

# SideTrail: Verifying Time-Balancing of Cryptosystems

Konstantinos Athanasiou<sup>1</sup>, Byron Cook<sup>2</sup>, Michael Emmi<sup>3</sup>, Colm MacCarthaigh<sup>2</sup>,  
Daniel Schwartz-Narbonne<sup>2</sup>, and Serdar Tasiran<sup>2</sup>

<sup>1</sup> Northeastern University

<sup>2</sup> Amazon Web Services

<sup>3</sup> SRI International

**Abstract.** Timing-based side-channel attacks are a serious security risk for modern cryptosystems. The time-balancing countermeasure used by several TLS implementations (*e.g.* s2n, GnuTLS) ensures that execution timing is negligibly influenced by secrets, and hence no *attacker-observable timing behavior* depends on secrets. These implementations can be difficult to validate, since time-balancing countermeasures depend on global properties across multiple executions. In this work we introduce the tool SIDE`TRAIL`, which we use to prove the correctness of time-balancing countermeasures in s2n, the open-source TLS implementation used across a range of products from AWS, including S3. SIDE`TRAIL` is used in s2n’s continuous integration process, and has detected three side-channel issues that the s2n team confirmed and repaired before the affected code was deployed to production systems.

## 1 Introduction

Timing-based side-channel attacks are a serious security risk for modern cryptosystems; the Lucky 13 attack is a recent example [1]. Current systems deploy one of two prevailing countermeasures to prevent such attacks. One possible mitigation against this threat is to apply the *constant-time* coding principle, where secrets must not influence control-flow paths, memory access patterns, or the cycle counts of instructions. This simplifies local reasoning about timing leaks: if secrets are not used in the prohibited manner, then the code does not exhibit timing side-channels. However, constant-time coding often requires replacing a program’s natural control-flow with complicated bitwise operations, and can require significant changes to standard data-structures and APIs, making it difficult to reason about functional correctness. The developers of OpenSSL recently applied a 500+ LOC patch to perform constant-time cipher block chaining (CBC) decoding; the complexity of which led to subsequent issues [2].

The second approach, dubbed *time-balancing*, ensures that execution time is negligibly influenced by secrets. This relaxation from constant-time enables simpler and more readable countermeasures: developers must *balance* a program’s executions to have similar timing footprints, allowing the use of standard operations that depend on secrets. The CBC code from s2n [3], for example, implements time-balancing in fewer than 20 additional lines, and s2n’s time-balanced HMAC has been proven functionally correct [4]. However, since time-balancing countermeasures depend on global properties across multiple executions, programmers easily miss subtle timing leaks [5].

In this work, we introduce SIDETRAIL (available at [6]), a deductive verifier for time-balancing countermeasures. SIDETRAIL uses an instruction-level precise timing model and encodes it via a *time counter*. It uses Boogie [7] to precisely reason about control flow and values, including the time counter. We automatically infer invariants over time-counter values with minimal user annotation, and use self-composition [8] to prove that the timing difference between every pair of executions with similar public inputs is below a given bound.

We have used SIDETRAIL to verify the correctness of the timing countermeasures for the data-packet processing stack in s2n. SIDETRAIL is used in s2n’s Travis-based continuous integration system [9], has detected three issues that the s2n team has confirmed and repaired before the affected code was used in production systems, and has proved correctness of the repaired countermeasures.

## 1.1 Related Work

Prior work has proposed verification for constant-time [10, 11, 12, 13, 14, 15, 16, 17, 18, 19, 20] and power-balancing [21] side-channel countermeasures including the Everest project which has proven the functional correctness of a side-channel resistant TLS implementation in the constant-time model [22, 23].

Different approaches have recently appeared in the context of *time-balancing* countermeasures [24] [25]. Blazer [24] uses a *partitioning strategy* instead of self-composition and scales to examples of up to 100 basic blocks in size. Themis [25] uses *Quantitative Cartesian Hoare Logic* to capture timing differences between executions and scales to real-world Java applications. Themis requires functions to be time-balanced; it otherwise reports spurious violations. In contrast, SIDETRAIL takes a *path-based* approach, and can handle time-balancing countermeasures, such as those used in s2n’s HMAC, which compose unbalanced functions to ensure that every execution path is time balanced. Both Blazer and Themis target Java programs while SIDETRAIL focuses on low-level C-code implementations. Unlike these two approaches, which approximate leakage using computational complexity and bytecode instruction counts respectively, SIDETRAIL supports fine grained cost models at the instruction level of LLVM’s Intermediate Representation language.

The self-composition we describe in Sec. 3 builds on those of existing works [26, 27, 28, 29], and our implementation follows ct-verif’s [19], replacing its cross-product construction with self-composition; our novelty is in its application to time balancing. Contrary to approaches providing upper bounds of information leakage metrics [30], SIDETRAIL employs relational verification.

