

1 The State of Sail

2 **Alasdair Armstrong**

3 Department of Computer Science and Technology, University of Cambridge, UK

4 **Thomas Bauereiss**

5 Department of Computer Science and Technology, University of Cambridge, UK

6 **Brian Campbell**

7 School of Informatics, University of Edinburgh, UK

8 **Alastair Reid**

9 ARM Ltd., Cambridge, UK

10 **Kathryn E. Gray**

11 Department of Computer Science and Technology, University of Cambridge (Formerly), UK

12 **Robert M. Norton**

13 Department of Computer Science and Technology, University of Cambridge, UK

14 **Prashanth Mundkur**

15 SRI International, Menlo Park, US

16 **Mark Wassell**

17 Department of Computer Science and Technology, University of Cambridge, UK

18 **Jon French**

19 Department of Computer Science and Technology, University of Cambridge, UK

20 **Christopher Pulte**

21 Department of Computer Science and Technology, University of Cambridge, UK

22 **Shaked Flur**

23 Department of Computer Science and Technology, University of Cambridge, UK

24 **Ian Stark**

25 School of Informatics, University of Edinburgh, UK

26 **Neel Krishnaswami**

27 Department of Computer Science and Technology, University of Cambridge, UK

28 **Peter Sewell**

29 Department of Computer Science and Technology, University of Cambridge, UK

30 — **Abstract** —

31 Sail is a custom domain-specific language for ISA semantics, in which we have developed formal
32 models for ARMv8-A, RISC-V, and MIPS, as well as CHERI-based capability extensions for
33 both RISC-V and MIPS. In particular, our model of ARMv8-A is automatically translated from
34 ARM-internal definitions and tested against the ARM Architecture validation suite. All the above
35 models contain enough system-level features to boot various operating systems, including Linux and
36 FreeBSD, but also various smaller microkernels and hypervisors.

37 In this short paper, we present the ways in which Sail enables us to bridge the gap between our
38 various ISA models and the myriad use cases for such models. By using Sail, we are able to generate
39 emulators for testing and validation, generate theorem prover definitions across multiple major tools
40 (Isabelle, HOL4, and Coq), translate Sail to SMT for automatic verification, and integrate with both
41 operational models for relaxed-memory concurrency via our RMEM tool.

42 We will also present our current work to extend Sail to support axiomatic concurrency models,
43 in the style of Alglave and Maranget's herd7 tool, with the intent being to explore the behaviour of
44 concurrent litmus tests that span the full behaviour of the architecture. As an illustrative example,
45 one could consider how instruction cache maintenance instructions interact with self-modifying code
46 in an axiomatic setting, or other interesting cases that are not well-covered by existing tools.



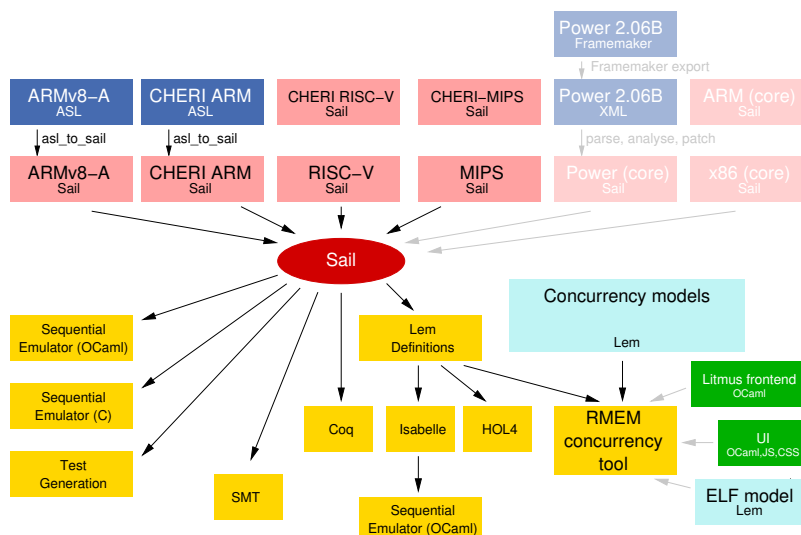
47 **2012 ACM Subject Classification** Theory of computation → Semantics and reasoning; Computer
 48 systems organization → Architectures; Software and its engineering → Assembly languages

49 **Keywords and phrases** Instruction Set Architectures, Semantics, Theorem Proving

50 **Digital Object Identifier** 10.4230/LIPIcs...

