\text{pc}) = x_{\omega}\}, \\ F_0 \times F_1 \end{array} \middle| \begin{array}{l} \{s \in P_1 \mid s(\text{pc}) = x_0\}, \\ \{s \in Q_1 \mid s(\text{pc}) = x_{\omega}\}, \\ F_1 \end{array} \right)$$

Finally, by applying Theorem 4, we obtain the correctness proof of the new program: $\text{ensures}_{fn_1}(\{s \in P_1 \mid s(\text{pc}) = x_0\}, \{s \in Q_1 \mid s(\text{pc}) = x_\omega\}, F_1)$. In Appendix D, we provide the steps required to promote a correctness proof that was originally written via the `ensures` operator—without an explicit number of steps—to a proof that uses the `ensuresn` operator. The majority of functional correctness proofs already available in the `s2n-bignum` library are written using the `ensures` operator. In total, the core of the equivalence checking proofs is 2629 lines of HOL Light code.

8 Obtaining Proofs for the `hol-bignum` Library

8.1 Case Study: Bignum Copy and Inversion Modulo Routine.

We apply the constant-time verification to the `s2n-bignum` library, notably on the copy program of large integers, cf. `bignum_copy`, and the inversion modulo a prime $p = 2^{255} - 19$, cf. `bignum_inv_p25519`. The following should provide guidance on which proof approach to apply depending on the program size and complexity.

The `bignum_copy` routine is relatively small, comprising 16 instructions that copy the content of buffer k to the buffer z , padding z with zeros if it is bigger than k . Despite its size, `bignum_copy` has the most complex program flow in the library, making it a good candidate for constant-time verification. The functional correctness proof is 180 lines. The constant-time proof, using Definition 7, is 276 lines: it does not require an explicit event trace and is fairly easy to prove correct. On the other hand, the unary constant-time proof using Definition 8 is 245 lines, and requires an explicit event trace. Although the event trace is small and intuitive, this parameter makes the proof more complex as it requires a nontrivial induction on list comprehensions. Notably, we can combine correctness and constant-time proofs together via Theorem 3 in a single, 277-line proof, which yields the lowest proof size overhead.

The `bignum_inv_p25519` routine instead is a 1033-instruction program that finds the inverse of a big integer modulo a prime $p = 2^{255} - 19$. The functional correctness proof is 2303 lines long. The constant-time proof, using the unary embedding of Definition 8 combining correctness and constant-time proofs, is 2633 lines long. Most of the additions in the proof are due to the explicit definition of the event trace, which contains 90 memory events alone. However, after defining the event trace, extending the correctness proof with the constant-time proof was effortless. All the mechanized proofs are available in the artifact.⁸ In future work, we plan to automate the generation of the event trace, which will significantly reduce the required level of manual effort.

8.2 Case Study: Elliptic Curves and Montgomery Reduction.

We utilized program equivalence to verify the functional correctness of optimized implementations for (1) *field and point operations* of NIST elliptic curves

⁸ <https://doi.org/10.5281/zenodo.15309209>

(specifically, curves P-256, P-384, and P-521), and (2) *Montgomery reduction*, an algorithm that allows efficient modular arithmetic when the modulus is large. These optimizations were achieved using an *autovectorizer*, a constraint solver-based instruction scheduler called SLOTHY [1], and the point operations of NIST curves were optimized using a custom memory instruction optimizer for the ARM architecture. We also have similar equivalence checking tactics for the x86 architecture. Overall, we checked the equivalence for 15 pairs of arithmetic routines, amounting to a total of 19k lines of proofs.

The autovectorizer replaces sequences of 64-bit scalar multiplication instructions, such as `mul` and `umulh`, with their equivalent NEON vector instructions. This optimization targets the ARM Neoverse N1 architecture, whose microarchitecture contains only one multiplication pipeline. The `mul/umulh` instructions stall this pipeline for a few cycles when executing scalar multiplication instructions. SLOTHY employs a constraint solver and cost model to find the optimal instruction scheduling, significantly reducing these stalls. Specifically, SLOTHY improves the scheduling of straight-line code in the main basic blocks of NIST curves' field operations, and also improves the software pipelining optimization in the main loop of the Montgomery reduction. The memory instruction optimizer performs two key tasks in the ARM architecture: store-to-load forwarding and dead store elimination. Store-to-load forwarding replaces load instructions with stored values, eliminating redundant memory accesses. Dead store elimination removes store instructions with results that are never used.