## 2 Time-Balancing

Time-balancing countermeasures provide the security assurance that program secrets have negligible influence on execution time, even in the presence of potentially malicious observers who can control the public inputs to the program. Formally, a program is  $\delta$ -secure if for every possible public-input value, the timing difference between every pair of executions with different secrets is at most  $\delta$ . In this work, we assume a standard

<pre> 1: <b>procedure</b> CBC-VULNERABLE 2:   pad := packet[len - 1] 3:   payload_len := len - pad 4:   update(mac, packet, payload_len) 5:   digest(mac) </pre>	<pre> 1: <b>procedure</b> CBC-TIMEBALANCED 2:   pad := packet[len - 1] 3:   payload_len := len - pad 4:   update(mac, packet, payload_len) 5:   update(dummyMAC, packet + payload_len, pad) 6:   digest(mac) </pre>
--	---

Fig. 1: A vulnerable TLS CBC algorithm (a), and its time-balanced version (b).

adversarial model: a network man in the middle (MITM), who has the ability to observe and modify both the contents and the timing of any packet on the network, but who cannot execute arbitrary code on the targeted endpoints [31]. This model is powerful enough to capture a wide range of real TLS attacks, ranging from Bleichenbacher’s attacks against PKCS #1 [32], to Brumley and Boneh’s attack against RSA decryption [33], to the Lucky 13 family of attacks against the TLS HMAC [1].

*Example 1.* The Lucky 13 family of attacks [1] takes advantage of a weakness in the specification for SSL/TLS; CBC mode uses an HMAC to prevent adversaries from modifying incoming cipher-texts, but neglects to protect the padding [34]. A MITM attacker can trick a TLS implementation into decrypting a selected encrypted byte (e.g. from a password) into the padding length field. Fig. 1a shows what happens in a naïve implementation: in line 4, `len - pad` bytes are hashed by the HMAC, whose timing strongly depends on the number of bytes hashed. Since `len` is known to the attacker, this creates a timing side-channel leaking the value of `pad`. A constant-time mitigation would be to rewrite the HMAC implementation so that its computation was independent of the value of `pad`. A simpler time-balanced countermeasure is shown in Fig. 1b: apply `update` on a dummy HMAC state `dummyMAC` (line 5), which ensures that no matter the value of `pad`, the HMAC will always process a total of `len` bytes, mitigating the timing side-channel.

*Verifying time-balancing.* Following previous approaches to verifying non-interference [26, 29], we reduce verification of  $\delta$ -security, which is a 2-safety property, i.e. a property over pairs of executions, to the verification of a standard safety property, i.e. over individual executions. For technical simplicity, we consider a program  $P$  as a sequence of instructions  $p_0; p_1; \dots; p_n$  whose variables  $V = Sec \uplus Pub$  are the disjoint union of secret and public variables, respectively; a program state  $s$  maps the variables to values. A configuration  $c$  is a tuple  $(s, p)$  of program state  $s$  and the next instruction  $p$  to be executed. An execution is defined as a configurations sequence  $c_1 c_2 \dots c_m$ .

Effective reasoning about  $\delta$ -security requires (i) a timing model that accurately captures programs’ timing behavior, and (ii) a verification technique able to relate the timing behavior of two distinct program executions. To capture timing, we introduce a leakage function  $\ell(c)$ , mapping configurations  $c$  to timing observations, i.e. the cost of executing the next program instruction using the values of variables in  $c$ . To keep track of the total cost of an execution we extend the set of variables with a *time counter*  $\mathbb{1}$  as  $V_L = V \uplus \{\mathbb{1}\}$  and write the *time counter instrumented* program  $P_L$  as  $l_1; p_1; l_2; p_2 \dots; l_n; p_n$ , in which each instruction  $l_i$  updates the time counter variable as  $\mathbb{1} := \mathbb{1} + \ell(s, p_i)$ . Finally to relate the timing cost of two execution paths we compose  $P_L$  with its renaming

$\hat{P}_L$  over variables  $\hat{V}_L$ , and form its self-composition  $P_L; \hat{P}_L$  ranging over variables  $V_L \cup \hat{V}_L$  [35, 8]. Accordingly,  $\delta$ -security can be specified as a safety property over the time-counter variables  $\mathbf{1}$  and  $\hat{\mathbf{1}}$ .

### 3 Implementation

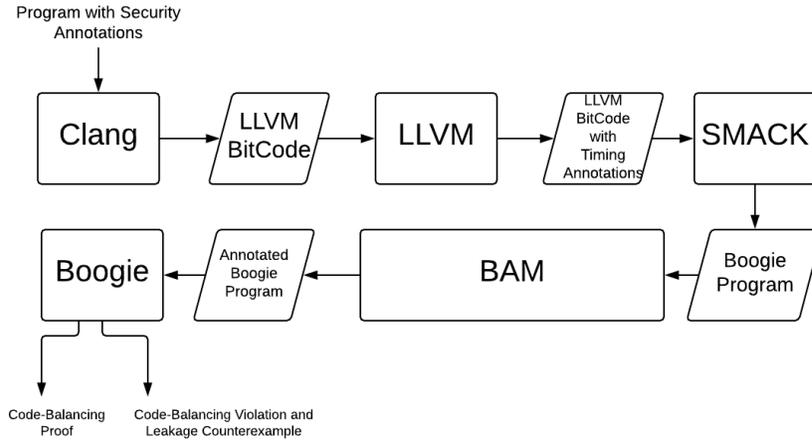


Fig. 2: SIDETRAIL Architecture

SIDETRAIL uses the SMACK verification infrastructure [36], which leverages Clang [37] and LLVM [38] to generate and optimize LLVM bitcode before translation to the Boogie [7] intermediate verification language. Using LLVM intermediate representation (IR) allows SIDETRAIL to reason precisely about the effect of compiler optimizations, which can affect the timing of time-balancing countermeasures. Our initial experience verifying simple time-balanced examples showcased the importance of correctly accounting for compiler optimizations. In some cases, the compiler can notice that the extra code added for time-balancing is side-effect free, and remove it, reintroducing the timing side-channel in the original algorithm.

