

Katamaran: semi-automated verification of ISA specifications

Steven Keuchel

Vrije Universiteit Brussel, Belgium
steven.keuchel@vub.be

Georgy Lukyanov

Newcastle University, United Kingdom
g.lukyanov2@newcastle.ac.uk

Dominique Devriese

Vrije Universiteit Brussel, Belgium
dominique.devriese@vub.be

1 Introduction

An instruction set architecture (ISA) is an abstract specification of the syntax and semantics of machine code. It defines an envelope of allowed behaviour for CPU designers and a set of assumptions that software designers can rely on. Instead of informal prose and pseudo-code [e.g. 1], rigorous, executable formalisations of ISAs disambiguate the contract and improve testability and support modification, experimentation and formal study [e.g. 3, 12, 19]. Such formalisations are a crucial requirement for formal verification of both hardware [e.g. 8] and software [e.g. 15].

We are interested in verifying that critical safety guarantees of the ISA are upheld by the semantics of all instructions. Our long-term goal is to verify security guarantees offered by ISAs, specifically features like Intel SGX [16], virtual memory or capability machines [6]. We want to verify these properties in a form that can be used to reason about programs, as a way to ultimately verify security properties of real systems.

For achieving this, we take inspiration from recent formulations of capability safety in capability machines and high-level languages [10, 20, 22, 24]. Contrary to, for example, Nienhuis et al. [18], such techniques directly enable reasoning across encapsulation boundaries. These approaches use (essentially) a general purpose program logic, and formulate capability safety as a *universal contract* that automatically holds for arbitrary programs. The universal contract expresses guarantees provided by the machine and can be used for manually verifying trusted programs that interact with untrusted programs. We believe that this approach generalises well beyond capability safety.

However, proving such results about a language or ISA currently requires a lot of manual reasoning. For example, an in-progress CoQ formalisation of capability safety for a simple capability machine with 19 instructions requires about 17kLOC of proofs [11]. Real ISAs can be much larger, for example 30kLOC of SAIL specifications for ARMv8.3 [3]. Consequently, scaling up ISA property proofs raises important proof engineering challenges. For the verification effort to scale reasonably in terms of the size and complexity of the specification and for making it robust to changes, proof automation is a necessity. Uninteresting parts of the proof should be dealt with automatically, but at the same time, a human should be able to intervene and prove certain cases

manually or provide heuristics to steer the automation. In this text, we present KATAMARAN: a semi-automated verification tool for SAIL, intended to accommodate these requirements. The tool is a semi-automated separation logic verifier for SAIL in the tradition of SMALLFOOT [23], VERIFAST [13] and others. We intend to mechanically verify it against an abstract separation logic interface, making the tool similar to BEDROCK [7], HOLFOOT [23], VERISTAR+VERISMAIL [2, 21], VST-FLOYD [5] and the mechanisation of FEATHERWEIGHT VERIFAST [14].

In the remaining sections, we give a more detailed motivation of how we want to verify of ISA properties in KATAMARAN (Section 2) and a high-level overview of the design and current status of the tool and our next plans (Section 3).

Our developments are publicly available on Github [9].

2 Motivation

To understand how we intend KATAMARAN to semi-automatically prove ISA properties, consider the procedure

```
execute_store: (d s: reg) → unit
```

in Figure 1, which defines the semantics of a store instruction in a simple capability machine. The code is written in μ SAIL, a subset of SAIL that we briefly describe in Section 3. Ignoring the contracts in red and blue for now, it reads a capability from a destination register d , checks its bounds and permissions, reads a word from a source register s , and stores it to memory using an omitted `write_mem` procedure.

The lines in red state the universal contract for instructions on the machine. The contract is specified in an underlying separation logic, in terms of an abstract predicate `safe: cap+int` → \mathbb{P} , essentially stating that the instruction starts out with safe values in all registers and should also finish with safe register contents. Additionally, under the rules of the underlying separation logic and the predicate `safe`, this contract also implies that the instruction can only access memory reachable through the register contents and will respect any registered invariants of the system (see, for example, Skorstengaard et al. [20] for more details).