Tactics for Program Equivalence Proofs. To automate the writing of equivalence proofs, we developed proof tactics that can be used for two different classes of optimizations: small localized updates and instruction reordering.

For local optimizations that update only small portions of the original program, such as autovectorization, we implemented the tactic `EQUIV_STEPS_TAC`. This tactic takes as input a list of line ranges and annotations describing whether each range is optimized or left identical. For the identical portion, the tactic performs lock-step symbolic simulation and eagerly abbreviates the common outputs of the instructions with fresh variables to avoid exponential explosion of the sizes of the output expressions. For optimized ranges, the tactic employs stuttering simulation, which executes the corresponding sections of each program step-by-step. To help `EQUIV_STEPS_TAC` automatically converge on complex cases, users can register custom bit-vector equality theorems for output expressions.

For optimizations involving instruction reordering, we implemented two additional tactics: `STEPS_ABBREV_TAC` and `STEPS_REWRITE_TAC`. The first tactic performs stuttering symbolic simulation for the first program, storing the symbolic output expressions to an OCaml array. The second tactic takes as input an instruction index mapping between the two programs, along with the symbolic output generated by `STEPS_ABBREV_TAC`. Then, it simulates the second program step-by-step, proving that the symbolic output of each instruction is equal to the symbolic expression in the first program, according to the instruction mapping.

Software Pipelining of Montgomery Reduction. The Montgomery reduction is heavily used in cryptographic operations performing modular exponentiations. Its original implementation in the s2n-bignum library includes a nested loop structure, where the outer loop consists of three basic blocks: loop entry, inner loop (which consists of a single basic block), and loop exit. A faster version was achieved by: caching repetitive calculations, vectorizing `mul` and `umulh` in all basic blocks, applying software pipelining optimizations to the inner loop, and rescheduling instructions using SLOTHY.

We verified the functional correctness of the optimized Montgomery reduction by *transitively* composing equivalence proofs with the original correctness proof after each optimization stage. For each optimization, we applied *sequential* composition of equivalences between each basic block pair and induced the equivalence of the whole loop. In the case of software pipelining, which transforms the control flow graph by adding loop prologue and epilogue blocks, equivalence between inner loops was proven and composed with equivalences for the loop entry and exit blocks.

Overall, the optimized field operations of NIST curves achieved throughput speedups up to 38%, and integrating these improvements into point operations alongside memory optimizations resulted in up to 23% throughput gains. These enhancements demonstrate the substantial impact of the new optimizations on performance of the s2n-bignum library.

9 Conclusion

This work presents a novel relational Hoare logic framework for verifying realistically modelled machine code, while preserving natural properties expected from Hoare-style reasoning. Fully formalized in HOL Light, the framework is applied in two case studies involving the s2n-bignum cryptographic library, a key component of a TLS/SSL implementation. Our results show that the logic scales to large assembly programs and yields practical value in the verification of cryptographic codebases.

While Mazzucato et al. [31] have investigated constant-time verification for libraries similar to s2n-bignum, their approach relies on abstraction-dependent methods through an untrusted computing base to decompile assembly into C. In contrast, our framework operates directly on the assembly level, ensuring higher reliability of the verification results as it reduces the trusted computing base to the minimal core of the HOL Light theorem prover and to the operational semantics implementations. As future work, we plan to increase coverage of relational properties on the s2n-bignum library and improve proof automation to handle repetitive tasks. As a natural extension to constant-time proofs, we aim to address speculative execution vulnerabilities [26, 15].

Acknowledgments. We would like to thank Hanno Becker for his help in improving the Montgomery reduction implementation, and the anonymous reviewers of CAV 2025 for their valuable feedback. Research reported in this publication was supported by an Amazon Research Award, Fall 2023.