In addition, using LLVM IR easily allows us to use an instruction-level-precise architectural timing model. We have extended SMACK to introduce the timing cost of LLVM instructions, and implement program transformations at the Boogie code-level, passing a time counter instrumented self-composition to the Boogie verifier [7]. If the program being verified is  $\delta$ -secure, SIDETRAIL returns with a proof; if there is a violation of  $\delta$ -security, SIDETRAIL provides a counter-example trace which includes the input values that triggered the exception, the trace leading to the exception, and the amount of leakage calculated for that trace. The SIDETRAIL flow is illustrated in Fig. 2 and Fig. 3.

```

#define DELTA 0      define i32 @plus(   procedure plus(   var l, _l: int;
                    i32, i32){         i0: int, i1: int) procedure wrapper(
                    %3 = add nsw i32     returns (r: int){  i0: int, i1: int,
                    %1, %0,              var i2: int;        _i0: int, _i1: int){
int plus(           }                 bb0:                assume i0 == _i0;
  int a, int b){    !TimeCost !19    l := l + 1;        call plus(i0, i1);
  public_in(a);     ret i32 %3,    i2 := i0 + i1;    call _plus(_i0, _i1);
  assert_leakage(   !TimeCost !27    l := l + 0;        assume l >= _l;
    DELTA);         }                 r := i2;           assert l - _l <= 0;
  return a + b;     }                 return;            return;
}
...                !19 = !{i64 1}
...                !27 = !{i64 0}
                    }
                    }

```

Fig. 3: Stages of SIDETRAIL translation, from left to right: (a) an annotated C-code add function; (b) the corresponding LLVM IR with timing annotations; (c) the translated Boogie code with time counter instrumentation; and (d) the Boogie code for the self-composition.

*Security Annotations:* SIDETRAIL requires a small number of annotations at the source-code level to specify a  $\delta$ -security proof. The programmer must annotate the appropriate entry-point arguments as `public` (unannotated arguments are treated as secrets) and specify the non-negative integer timing-difference bound ( $\delta$ ). The `public_in` and `assert_leakage` annotations of Fig. 3a serve these purposes.

*Timing Model:* SIDETRAIL uses LLVM’s Cost Model Analysis for its instruction-level precise timing model. The analysis approximates the timing cost of each instruction when lowered to machine code by mapping it to a positive integer, and assumes that all memory accesses hit their respective caches; we discuss the soundness of this assumption in Sec. 4.3. Fig. 3b shows how SIDETRAIL annotates the LLVM IR `add` (time cost: 1) and `ret` (time cost: 0) instructions with timing metadata, represented as metadata pointers `!19` and `!27` respectively. Fig. 3c shows how the timing metadata are carried over to the Boogie code. A *time-modeling* transformation, implemented as a Boogie-code transformation, introduces the integer-type time counter variable `l` and updates it in lockstep with program instructions.

*Loop Invariants:* To capture how the values of the time counter variables are updated throughout a loop’s execution, SIDETRAIL automatically inserts *loop-timing-invariants* in the Boogie code, based on annotations provided by the user.

A loop’s cost depends on two factors: the number of times it iterates, and the cost of each iteration. The number of iterations can be captured by annotating a simple continuation invariant — for example, loops of the form `for` (`i=0; i<n; ++i`) should be user-annotated with a continuation invariant (`i<=n`). In the common case where the execution time of the loop-body does not vary across iterations, SIDETRAIL can automatically infer the cost of each iteration. If the loop body contains nested control statements the user must provide annotations that describe how many times each branch of the control statement is visited, although we note that we encountered only a single loop with nested control in our experiments. In either case, SIDETRAIL automatically infers timing-invariants of the form  $(l = l\_prior + i * body\_cost)$ , where `l_prior`

is the value of `l` before entering the loop and `body_cost` is the timing cost of executing the loop `body` once, and inserts them in the Boogie code.

*Self-Composition:* We implement the self-composition-based reduction of  $\delta$ -security to assertion checking as an additional Boogie-code transformation, demonstrated in Fig. 3d. We duplicate the program to be verified, making a renamed copy of all functions (`plus` becomes `_plus`), and then transform these duplicated functions to use renamed copies of global variables including the time counter, and to perform nested procedure calls on the renamed procedures. Finally, a `wrapper` procedure enforces the equality of public inputs, makes two consecutive calls to the entry function of the program and its renamed copy, and adds an assertion to check  $\delta$ -security.

*Inter-procedural Analysis:* SIDETRAIL supports inter-procedural analysis through function inlining, allowing the analysis of arbitrary entry points which may invoke individually unbalanced functions (`s2n` uses such functions in its path-balanced HMAC). As we discuss in Sec. 4, this approach is able to handle industrial codebases such as the `s2n` HMAC. SIDETRAIL also supports modular verification through timing stubs, described below.