To verify the contract, KATAMARAN applies *symbolic execution* [4] and automatically computes the blue intermediate assertions by applying program logic rules. These intermediate assertions consist of pure logical facts (the “path condition”) and spatial separation logic assertions (the “symbolic heap”).

```

111 { * r ↦ wr * safe(cr) }
112   rereg
113   let: c : cap := call read_reg_cap d in
114   { ( * r ↦ wr * safe(cr) ) * ⌈c = wd ∧ c = cap(p, b, e, a)⌋ }
115   rereg
116   let: wa : bool := call write_allowed c.perm in
117   let: wb : bool := call within_bounds c in
118   assert (wa && wb) ;;
119   { ... c = cap(p, b, e, a) ∧ p ⊑ W ∧ b ≤ a < e⌈ }
120   let: w : cap + int := call read_reg_word s in
121   { ... c = cap(p, b, e, a) ∧ p ⊑ W ∧ b ≤ a < e ∧ w = ws⌈ }
122   call write_mem c.cursor w ;; call update_pc
123   { * r ↦ wr * safe(cr) }
124   rereg

```

Figure 1. Contract checking for `execute_store` `d s`

Function calls are symbolically executed by instantiating function contracts. For instance, `execute_store` uses the underlying function

```

125 { safe(cap(p, b, e, -)) * safe(w) * ⌈b ≤ a < e ∧ p ⊑ W⌋ }
126   write_mem a w
127 { safe(cap(p, b, e, -)) * safe(w) }

```

The precondition of a function contract usually needs to be massaged to match the symbolic heap and path condition when the call is made. This means instantiating quantified variables, applying logic rules – like the *frame rule* – and lemmas for both pure and impure assertions. KATAMARAN attempts to automatically perform this matching and otherwise delegates to user-provided heuristics that may tell it to apply specific lemmas for making the matching go through.

Like other tools, handling of straight-line code is fully automatic and branches are dealt with by exploring all paths separately. Easy pure assertions are solved immediately and others are collected as *verification conditions (VCs)*, to be proved separately (see Section 3).

We conjecture that KATAMARAN should be able to automatically discharge most uninteresting parts of ISA property proofs because typical SAIL specifications tend to contain a lot of mostly straight-line code, (almost) no loops, no higher-order functions and no dynamic memory allocation. Nevertheless, we plan to provide escape hatches for manually discharging the obligations that KATAMARAN cannot handle.

3 Overview and Status

Finally, we provide an overview of KATAMARAN’s workflow, depicted in Figure 2, and its implementation status.

Machine specifications As a specification language, KATAMARAN reuses SAIL which is translated to μ SAIL – a subset of SAIL deeply-embedded into COQ. We believe that all of SAIL is compilable to μ SAIL via transformations like monomorphisation and desugaring. μ SAIL’s features and a comparison to SAIL are documented in our code repository [9]. For the moment, the translation is still vaporware and we are writing examples as μ SAIL ASTs directly. One notable difference to SAIL’s existing COQ backend is the deep rather than shallow

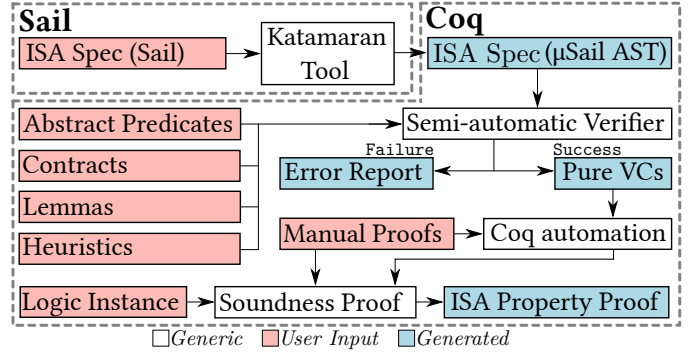


Figure 2. KATAMARAN Overview

embedding which allows for more parts of the verifier to be implemented in GALLINA rather than LTAC.

ISA property specifications ISA properties are specified by way of *contracts*, like the one in Figure 1, which may use predicates like `safe`. KATAMARAN treats predicates and also – like BEDROCK [7] VERiSTAR [21] – the assertion logic abstractly, relying on a generic logic interface that the user instantiates. This modularity allows the user to choose a separation logic rich enough to define all the predicates.

Notably absent from SAIL and μ SAIL are definitions of what constitutes memory of a specified machine. KATAMARAN leaves it up to the user to define a *memory model* as part of her logic instantiation and provide primitive access to it. For instance, the `write_mem` function can be defined and its contract proved sound externally in the logic, and given as inputs to KATAMARAN.