Bibliography

- [1] Abdulrahman, A., Becker, H., Kannwischer, M.J., Klein, F.: Fast and Clean: Auditable high-performance assembly via constraint solving. *Cryptology ePrint Archive*, Paper 2022/1303 (2022)
- [2] Affeldt, R.: On construction of a library of formally verified low-level arithmetic functions. In: *Proceedings of the 27th Annual ACM Symposium on Applied Computing*, pp. 1326–1331, SAC '12, Association for Computing Machinery, New York, NY, USA (Mar 2012), ISBN 978-1-4503-0857-1, <https://doi.org/10.1145/2245276.2231986>
- [3] Antonopoulos, T., Gazzillo, P., Hicks, M., Koskinen, E., Terauchi, T., Wei, S.: Decomposition instead of self-composition for proving the absence of timing channels. In: *Proceedings of the 38th ACM SIGPLAN Conference on Programming Language Design and Implementation*, pp. 362–375, PLDI 2017, Association for Computing Machinery, New York, NY, USA (Jun 2017), ISBN 978-1-4503-4988-8, <https://doi.org/10.1145/3062341.3062378>
- [4] Bartels, B., Jähnig, N.: Mechanized, Compositional Verification of Low-Level Code. In: Badger, J.M., Rozier, K.Y. (eds.) *NASA Formal Methods*, pp. 98–112, Springer International Publishing, Cham (2014), ISBN 978-3-319-06200-6, https://doi.org/10.1007/978-3-319-06200-6_8
- [5] Barthe, G., Crespo, J.M., Kunz, C.: Relational Verification Using Product Programs. In: Butler, M., Schulte, W. (eds.) *FM 2011: Formal Methods*, pp. 200–214, Springer, Berlin, Heidelberg (2011), ISBN 978-3-642-21437-0, https://doi.org/10.1007/978-3-642-21437-0_17
- [6] Barthe, G., Gregoire, B., Laporte, V.: Secure compilation of side-channel countermeasures: The case of cryptographic “constant-time”. In: *2018 IEEE 31st Computer Security Foundations Symposium (CSF)*, pp. 328–343, IEEE (2018)
- [7] Barthe, G., Rezk, T., Saabas: Proof obligations preserving compilation. In: *International Workshop on Formal Aspects in Security and Trust*, pp. 112–126, Springer (2005)
- [8] Benton: A typed, compositional logic for a stack-based abstract machine. In: *Asian Symposium on Programming Languages and Systems*, pp. 364–380, Springer (2005)
- [9] Benton, N.: Simple relational correctness proofs for static analyses and program transformations. In: *Proceedings of the 31st ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pp. 14–25, POPL '04, Association for Computing Machinery, New York, NY, USA (Jan 2004), ISBN 978-1-58113-729-3, <https://doi.org/10.1145/964001.964003>
- [10] Benton, N.: Semantic Equivalence Checking for HHVM Bytecode. In: *Proceedings of the 20th International Symposium on Principles and Practice of Declarative Programming*, pp. 1–8, PPDP '18, Association for Computing Machinery, New York, NY, USA (Sep 2018), ISBN 978-1-4503-6441-6, <https://doi.org/10.1145/3236950.3236975>