*Timing Stubs:* SIDETRAIL allows the user to specify the expected leakage from a function by providing support for *timing stubs*. Users specify these stubs by adding `assume_leakage(expr)` statements to the body of a function, where `expr` is any expression computable within the function. When the *time-modeling* transformation encounters this call, it increases the time counter variable by `expr`. This allows stubs to represent complex timing behaviour which may depend on properties of both the input and current state of the function. It is the responsibility of the user to ensure that the stub correctly models the timing behaviour of the concrete implementation; we discuss how SIDETRAIL can be used to verify the correctness of stubs in Sec. 4.2.

## 4 Case Study of the `s2n` TLS Library

`s2n` is an open-source TLS library used by Amazon, including S3, and AWS services [39]. Its design goals are that it be: “small, fast, with simplicity as a priority” [40]. Its time-balanced CBC mode requires fewer than 20 additional lines of code, compared to the 500+ LOC patch to perform constant-time CBC decoding in OpenSSL. We have used SIDETRAIL to verify the correctness of the timing countermeasures for the whole data-packet processing stack (which includes CBC verification as a sub-component) in the current `s2n` release. In the process we have discovered three previously unknown issues that violate  $\delta$ -security. We note `s2n` has a belt-and-suspenders security model with randomized delay on error and a secure default configuration [41], which would have prevented these issues from affecting data in production.

The proofs described below are automatically rerun as part of `s2n`’s Travis based continuous integration system. This ensures that code changes are only accepted to the `s2n` repository after they have been validated using SIDETRAIL.

```

int s2n_hash_update(struct s2n_hash_state *state, const void *data, uint32_t size)
{
    assume_leakage(PER_BYTE_COST * size);
    state->currently_in_hash_block += size;
    int num_filled_blocks = state->currently_in_hash_block / BLOCK_SIZE;
    assume_leakage(num_filled_blocks * PER_BLOCK_COST);
    state->currently_in_hash_block = state->currently_in_hash_block % BLOCK_SIZE;
    return SUCCESS;
}

```

Fig. 4: Hash function timing stub.

#### 4.1 Timing Stubs

s2n is written in a modular fashion, and does not implement its own cryptographic primitives – instead they are linked from the system `libcrypto`. We provide timing stubs for each cryptographic primitive used by s2n, following the approach described in Sec. 3. An example stub for `hash_update` is shown in 4. This stub has two components: a per-byte cost, representing the cost of `memcpy`ing the data, and a per-block cost, representing the cost of a hash compression round. The stubs were validated by using `SIDETRAIL` to verify that the stub had the same timing behaviour as a C implementation, as described in Experiment 3 (Sec. 4.2).

#### 4.2 Experiments

Except as noted below, all experiments had approximately 400 lines of initial source, 5 security annotations, 3 loop invariants, 100 lines of additional code (stubs + test harness), and expanded to an order of 1000 SMT clauses. All of the experiments listed below completed in less than 8 minutes on a 3.1 GHz Intel Core i7 with 16GB of RAM running OSX 10.11, using Z3 4.6.0’s integer arithmetic theory.<sup>4</sup> The code for all experiments is available online.<sup>5</sup>

For each experiment, we determined the precise amount of  $\delta$ -leakage by varying  $\delta$  to find the boundary where verification moved from unsuccessful to successful. Our experience in this process suggests that refuting values of  $\delta$  that are too small is typically faster than verifying values of the correct size.

**Properties Verified** We performed four related verification experiments using `SIDETRAIL`. Experiment 1 validated our ability to detect a previously reported timing issue in the CBC mode of s2n. Experiment 2 validated the correctness of the current implementation of this code. Experiment 3 validated the correctness of the timing stubs used in Experiment 2. Experiment 4 extends this proof to provide an end to end guarantee for data-packet processing.

<sup>4</sup> Using integer arithmetic provides performance benefits at the cost of losing information about the underlying C types. In a few cases, we needed to annotate back this information, for example by adding `assume(x >= 0 && x < 256)` after an assignment to a `uint8_t x`;

<sup>5</sup> <https://github.com/danielsn/s2n/tree/sidetrail-vstte-artifact/tests/sidewinder>

*[Experiment 1]* One of the motivations for this work is a previously reported and repaired timing side-channel issue [42] in s2n’s time-balanced CBC decoder. The issue is caused by an off-by-one error in the code that tracks how many bytes have been hashed and triggers for a particular edge-case in the padding size. SIDETRAIL reports an error trace that includes the concrete value of the padding size, as well as the expected leakage (equivalent to one hash compression round, approximately 1 microsecond).

*[Experiment 2]* In this experiment, we verify the correctness of the CBC mode time-balancing countermeasures for all protocol versions (SSL3, TLS[1.0–1.2]) and hash functions (MD5, SHA1, SHA-224, SHA-256, SHA-384, SHA-512) supported by s2n. In order to handle this wide range of modes and functions, both the test-harness and the hash-stubs were written in a generic fashion, allowing different modes to be tested by setting the appropriate compile-time constants. The experiment could then be rerun with varying parameters to cover all modes.