ISA property proofs KATAMARAN’s semi-automatic μ SAIL verifier will attempt to check the desired ISA property. The user may control the behaviour of the verifier through *lemmas*, for instance to fold/unfold recursive predicates, and *heuristics* that use lemmas to guide the spatial reasoning.

If successful, the verifier outputs a list of pure VCs for each possible execution path. If these can be proven (by existing COQ automation or, if necessary, manual COQ proofs), the contracts are provable in the underlying logic. If verification fails, KATAMARAN will terminate with an error message.

At the time of writing, the symbolic executor is the most mature component of KATAMARAN, although it is still missing a soundness proof and some features, like for instance the automatic application of lemmas via heuristics. Lemmas can be invoked manually though by a form of ad-hoc ghost statement in the syntax.

Our next plans for KATAMARAN are to develop it further and apply it to a number of increasingly challenging applications. This starts with simple examples, and ends with verifying capability safety in full-fledged capability machine ISAs. Intermediate goals notably include properties of artificial ISAs from the literature [e.g. 20] and a form of memory safety of Redfin [17] – a specialised ISA for aerospace applications designed with formal verification in mind.

References

- 221
222
223
224
225
226
227
228
229
230
231
232
233
234
235
236
237
238
239
240
241
242
243
244
245
246
247
248
249
250
251
252
253
254
255
256
257
258
259
260
261
262
263
264
265
266
267
268
269
270
271
272
273
274
275
- [1] AMD. 2019. AMD64 Architecture Programmer's Manual Volume 3: General-Purpose and System Instructions. <https://www.amd.com/system/files/TechDocs/24594.pdf>
- [2] Andrew W. Appel. 2011. VeriSmall: Verified Smallfoot Shape Analysis. In *Certified Programs and Proofs*. Springer Berlin Heidelberg, 231–246.
- [3] Alasdair Armstrong, Thomas Bauereiss, Brian Campbell, Alastair Reid, Kathryn E. Gray, Robert M. Norton, Prashanth Mundkur, Mark Was-sell, Jon French, Christopher Pulte, Shaked Flur, Ian Stark, Neel Krishnaswami, and Peter Sewell. 2019. ISA Semantics for ARMv8-a, RISC-v, and CHERI-MIPS. *Proceedings of the ACM on Programming Languages* 3, POPL (Jan. 2019), 71:1–71:31. <https://doi.org/10.1145/3290384>
- [4] Josh Berdine, Cristiano Calcagno, and Peter W. O'Hearn. 2005. Symbolic Execution with Separation Logic. In *Programming Languages and Systems*. Springer Berlin Heidelberg, 52–68.
- [5] Qinxiang Cao, Lennart Beringer, Samuel Gruetter, Josiah Dodds, and Andrew W. Appel. 2018. VST-Floyd: A Separation Logic Tool to Verify Correctness of C Programs. *Journal of Automated Reasoning* 61, 1 (June 2018), 367–422. <https://doi.org/10.1007/s10817-018-9457-5>
- [6] Nicholas P. Carter, Stephen W. Keckler, and William J. Dally. 1994. Hardware Support for Fast Capability-Based Addressing (*ASPLOS VI*). ACM. <https://doi.org/10.1145/195473.195579>
- [7] Adam Chlipala. 2011. Mostly-automated Verification of Low-level Programs in Computational Separation Logic. *SIGPLAN Not.* 46, 6 (June 2011), 234–245. <https://doi.org/10.1145/1993316.1993526>
- [8] Joonwon Choi, Muralidaran Vijayaraghavan, Benjamin Sherman, Adam Chlipala, and Arvind. 2017. Kami: A Platform for High-Level Parametric Hardware Specification and Its Modular Verification. *Proceedings of the ACM on Programming Languages* 1, ICFP (Aug. 2017), 24:1–24:30. <https://doi.org/10.1145/3110268>
- [9] Katamaran Developers. 2020. Katamaran Code Repository. <https://github.com/skeuchel/katamaran>.