51 **1 Overview**

52 Sail is a custom pseudocode-like language for specifying the semantics of instruction set
 53 architectures (ISAs). It is a first-order imperative language with a semantics that is as
 54 straightforward as possible. We have aimed to strike a balance in designing a language that
 55 is both expressive enough to idiomatically express multiple instruction set architectures,
 56 which simultaneously being as inexpensive as possible to allow translation to each desired
 57 target. This balance is partially achieved with a type system that allows dependent types
 58 for bitvector widths and integer ranges, the typing information from which can be exploited
 59 by various generic rewrites and backend-specific optimisations e.g. for monomorphisation.
 60 Figure 1 gives an overview of our currently supported instruction set architectures and target
 61 uses, updated with changes since our previous paper [6].



■ **Figure 1** Sail ISA semantics (top) and target use cases (bottom). The greyed out ISAs are from previous work we are not actively working on

62 This short paper describes the current state of each of our Sail models, and describes
 63 ongoing work to enhance Sail with automatic verification and axiomatic concurrency support
 64 via a translation from Sail into SMTLIB definitions for the Z3 and CVC4 SMT solvers. The
 65 below table summarises the state of most of our models including our CHERI extensions.

	source	KLoS	provers	boots
ARMv8.5-A	ASL	125	Isa, HOL4, Coq*	Linux, Hafnium
MIPS	hand	2	Isa, HOL4, Coq	FreeBSD
CHERI MIPS	hand	+2	Isa, HOL4, Coq	FreeBSD, CheriBSD
RISC-V	hand	5	Isa, HOL4, Coq	Linux, FreeBSD, FreeRTOS, Hafnium
CHERI RISC-V	hand	+2	Isa, HOL4, Coq	
CHERI ARM	ASL		Isa	

67 **ARMv8.5-A** Our ARMv8-A model is translated from ARM’s internal architecture speci-
 68 fication language ASL. Our translation has been validated by running our translation against
 69 the ARM internal architecture validation suite, as previously discussed in [6]. Recently we
 70 have been improving Sail’s translation into Coq. We have continued to work on emulation
 71 performance for ARM, which was previously slower than our other models due to the size
 72 of the specification, and it now boots Linux in just under 2 minutes (on a Ryzen 5 2600X),
 73 corresponding to approximately 200 000 instructions per second, roughly a four-fold improve-
 74 ment over [6]. We have also been working on developing infrastructure for formally proving
 75 properties of a CHERI ARM specification, but this work is in early stages.

76 **(CHERI) RISC-V** We have extended our RISC-V model with CHERI capability support.
 77 We have ensured our RISC-V model is extensible, so the CHERI extension (and other
 78 extensions) are able to exist as a separate repository which builds upon the base model. We
 79 have further validated the RISC-V spec by running FreeRTOS and a port of the Hafnium
 80 hypervisor atop the RISC-V model, in addition to Linux, FreeBSD, and seL4.

81 **(CHERI) MIPS** Our CHERI-MIPS model continues to be extended with new CHERI
 82 instructions, and is now an official part of the CHERI ISAv7 [10] architecture manual.

83 2 Automatic property verification with Sail-SMT

84 In addition to translating Sail to interactive theorem provers, we have more recently im-
 85 plemented a translation from Sail into SMT. This enables QuickCheck-like properties to
 86 be stated and verified in Sail itself (provided they are free of loops, in which case they are
 87 checked up to some iteration bound). For example, Figure 2 shows a property from our
 88 CHERI RISC-V specification, which verifies that if `setCapBounds` claims to have set `c`’s bounds
 89 to base and top exactly, then `getCapBounds` will return the same bounds as were set. While this
 90 property seems simple, capability bounds are stored using a fairly intricate floating-point-like
 91 compressed format, and there are additional subtle edge cases at the top and bottom of the
 92 address space. Our SMT translation was able to discover bugs in our implementations of
 93 such capability manipulation functions which had not been found via random testing.