- [11] Beringer, L.: Relational Decomposition. In: van Eekelen, M., Geuvers, H., Schmaltz, J., Wiedijk, F. (eds.) *Interactive Theorem Proving*, pp. 39–54, Springer, Berlin, Heidelberg (2011), ISBN 978-3-642-22863-6, https://doi.org/10.1007/978-3-642-22863-6_6
- [12] Blatter, L., Kosmatov, N., Prevosto, V., Le Gall, P.: Certified Verification of Relational Properties. In: *Integrated Formal Methods: 17th International Conference, IFM 2022, Lugano, Switzerland, June 7–10, 2022, Proceedings*, pp. 86–105, Springer-Verlag, Berlin, Heidelberg (Jun 2022), ISBN 978-3-031-07726-5, https://doi.org/10.1007/978-3-031-07727-2_6
- [13] Bond, B., Hawblitzel, C., Kapritsos, M., Leino, K.R.M., Lorch, J.R., Parno, B., Rane, A., Setty, S., Thompson, L.: Vale: Verifying {High-Performance} Cryptographic Assembly Code. In: *26th USENIX Security Symposium (USENIX Security 17)*, pp. 917–934 (2017), ISBN 978-1-931971-40-9
- [14] Bosamiya, J., Gibson, S., Li, Y., Parno, B., Hawblitzel, C.: Verified Transformations and Hoare Logic: Beautiful Proofs for Ugly Assembly Language. In: *Software Verification: 12th International Conference, VSTTE 2020, and 13th International Workshop, NSV 2020, Los Angeles, CA, USA, July 20–21, 2020, Revised Selected Papers*, pp. 106–123, Springer-Verlag, Berlin, Heidelberg (Jul 2020), ISBN 978-3-030-63617-3, https://doi.org/10.1007/978-3-030-63618-0_7
- [15] Cauligi, S., Disselkoen, C., Moghimi, D., Barthe, G., Stefan, D.: SoK: Practical Foundations for Software Spectre Defenses. In: *2022 IEEE Symposium on Security and Privacy (SP)*, pp. 666–680, IEEE Computer Society (May 2022), ISBN 978-1-66541-316-9, <https://doi.org/10.1109/SP46214.2022.9833707>
- [16] Clarkson, M.R., Schneider, F.B.: Hyperproperties. In: *2008 21st IEEE Computer Security Foundations Symposium*, pp. 51–65 (2008), <https://doi.org/10.1109/CSF.2008.7>
- [17] Erbsen, A., Philipoom, J., Gross, J., Sloan, R., Chlipala, A.: Simple high-level code for cryptographic arithmetic - with proofs, without compromises. In: *2019 IEEE Symposium on Security and Privacy (SP)*, pp. 1202–1219 (2019), <https://doi.org/10.1109/SP.2019.00005>
- [18] Harrison, J.: HOL Light: An Overview. In: Berghofer, S., Nipkow, T., Urban, C., Wenzel, M. (eds.) *Theorem Proving in Higher Order Logics*, pp. 60–66, Springer, Berlin, Heidelberg (2009), ISBN 978-3-642-03359-9, https://doi.org/10.1007/978-3-642-03359-9_4
- [19] Harrison, J.: *HOL Light Tutorial (for Version 2.20)* (2011)
- [20] Hoare, C.: An axiomatic basis for computer programming. *Communications of the ACM* **12**(10), 576–580 (1969)
- [21] Jähnig, N., Gothel, T., Glesner: A denotational semantics for communicating unstructured code (2015)
- [22] Jähnig, N., Gothel, T., Glesner: Refinement-based verification of communicating unstructured code. In: *International Conference on Software Engineering and Formal Methods*, pp. 61–75, Springer (2016)
- [23] Jensen, J.B., Benton, N., Kennedy, A.: High-level separation logic for low-level code. In: *Proceedings of the 40th Annual ACM SIGPLAN-SIGACT*