- [10] Dominique Devriese, Lars Birkedal, and Frank Piessens. 2016. Reasoning about Object Capabilities Using Logical Relations and Effect Parametricity. In *European Symposium on Security and Privacy*. IEEE. <https://doi.org/10.1109/EuroSP.2016.22>
- [11] Aina Linn Georges, Alix Trieu, and Lars Birkedal. [n.d.]. Mechanized Reasoning about a Capability Machine. <https://cs.au.dk/~birke/papers/iris-capabilities-prisc-conf.pdf> Principles of Secure Compilation, 2020.
- [12] Shilpi Goel, Warren A. Hunt, and Matt Kaufmann. 2017. Engineering a Formal, Executable X86 ISA Simulator for Software Verification. In *Provably Correct Systems*. Springer International Publishing. https://doi.org/10.1007/978-3-319-48628-4_8
- [13] Bart Jacobs, Jan Smans, and Frank Piessens. 2010. A Quick Tour of the VeriFast Program Verifier. In *Programming Languages and Systems*. Lecture Notes in Computer Science, Vol. 6461. Springer Berlin Heidelberg, 304–311.
- [14] Bart Jacobs, Frédéric Vogels, and Frank Piessens. 2015. Featherweight VeriFast. *Logical Methods in Computer Science* Volume 11, Issue 3 (Sept. 2015). [https://doi.org/10.2168/LMCS-11\(3:19\)2015](https://doi.org/10.2168/LMCS-11(3:19)2015)
- [15] Xavier Leroy. 2009. Formal Verification of a Realistic Compiler. *Commun. ACM* 52, 7 (2009), 107–115.
- [16] Frank McKeen, Ilya Alexandrovich, Alex Berenzon, Carlos V. Rozas, Hisham Shafi, Vedvyas Shanbhogue, and Uday R. Savagaonkar. 2013. Innovative Instructions and Software Model for Isolated Execution (*HASP '13*). ACM. <https://doi.org/10.1145/2487726.2488368>
- [17] Andrey Mokhov, Georgy Lukyanov, and Jakob Lechner. 2019. Formal Verification of Spacecraft Control Programs (Experience Report) (*Haskell '19*). ACM. <https://doi.org/10.1145/3331545.3342593>
- [18] Kyndylan Nienhuis, Alexandre Joannou, Anthony Fox, Michael Roe, Thomas Bauereiss, Brian Campbell, Matthew Naylor, Robert M Norton, Simon W Moore, Peter G Neumann, et al. 2019. *Rigorous engineering for hardware security: formal modelling and proof in the CHERI design and implementation process*. Technical Report. University of Cambridge,
- Computer Laboratory. Accepted for publication at IEEE S&P 2020.
- [19] Alastair Reid, Rick Chen, Anastasios Deligiannis, David Gilday, David Hoyes, Will Keen, Ashan Pathirane, Owen Shepherd, Peter Vrabel, and Ali Zaidi. 2016. End-to-End Verification of Processors with ISA-Formal. In *Computer Aided Verification (Lecture Notes in Computer Science)*. Springer International Publishing, 42–58. https://doi.org/10.1007/978-3-319-41540-6_3
- [20] Lau Skorstengaard, Dominique Devriese, and Lars Birkedal. 2017. Reasoning About a Capability Machine with Local Capabilities - Provably Safe Stack and Return Pointer Management (without OS Support). (2017). <http://people.cs.kuleuven.be/dominique.devriese/wellbracketed-locally.pdf> In submission.
- [21] Gordon Stewart, Lennart Beringer, and Andrew W. Appel. 2012. Verified Heap Theorem Prover by Paramodulation (*ICFP '12*). ACM. <https://doi.org/10.1145/2364527.2364531>
- [22] David Swasey, Deepak Garg, and Derek Dreyer. 2017. Robust and Compositional Verification of Object Capability Patterns. In *OOPSLA*. ACM. <https://people.mpi-sws.org/~swasey/papers/ocpl/ocpl-20170418.pdf>
- [23] Thomas Tuerk. 2009. A Formalisation of Smallfoot in HOL. In *Theorem Proving in Higher Order Logics*. Springer Berlin Heidelberg, 469–484.
- [24] Thomas Van Strydonck, Frank Piessens, and Dominique Devriese. 2019. Linear Capabilities for Fully Abstract Compilation of Separation-Logic-Verified Code. *Proc. ACM Program. Lang.* ICFP (2019). accepted.
- 276
277
278
279
280
281
282
283
284
285
286
287
288
289
290
291
292
293
294
295
296
297
298
299
300
301
302
303
304
305
306
307
308
309
310
311
312
313
314
315
316
317
318
319
320
321
322
323
324
325
326
327
328
329
330