We experienced two types of scalability issues in this experiment. Firstly, s2n does not contain its own cryptographic primitives, leading us to use a modular verification approach using timing stubs as described in Sec. 4.1. Secondly, the CBC code makes branching decisions based on non-linear operators (`mod` and `div`), which the backend SMT solver had difficulty solving in a reasonable time. We worked around this issue by replacing the `div` operation with a handwritten variant that uses pre-calculated values for the block-sizes being tested. Our proofs for the various modes all showed a small  $\delta$ -leakage, caused by the extra call to `hash_digest` necessary to affect time-balancing. Since this  $\delta$  (approximately 0.03 microseconds) is significantly smaller than the time granularity visible to a network based attacker (estimated at 1 microsecond based on timing experiments done between co-located machines, which is in agreement with [31]), the code is successfully balanced.

*[Experiment 3]* The soundness of Experiment 2 depends on correct modeling of the timing behaviour of the timing stubs. In this experiment we verify that our SHA-1 timing stub accurately captures the timing cost of an open-source C implementation [43]. Our proof methodology is similar to the one used for  $\delta$ -security, but in this case we form a composition of the timing stub (Fig. 4) with the C implementation and assert that their time counter variables are equal (i.e.  $\delta = 0$ ). The SHA-1 C code contains a loop with a nested control statement, which requires us to indicate the times each branch is exercised via an additional invariant. After experimenting to find the correct values for `PER_BYTE_COST` and `PER_BLOCK_COST`, we verify that the  $\delta$  between the stub and the C implementation is 0.

*[Experiment 4]* (450LOC Source, 150LOC stubs, 20LOC Annotations, 1860 SMT clauses) In this experiment we verify that data-packet processing is time-balanced, for all protocols (SSL3, TLS[1.0–1.2]) and modes (AEAD, CBC, Composite, and Stream) currently supported by s2n. This provides end-to-end confidence that from the time a data-packet is decrypted, until when the bytes are returned to the client, all paths are time-balanced.

The proof decomposes packet processing into three phases. Phase one, packet parsing, operates on public data, such as header fields, which are already known to a MITM attacker, and hence can be treated as public inputs. Phase two, decryption, is handled by stubs which havoc the decrypted data buffer, making the values non-deterministically different across the self-composed program executions. Phase three validates the decrypted data and returns it to the user. This phase uses the decrypted data from stage two, and hence is the stage which could potentially leak confidential information via a timing side-channel. We leverage the  $\delta$ -security proof and the concrete value of  $\delta$  reported from Experiment 2 to validate the timing stub for HMAC verification. All modes reported a  $\delta < 0.05$  microseconds.

**Previously Unknown Issues Discovered** Using SIDETRAIL, we reported three previously unknown issues in the s2n time-balancing countermeasures, which have been acknowledged and fixed by the s2n team, and validated the repaired code.

*[Issue 1]* s2n time-balancing code counts the number of bytes in the hash block, and uses this value as part of its time-balancing countermeasures. As an optimization, s2n took advantage of the fact that the standard HMAC specification [44] pads keys to be multiples of the hash-block in length, which means that hashing a padded key does not affect the number of bytes remaining in the hash block. Unfortunately, the SSLv3 specification follows this recommendation for MD5, but in SHA1 mode uses a padded key that is 4 blocks short of the hash block size [45], causing the time-balancing code to incorrectly count the number of bytes.

*[Issue 2]* s2n's HMAC had two variables with similar names: `hash_block_size` and `block_size`. Due to a typo, the wrong variable was used in a modular operation which determines the need for time-balancing operations. Interestingly, this issue only exposes itself in SSLv3 mode; in all other modes, `block_size` and `hash_block_size` have the same value. This issue was accepted and repaired by the s2n team, who also renamed the variables to have more descriptive names.

*[Issue 3]* As discussed in Issue 1, s2n uses a count of the number of bytes hashed as part of its time-balancing countermeasures. In addition to the bytes directly added by calls to `hash_update()`, s2n must track the number of bytes added behind the scenes by the hash algorithm. In particular, before generating a hash digest, most hash algorithms append padding, which typically includes an integer specifying the total number of bytes that have been hashed. Most hash functions used in TLS (e.g. MD5, SHA1) use an 64-bit (i.e. 8-byte) integer for this purpose, and hence s2n's time-balancing code adds 8 to the number of bytes that have been hashed when determining the need for time-balancing operations. However, some hash algorithms, such as SHA-384 and SHA-512, append a 128 bit (16-byte) integer, which would cause s2n time-balancing code to miscount the number of bytes hashed when using these algorithms.

### 4.3 Discussion on the Timing model

SIDETRAIL assumes an architectural timing model, which abstracts away micro-architectural features such as caching and branch prediction, similarly to other approaches [24, 25]. This model captures the capabilities of a MITM network attacker against a TLS endpoint who can measure final execution time (with limited precision due to network jitter) but cannot affect or directly observe machine state. Consequently the attacker can only observe the cumulative number of cache hits/misses, and cannot influence them by altering the state of the cache. In the context of our s2n case study, the TLS code we verify preloads a data-packet whose maximum size is 16KB into the cache, and then spends approximately 50 microseconds doing a linear scan across it (based on the speed of the hash functions used in TLS). Since the data-packet fits in the L1 cache (typically at least 32KB), all memory accesses results in cache hits. Additionally, 50 microseconds are considerably less than the Linux quantum which is on the order of milliseconds, meaning that cache interference effects from context switches are minimal.