- Symposium on Principles of Programming Languages, pp. 301–314, POPL '13, Association for Computing Machinery, New York, NY, USA (Jan 2013), ISBN 978-1-4503-1832-7, <https://doi.org/10.1145/2429069.2429105>
- [24] Kang, J., Kim, Y., Song, Y., Lee, J., Park, S., Shin, M.D., Kim, Y., Cho, S., Choi, J., Hur, C.K., Yi, K.: Crellvm: Verified credible compilation for LLVM. In: Proceedings of the 39th ACM SIGPLAN Conference on Programming Language Design and Implementation, pp. 631–645, PLDI 2018, Association for Computing Machinery, New York, NY, USA (Jun 2018), ISBN 978-1-4503-5698-5, <https://doi.org/10.1145/3192366.3192377>
- [25] Klein, G., Elphinstone, K., Heiser, G., Andronick, J., Cock, D., Derrin, P., Elkaduwe, D., Engelhardt, K., Kolanski, R., Norrish, M.: Formal verification of an os kernel. In: Proceedings of the ACM SIGOPS 22nd Symposium on Operating Systems Principles, pp. 207–220 (2009)
- [26] Kocher, P., Horn, J., Fogh, A., Genkin, D., Gruss, D., Haas, W., Hamburg, M., Lipp, M., Mangard, S., Prescher, T., Schwarz, M., Yarom, Y.: Spectre Attacks: Exploiting Speculative Execution. In: 2019 IEEE Symposium on Security and Privacy (SP), pp. 1–19 (May 2019), ISSN 2375-1207, <https://doi.org/10.1109/SP.2019.00002>
- [27] Kumar, R., Myreen, M., Norrish, M., Owens, S.: Cakeml: A verified implementation of ml. ACM SIGPLAN Notices **49**(1), 179–191 (2014)
- [28] Lehner, H., Muller: Formal translation of bytecode into BoogiePL. Electronic Notes in Theoretical Computer Science **190**(1), 35–50 (2007)
- [29] Lundberg, D., Guanciale, R., Lindner, A., Dam, M.: Hoare-Style Logic for Unstructured Programs. In: de Boer, F., Cerone, A. (eds.) Software Engineering and Formal Methods, pp. 193–213, Springer International Publishing, Cham (2020), ISBN 978-3-030-58768-0, https://doi.org/10.1007/978-3-030-58768-0_11
- [30] Marti, N.: Formal Verification of Low-Level Software. Ph.D. thesis, University of Tokyo (2008)
- [31] Mazzucato, D., Champion, M., Urban, C.: Quantitative static timing analysis. In: 31st Static Analysis Symposium (SAS 2024), Roberto Giacobazzi and Alessandra Gorla and Marco Champion, Pasadena, CA, United States (2024), https://doi.org/10.1007/978-3-031-74776-2_11
- [32] Myreen, M., Curello, G.: Formal verification of machine-code programs. In: Lundberg, D. (ed.) International Conference on Certified Programs and Proofs, p. 20, University of Cambridge, Computer Laboratory (2009)
- [33] Myreen, M., Gordon: Verification of machine code implementations of arithmetic functions for cryptography. In: Theorem Proving in Higher Order Logics: Emerging Trends Proceedings, Dept. of Computer Science, University of Kaiserslautern (2007)
- [34] Myreen, M., Gordon, M., Slind, K.: Machine-code verification for multiple architectures—an application of decompilation into logic. In: 2008 Formal Methods in Computer-Aided Design, pp. 1–8, IEEE (2008)
- [35] Myreen, M., Gordon, M., Slind, K.: Decompilation into logic—improved. In: 2012 Formal Methods in Computer-Aided Design (FMCAD), pp. 78–81, IEEE (2012)

- [36] Myreen, M.O., Fox, A.C.J., Gordon, M.J.C.: Hoare Logic for ARM Machine Code. In: Arbab, F., Sirjani, M. (eds.) International Symposium on Fundamentals of Software Engineering, pp. 272–286, Springer, Berlin, Heidelberg (2007), ISBN 978-3-540-75698-9, https://doi.org/10.1007/978-3-540-75698-9_18
- [37] Myreen, M.O., Gordon, M.J.C.: Hoare logic for realistically modelled machine code. In: Proceedings of the 13th International Conference on Tools and Algorithms for the Construction and Analysis of Systems, pp. 568–582, TACAS’07, Springer-Verlag, Berlin, Heidelberg (Mar 2007), ISBN 978-3-540-71208-4
- [38] Naumann, D.A.: Thirty-Seven Years of Relational Hoare Logic: Remarks on Its Principles and History. In: Leveraging Applications of Formal Methods, Verification and Validation: Engineering Principles: 9th International Symposium on Leveraging Applications of Formal Methods, ISOFA 2020, Rhodes, Greece, October 20–30, 2020, Proceedings, Part II, pp. 93–116, Springer-Verlag, Berlin, Heidelberg (Oct 2020), ISBN 978-3-030-61469-0, https://doi.org/10.1007/978-3-030-61470-6_7
- [39] O’Hearn, P., Reynolds, J., Yang, H.: Local reasoning about programs that alter data structures. In: Proceedings of Computer Science Logic (2001)
- [40] Pit-Claudel, C., Philipoom, J., Jamner, D., Erbsen, A., Chlipala, A.: Relational compilation for performance-critical applications: Extensible proof-producing translation of functional models into low-level code. In: Proceedings of the 43rd ACM SIGPLAN International Conference on Programming Language Design and Implementation, pp. 918–933, PLDI 2022, Association for Computing Machinery, New York, NY, USA (Jun 2022), ISBN 978-1-4503-9265-5, <https://doi.org/10.1145/3519939.3523706>
- [41] Protzenko, J., Parno, B., Fromherz, A., Hawblitzel, C., Polubelova, M., Bhargavan, K., Beurdouche, B., Choi, J., Delignat-Lavaud, A., Fournet, C., Kulatova, N., Ramananandro, T., Rastogi, A., Swamy, N., Wintersteiger, C.M., Zanella-Beguelin, S.: EverCrypt: A Fast, Verified, Cross-Platform Cryptographic Provider. In: 2020 IEEE Symposium on Security and Privacy (SP), pp. 983–1002 (May 2020), ISSN 2375-1207, <https://doi.org/10.1109/SP40000.2020.00114>
- [42] Ray, S., Hunt, W.A., Matthews, J., Moore, J.S.: A Mechanical Analysis of Program Verification Strategies. *Journal of Automated Reasoning* **40**(4), 245–269 (May 2008), ISSN 1573-0670, <https://doi.org/10.1007/s10817-008-9098-1>
- [43] Reynolds, J.: Separation logic: A logic for shared mutable data structures. In: Proceedings of 17th IEEE Symposium on Logic in Computer Science (LICS (2002))
- [44] Rinard, M.: *Credible Compilation* (1999)
- [45] Sewell, T., Myreen, M., Klein, G.: Translation validation for a verified os kernel. In: Proceedings of the 34th ACM SIGPLAN Conference on Programming Language Design and Implementation, pp. 471–482 (2013)
- [46] Tan, J., Tay, H., Gandhi, R., Narasimhan, P.: Auspice: Automatic Safety Property Verification for Unmodified Executables. *VSSSTE*, In (2015)

- [47] Wang, A.: An axiomatic basis for proving total correctness of goto-programs. *BIT Numerical Mathematics* **16**(1), 88–102 (Mar 1976), ISSN 1572-9125, <https://doi.org/10.1007/BF01940782>

$$\begin{array}{c}
\frac{P' \subseteq P \quad \text{ensures}(P, Q, F)}{\text{ensures}(P', Q, F)} \text{ PRE} \qquad \frac{Q \subseteq Q' \quad \text{ensures}(P, Q, F)}{\text{ensures}(P, Q', F)} \text{ POST} \\
\\
\frac{F' \subseteq F \quad \text{ensures}(P, Q, F)}{\text{ensures}(P, Q, F')} \text{ FRAME} \\
\\
\frac{\text{ensures}(P, R, F) \quad \text{ensures}(R, Q, F')}{\text{ensures}(P, Q, F \circ F')} \text{ SEQ} \\
\\
\frac{\text{ensures}(P \cap B, Q, F) \quad \text{ensures}(P \cap \overline{B}, Q, F)}{\text{ensures}(P, Q, F)} \text{ BRANCH} \\
\\
\frac{\begin{array}{c} \text{ensures}(P, I(0), F) \\ \forall i \in \mathbb{N}. i < k \implies \text{ensures}(I(i), I(i+1), F) \\ \text{ensures}(I(k), Q, F) \end{array}}{\text{ensures}(P, Q, F)} \text{ LOOP}
\end{array}$$

Fig. 2. Derivation rules of the `ensures` predicate in \mathcal{L}_1 .

A Derivation Rules of `ensures`

Figure 2 shows the derivation rules of `ensures`, including the precondition weakening (PRE), postcondition strengthening (POST), frame monotonicity (FRAME), sequencing (SEQ), branching (BRANCH), and loop elimination rule (LOOP). Regarding LOOP, we require the invariant $I : \mathbb{N} \rightarrow \wp(\Sigma)$ to be defined for each iteration of the loop, and the postcondition Q to be satisfied after the k -th iteration for any k strictly greater than 0.

Note that, given a property of states P , we denote by \overline{P} the complement of P in the state space, i.e., $\overline{P} = \Sigma \setminus P$.

B Challenges in Extending Unary Logic

Defining relational Hoare triples in \mathcal{L}_2 is nontrivial and requires careful consideration to maintain the natural properties expected from a Hoare logic. Below, we present the challenges that arise when extending unary Hoare logic to a relational one and outline how we address these challenges.

Product Relation. A straightforward approach to reasoning about two programs is to define their operational semantics as the product of their individual transitions:

$$(s_0, s_1) \xrightarrow{\tau^2} (s'_0, s'_1) \stackrel{\text{def}}{\iff} s_0 \xrightarrow{\tau} s'_0 \vee s_1 \xrightarrow{\tau} s'_1$$

However, this approach is problematic. The operational semantics τ^2 may advance one program indiscriminately, potentially reaching a state where the

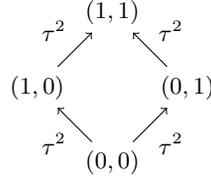


Fig. 3. Product relation τ^2

postcondition Q no longer holds for the pair, even if both programs satisfy Q individually under the precondition P .