### 4.4 Verification Inspired Refactoring

Verifying an industrial code-base requires forming a clear understanding of the code being verified. As part of this effort, we discovered a number of refactorings that made the code more modular, clean, and easy to verify. These code changes were shared with the s2n team as GitHub PRs, and have been merged into the mainline code-base.

The proposed changes ranged from small optimizations such as removing an unnecessary loop, to refactoring of larger portions of the code-base. For example, as part of Experiment 2, we discovered that different HMAC modes had duplicated functionality; merging this functionality into a common function simplified both the code, and the proof effort. Conversely, in Experiment 4, the data-packet processing code interleaved functionality from several different modes, requiring a large number of local variables and making it difficult to write a test harness that covered all cases. We split the code into four simpler functions, with four simple test harnesses. The s2n team accepted the PR with the comment “this looks much better, thanks.” The most interesting fix was to s2n’s error handling which follows a disciplined methodology. As we analysed several error handling code paths, we realized that every case followed the same template, and could be simplified with a macro that made error-handling easier to annotate (since we only needed to add the annotation in one place). As an added benefit, this change removed 400 LOC from an approximately 6 KLOC code-base.

Formal verification both inspired and enabled these changes. Our refactoring touched large portions of the overall code-base, and made changes to security-critical functionality. Without automated formal proofs to give us confidence that our changes would not introduce new timing regressions, the amount of effort to manually validate the changes would have made these changes impractical.

## 5 Future work

As future work we first aim to improve the accuracy of SIDETRAIL’s timing model. Modeling the behaviour of micro-architectural components, such as the cache and the

branch predictor, will extend the class of attackers that SIDETRAIL can reason about with on-machine active attackers. Additionally, replacing LLVM's instruction cost model with a model based on micro-benchmarks will increase the precision of SIDETRAIL, especially for operations such as `div` and `mod` which take a different number of cycles depending on the values of their input. We have demonstrated SIDETRAIL's capability to utilize and validate timing stubs. As a second direction of future work, we envision extending the tool's usability by inferring timing stubs in an automated fashion.

## 6 Conclusion

Ideal cryptographic practice is to design algorithms for an easy and straightforward implementation that is naturally constant-time. For legacy algorithms that were not designed with this restriction in mind, developers must use alternate approaches such as time-balancing. SIDETRAIL allows developers of industrial cryptographic code-bases such as `s2n` to verify the correctness of these mitigations, and to detect issues and regressions when they occur.

## References

- [1] AlFardan, N.J., Paterson, K.G.: Lucky thirteen: Breaking the TLS and DTLS record protocols. In: 2013 IEEE Symposium on Security and Privacy, SP 2013, Berkeley, CA, USA, May 19-22, 2013, IEEE Computer Society (2013) 526–540
- [2] Somorovsky, V.J.: Curious Padding oracle in OpenSSL (CVE-2016-2107). <https://web-in-security.blogspot.co.uk/2016/05/curious-padding-oracle-in-openssl-cve.html> (2016) [Online; accessed 15-Jan-2018].
- [3] Amazon Web Services: `s2n`: an implementation of the TLS/SSL protocols. <https://github.com/awslabs/s2n> (2018)
- [4] Dodds, J.: Part one: Verifying `s2n` HMAC with SAW. <https://galois.com/blog/2016/09/verifying-s2n-hmac-with-saw/> (2016) [Online; accessed 15-Jan-2018].
- [5] Albrecht, M.R., Paterson, K.G.: Lucky Microseconds: A Timing Attack on Amazon's `s2n` Implementation of TLS. In: Proceedings of the 35th Annual International Conference on Advances in Cryptology — EUROCRYPT 2016 - Volume 9665, New York, NY, USA, Springer-Verlag New York, Inc. (2016) 622–643
- [6] Sidewinder: Time-balanced Verification Tests. <https://github.com/awslabs/s2n/tree/master/tests/sidewinder> (2018)
- [7] Barnett, M., Chang, B.E., DeLine, R., Jacobs, B., Leino, K.R.M.: Boogie: A modular reusable verifier for object-oriented programs. In de Boer, F.S., Bonsangue, M.M., Graf, S., de Roever, W.P., eds.: Formal Methods for Components and Objects, 4th International Symposium, FMCO 2005, Amsterdam, The Netherlands, November 1-4, 2005, Revised Lectures. Volume 4111 of Lecture Notes in Computer Science., Springer (2005) 364–387
- [8] Barthe, G., D'Argenio, P.R., Rezk, T.: Secure information flow by self-composition. In: Computer Security Foundations Workshop, 2004. Proceedings. 17th IEEE, IEEE (2004) 100–114