For example, consider the operational semantics $\tau = \{(0, 1)\}$ where from state 0 we can reach state 1 in a single step. The product relation is τ^2 , depicted in Figure 3. The postcondition $Q = \{(0, 1)\}$ should eventually be satisfied by the starting pair of states $(0, 0)$ as the first program already satisfies its postcondition at the initial state 0, and the second program should eventually reach the state 1. However, from the initial state $(0, 0)$, one possible transition of the product relation is to the state $(1, 0)$ from which the postcondition Q is not reachable anymore. This behavior prevents the logic from proving intuitive properties.

Lockstep Simulation. An alternative approach is a lockstep simulation where both programs advance simultaneously:

$$(s_0, s_1) \xrightarrow{\tau^{\text{lock}}} (s'_0, s'_1) \stackrel{\text{def}}{\iff} s_0 \xrightarrow{\tau} s'_0 \wedge s_1 \xrightarrow{\tau} s'_1$$

While suitable for programs with identical control flow paths, this approach fails for cases where programs traverse diverging control flows, such as in equivalence checking for optimized and unoptimized implementations.

Nested Eventually Operators. Another candidate could be to compose two nested eventually operators to account for the two programs' behavior. The nested eventually operator would be defined as, for any two states s_0, s_1 :

$$(s_0, s_1) \in P \implies s_0 \in \text{eventually} \left(\left\{ s'_0 \mid s_1 \in \text{eventually} \left(\dot{Q}_{\pi_0=s'_0} \right) \right\} \right)$$

However, this approach is not be able to express two natural properties of Hoare logic. First, it loses symmetry as the nesting of eventually operators imposes ordering, meaning that the resulting logic would not commute between the two programs. For instance, in the constant-time proofs where the two programs are the same, a postcondition Q could be satisfiable while its inverse $\{(s_1, s_0) \mid (s_0, s_1) \in Q\}$ may not be, which is unnatural as the two programs are equal and would lead to a significant overhead in the verification framework. Second, it does not provide a compositional definition. Hence, we would not be able to split the verification of two programs into smaller, independent fragments. Increasing the complexity of the verification process.

C Additional Properties of `ensures2`

As a continuation of the discussion in Section 5, we present additional properties of the `ensures2` predicate in \mathcal{L}_2 .

As expected from a Hoare logic, the `ensures2` predicate can be weakened in the precondition, strengthened in the postcondition, and extended in the frame, as shown in the rules PRE, POST, and FRAME, respectively.

$$\frac{\text{ensures2}_{fn_0,fn_1}(P, Q^\times, F^\times) \quad P' \subseteq P}{\text{ensures2}_{fn_0,fn_1}(P', Q, F^\times)} \text{ PRE}$$

$$\frac{\text{ensures2}_{fn_0,fn_1}(P, Q, F^\times) \quad Q \subseteq Q'}{\text{ensures2}_{fn_0,fn_1}(P, Q', F^\times)} \text{ POST}$$

$$\frac{\text{ensures2}_{fn_0,fn_1}(P, Q, F) \quad F \subseteq F'}{\text{ensures2}_{fn_0,fn_1}(P, Q, F')} \text{ FRAME}$$

Additionally, we can derive a stronger Hoare triple by combining a weaker one with a restriction $f \subseteq \Sigma \times \Sigma$ provided that f contains all and only the pairs of states that are related by the frame condition. Formally:

$$\frac{\text{ensures2}_{fn_0,fn_1}(P^\times, Q^\times, F^\times) \quad (\forall s_0, s_1, s'_0, s'_1. ((s_0, s_1), (s'_0, s'_1)) \in F \implies ((s_0, s_1) \in f \iff (s'_0, s'_1) \in f))}{\text{ensures2}_{fn_0,fn_1}(P \cap f, Q \cap f, F)}$$

D Promotion of Unary Hoare Triples without Steps

If the original correctness proof is written in the `ensures` form, not `ensuresn`, the proof must be promoted to the `ensures` form first. For this reason, we introduce the following `eventuallynatpc` property, stronger than `eventuallyn`, which allows us to promote a proof made via the `eventually` operator to `eventuallyn` for a program that halts at a specific program counter.

Definition 10 (Eventually n at pc). *Given a number of steps $n \in \mathbb{N}$, two addresses $x_0, x_\omega \in \mathbb{N}$ for the initial and final program counters, and a precondition $P \subseteq \Sigma$, we define:*

$$\text{eventuallynatpc}_n^{x_0, x_\omega}(P) \stackrel{\text{def}}{=} \left\{ s \in \Sigma \mid \begin{array}{l} \forall F \subseteq \Sigma \times \Sigma. s(\text{instr}) = x_0 \wedge s \in P \\ \implies s \in \text{eventually}^\tau(\{s' \mid s'(\text{instr}) = x_\omega \wedge (s, s') \in F\}) \\ \implies s \in \text{eventuallyn}_n^\tau(\{s' \mid s'(\text{instr}) = x_\omega \wedge (s, s') \in F\}) \end{array} \right\}$$

Note that this `eventuallynatpc` predicate does not hold for an arbitrary assembly program in general because the program execution may go through

the postcondition (which is a part of frame F) at x_ω , jump back prior to x_ω and eventually meet the F . In general, there can be multiple n s that meet the postcondition in F .

For example, let us assume that $\text{prog}(P)$ is the single line program adding two operators: `add x1, x2, x3`; we define x_ω equal to x_0+1 . It would be tempting to state that `eventuallynatpc` holds when $n = 1$ because the program counter arrives at x_ω after a single step. However, $n = 1$ does not satisfy `eventuallynatpc` because there is a possibility that the next instruction is a decodable instruction that jumps back before the addition. In this case, any number of steps $n \geq 1$ will arrive at x_ω .

To address this problem, we introduce a *stopper*: a byte sequence that fails to decode. In `s2n-bignum`, we use the 4-byte zeros (`0x00000000`) as a stopper sequence. In `eventuallynatpc`, we constrain the program of interest to end with this stopper sequence. With this stopper sequence appended, we can prove `eventuallynatpc` for a reasonable program and n . For the previous example, $n = 1$ would be valid.

The next result highlights the promotion of `ensures` to `ensuresn` given a proof of `eventuallynatpc`.

Lemma 6 (Promotion of ensures to ensuresn).

$$\begin{aligned} \text{eventuallynatpc}_n^{x_0, x_\omega}(P) &\implies \\ \forall Q, F. \text{ensures}^\tau(P^{pc}, Q^{pc}, F) &\implies \text{ensuresn}_{\lambda s.n}^\tau(P^{pc}, Q^{pc}, F) \end{aligned}$$

where $P^{pc} = \{s \in P \mid s(\text{pc}) = x_0\}$ and $Q^{pc} = \{s \in Q \mid s(\text{pc}) = x_\omega\}$.

As `eventuallynatpc` $_n^{x_0, x_\omega}(P)$ must end with the stopper sequence, also the precondition P of both `ensures`(P, Q, F) and `ensuresn` $_{fn}$ (P, Q, F) in the lemma above should handle the stopper sequence as well. We need to consider take care of both when lifting the existing proofs.

We lift existing proofs of `ensures`(P, Q, F), c.f. the left-hand side of the implication in Lemma 6, that do not mention the stopper sequence by the fact that `ensures` $^\tau(P \wedge P', Q, F) \implies \text{ensures}^\tau(P, Q, F)$ holds, where P' ensures that the stopper sequence ends the program.

The case of the right-hand side of the implication in Lemma 6 seems to be more complicated. Indeed, we cannot apply the same trick as above because `ensuresn` $_{\lambda s.n}^\tau(P \wedge P', Q, F) \implies \text{ensuresn}_{\lambda s.n}^\tau(P, Q, F)$ does not hold. However, we notice that when combining `ensuresn` $_{fn}$ (P, Q, F) with a program equivalence proof (defined as `ensures2`), the stopper sequence does not appear in the final specification. The reason is that this combination creates a hybrid ensure with its P, Q, R left under an existential quantifier. Thus, after applying Theorem 4, the condition regarding the stopper sequence is removed from the final precondition.