- [9] Amazon Web Services: s2n Travis CI integration page. <https://travis-ci.org/awslabs/s2n/> (2018)
- [10] Agat, J.: Transforming out timing leaks. In Wegman, M.N., Reps, T.W., eds.: POPL 2000, Proceedings of the 27th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, Boston, Massachusetts, USA, January 19-21, 2000, ACM (2000) 40–53
- [11] Molnar, D., Piotrowski, M., Schultz, D., Wagner, D.A.: The program counter security model: Automatic detection and removal of control-flow side channel attacks. In Won, D., Kim, S., eds.: Information Security and Cryptology - ICISC 2005, 8th International Conference, Seoul, Korea, December 1-2, 2005, Revised Selected Papers. Volume 3935 of Lecture Notes in Computer Science., Springer (2005) 156–168
- [12] Svenningsson, J., Sands, D.: Specification and verification of side channel declassification. In Degano, P., Guttman, J.D., eds.: Formal Aspects in Security and Trust, 6th International Workshop, FAST 2009, Eindhoven, The Netherlands, November 5-6, 2009, Revised Selected Papers. Volume 5983 of Lecture Notes in Computer Science., Springer (2009) 111–125
- [13] Stefan, D., Buiras, P., Yang, E.Z., Levy, A., Terei, D., Russo, A., Mazières, D.: Eliminating cache-based timing attacks with instruction-based scheduling. In Crampton, J., Jajodia, S., Mayes, K., eds.: Computer Security - ESORICS 2013 - 18th European Symposium on Research in Computer Security, Egham, UK, September 9-13, 2013. Proceedings. Volume 8134 of Lecture Notes in Computer Science., Springer (2013) 718–735
- [14] Almeida, J.B., Barbosa, M., Pinto, J.S., Vieira, B.: Formal verification of side-channel countermeasures using self-composition. *Sci. Comput. Program.* **78**(7) (2013) 796–812
- [15] Barthe, G., Betarte, G., Campo, J.D., Luna, C.D., Pichardie, D.: System-level non-interference for constant-time cryptography. In Ahn, G., Yung, M., Li, N., eds.: Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security, Scottsdale, AZ, USA, November 3-7, 2014, ACM (2014) 1267–1279
- [16] Zhang, D., Wang, Y., Suh, G.E., Myers, A.C.: A hardware design language for timing-sensitive information-flow security. In Öztürk, Ö., Ebcioğlu, K., Dwarkadas, S., eds.: Proceedings of the Twentieth International Conference on Architectural Support for Programming Languages and Operating Systems, ASPLOS '15, Istanbul, Turkey, March 14-18, 2015, ACM (2015) 503–516
- [17] Doychev, G., Köpf, B., Mauborgne, L., Reineke, J.: Cacheaudit: A tool for the static analysis of cache side channels. *ACM Trans. Inf. Syst. Secur.* **18**(1) (2015) 4:1–4:32
- [18] Rodrigues, B., Pereira, F.M.Q., Aranha, D.F.: Sparse representation of implicit flows with applications to side-channel detection. In Zaks, A., Hermenegildo, M.V., eds.: Proceedings of the 25th International Conference on Compiler Construction, CC 2016, Barcelona, Spain, March 12-18, 2016, ACM (2016) 110–120
- [19] Almeida, J.B., Barbosa, M., Barthe, G., Dupressoir, F., Emmi, M.: Verifying constant-time implementations. In: 25th USENIX Security Symposium, USENIX Security 16, Austin, TX, USA, August 10-12, 2016. (2016) 53–70

- [20] Blazy, S., Pichardie, D., Trieu, A.: Verifying constant-time implementations by abstract interpretation. In: European Symposium on Research in Computer Security, Springer (2017) 260–277
- [21] Fang, X., Luo, P., Fei, Y., Leeser, M.: Leakage evaluation on power balance countermeasure against side-channel attack on fpgas. In: 2015 IEEE High Performance Extreme Computing Conference, HPEC 2015, Waltham, MA, USA, September 15-17, 2015, IEEE (2015) 1–6
- [22] Bond, B., Hawblitzel, C., Kapritsos, M., Leino, K.R.M., Lorch, J.R., Parno, B., Rane, A., Setty, S.T.V., Thompson, L.: Vale: Verifying high-performance cryptographic assembly code. In: 26th USENIX Security Symposium, USENIX Security 2017, Vancouver, BC, Canada, August 16-18, 2017. (2017) 917–934
- [23] Bhargavan, K., Bond, B., Delignat-Lavaud, A., Fournet, C., Hawblitzel, C., Hritcu, C., Ishtiaq, S., Kohlweiss, M., Leino, R., Lorch, J.R., Maillard, K., Pan, J., Parno, B., Protzenko, J., Ramananandro, T., Rane, A., Rastogi, A., Swamy, N., Thompson, L., Wang, P., Béguelin, S.Z., Zinzindohoue, J.K.: Everest: Towards a verified, drop-in replacement of HTTPS. In Lerner, B.S., Bodík, R., Krishnamurthi, S., eds.: 2nd Summit on Advances in Programming Languages, SNAPL 2017, May 7-10, 2017, Asilomar, CA, USA. Volume 71 of LIPIcs., Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik (2017) 1:1–1:12
- [24] 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. PLDI 2017, New York, NY, USA, ACM (2017) 362–375
- [25] Chen, J., Feng, Y., Dillig, I.: Precise detection of side-channel vulnerabilities using quantitative cartesian hoare logic. In: Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security. CCS '17, New York, NY, USA, ACM (2017) 875–890
- [26] Terauchi, T., Aiken, A.: Secure information flow as a safety problem. In Hankin, C., Siveroni, I., eds.: Static Analysis, 12th International Symposium, SAS 2005, London, UK, September 7-9, 2005, Proceedings. Volume 3672 of Lecture Notes in Computer Science., Springer (2005) 352–367
- [27] Zaks, A., Pnueli, A.: Covac: Compiler validation by program analysis of the cross-product. In Cuéllar, J., Maibaum, T.S.E., Sere, K., eds.: FM 2008: Formal Methods, 15th International Symposium on Formal Methods, Turku, Finland, May 26-30, 2008, Proceedings. Volume 5014 of Lecture Notes in Computer Science., Springer (2008) 35–51
- [28] Barthe, G., Crespo, J.M., Kunz, C.: Relational verification using product programs. In Butler, M.J., Schulte, W., eds.: FM 2011: Formal Methods - 17th International Symposium on Formal Methods, Limerick, Ireland, June 20-24, 2011. Proceedings. Volume 6664 of Lecture Notes in Computer Science., Springer (2011) 200–214
- [29] Barthe, G., D’Argenio, P.R., Rezk, T.: Secure information flow by self-composition. *Mathematical Structures in Computer Science* **21**(6) (2011) 1207–1252
- [30] Pasareanu, C.S., Phan, Q.S., Malacaria, P.: Multi-run side-channel analysis using symbolic execution and max-SMT. In: Computer Security Foundations Symposium (CSF), 2016 IEEE 29th, IEEE (2016) 387–400

- [31] Crosby, S.A., Wallach, D.S., Riedi, R.H.: Opportunities and limits of remote timing attacks. *ACM Trans. Inf. Syst. Secur.* **12**(3) (January 2009) 17:1–17:29
- [32] Bleichenbacher, D.: Chosen ciphertext attacks against protocols based on the rsa encryption standard pkcs #1. In: *Proceedings of the 18th Annual International Cryptology Conference on Advances in Cryptology. CRYPTO '98*, London, UK, UK, Springer-Verlag (1998) 1–12
- [33] Brumley, D., Boneh, D.: Remote timing attacks are practical. In: *Proceedings of the 12th Conference on USENIX Security Symposium - Volume 12. SSYM'03*, Berkeley, CA, USA, USENIX Association (2003) 1–1
- [34] Rescorla, E., Dierks, T.: The Transport Layer Security (TLS) Protocol Version 1.2. RFC 5246 (August 2008)
- [35] Terauchi, T., Aiken, A.: Secure information flow as a safety problem. In: *SAS. Volume 3672.*, Springer (2005) 352–367
- [36] Rakamaric, Z., Emmi, M.: SMACK: decoupling source language details from verifier implementations. In Biere, A., Bloem, R., eds.: *Computer Aided Verification - 26th International Conference, CAV 2014, Held as Part of the Vienna Summer of Logic, VSL 2014, Vienna, Austria, July 18-22, 2014. Proceedings. Volume 8559 of Lecture Notes in Computer Science.*, Springer (2014) 106–113
- [37] llvm: clang: a C language family frontend for LLVM. <https://clang.llvm.org/> (2018)
- [38] Lattner, C., Adve, V.: LLVM: A Compilation Framework for Lifelong Program Analysis & Transformation. In: *Proceedings of the 2004 International Symposium on Code Generation and Optimization (CGO'04)*, Palo Alto, California (Mar 2004)
- [39] Schmidt, S.: s2n Is Now Handling 100 Percent of SSL Traffic for Amazon S3. <https://aws.amazon.com/blogs/security/s2n-is-now-handling-100-percent-of-of-ssl-traffic-for-amazon-s3/> (2017) [Online; accessed 15-Jan-2018].
- [40] Schmidt, S.: Introducing s2n, a New Open Source TLS Implementation. <https://aws.amazon.com/blogs/security/introducing-s2n-a-new-open-source-tls-implementation/> (2015) [Online; accessed 15-Jan-2018].
- [41] MacCarthaigh, C.: s2n and Lucky 13. <https://aws.amazon.com/blogs/security/s2n-and-lucky-13/> (2015) [Online; accessed 15-Jan-2018].
- [42] Almeida, J.B., Barbosa, M., Barthe, G., Dupressoir, F.: Verifiable side-channel security of cryptographic implementations: constant-time mce-cbc. In: *International Conference on Fast Software Encryption*, Springer (2016) 163–184
- [43] Brad Conte: Basic implementations of standard cryptography algorithms, like AES and SHA-1. <https://github.com/B-Con/crypto-algorithms> (2018) Commit: 02b66ec38b474445d10a5d1f0114bc0e8326707e.
- [44] Krawczyk, H., Bellare, M., Canetti, R.: HMAC: Keyed-Hashing for Message Authentication. RFC 2104, RFC Editor (February 1997) <http://www.rfc-editor.org/rfc/rfc2104.txt>.
- [45] Freier, A., Karlton, P., Kocher, P.: The Secure Sockets Layer (SSL) Protocol Version 3.0. RFC 6101, RFC Editor (August 2011) <http://www.rfc-editor.org/rfc/rfc6101.txt>.