94 A major advantage we have found in this style of lightweight verification with SMT
 95 solvers is that it can be used by hardware-designers developing ISA extensions who have
 96 no experience with interactive theorem proving tools. Another use for hardware-designers
 97 is to write a Sail version of a function that closely mimics a Bluespec (or other HDL)
 98 implementation that is complicated due to e.g. timing requirements, and automatically prove
 99 it equivalent to a simple Sail implementation.

```
function set_bounds_exact(c : Capability, base : bits(64), top : bits(65)) -> bool = {
  let (exact, c') = setCapBounds(c, base, top);
  let (base', top') = getCapBounds(c');
  ~(exact) | (unsigned(base) >= unsigned(top))
  | (base' == unsigned(base) & top' == unsigned(top))
}
```

■ **Figure 2** An automatically verified property from CHERI RISC-V

100 The basic approach is similar to that used by existing model-checking tools such as
 101 CBMC [1], and the approach used for ARM’s ASL language [9]. We first translate the Sail
 102 source into an intermediate representation (IR) which is shared by the C backend, this is
 103 then converted into a SSA based control-flow graph, which is then turned into a sequence of
 104 SMTLIB definitions which can be used with either Z3 or CVC4.

105 **3** Axiomatic relaxed-memory concurrency with Sail

106 Previous work on concurrent behaviours of instruction set architectures using Sail was based
 107 on our RMEM tool [8] which provides operational-semantics for various memory models.
 108 However, many architectures, such as RISC-V specify their memory model in an axiomatic-
 109 style, where the memory model is described in terms of axioms that restrict the set of possible
 110 candidate executions. Alglave et al’s diy7 [3, 4] tool suite, in particular the herd7 [5] tool,
 111 already provides a framework for evaluating the relaxed-memory behaviour for small assembly
 112 programs (litmus tests) over several architectures, using a language called *cat* [2]. However
 113 the ISA semantics used by herd is hard-coded in OCaml for each supported architecture
 114 within the tool, plus additional architecture specific infrastructure for e.g. assembly parsing.

115 By combining our Sail to SMT translation with the existing infrastructure for litmus
 116 tests and *cat* files provided by the diy7 tools, we aim to produce a tool similar to herd7,
 117 except using the Sail instruction semantics and assembly parsing infrastructure (which can
 118 also be specified within Sail). This would give us a architecture-agnostic tool that can
 119 combine an arbitrary memory model specified in *cat*, with an ISA specified in Sail. While
 120 our implementation is still very experimental, initial results are promising, and prior work
 121 such as Lau et al’s Cerberus-BMC [7] for C11 concurrency demonstrate that the use of a
 122 SMT solver in this area is practicable.

123 — References —

- 124 **1** CBMC: Bounded Model Checking for Software, 2017. <http://www.cprover.org/cbmc/>.
- 125 **2** Jade Alglave, Patrick Cousot, and Luc Maranget. Syntax and semantics of the weak consistency
 126 model specification language *cat*. *CoRR*, abs/1608.07531, 2016. URL: [http://arxiv.org/abs/
 127 1608.07531](http://arxiv.org/abs/1608.07531), arXiv:1608.07531.
- 128 **3** Jade Alglave and Luc Maranget. The diy7 tool. <http://diy.inria.fr/>, 2017.
- 129 **4** Jade Alglave, Luc Maranget, Susmit Sarkar, and Peter Sewell. Fences in weak memory
 130 models. In *Proceedings of CAV 2010: the 22nd International Conference on Computer Aided
 131 Verification, LNCS 6174*, 2010. doi:10.1007/978-3-642-14295-6_25.
- 132 **5** Jade Alglave, Luc Maranget, and Michael Tautschnig. Herding Cats: Modelling, Simulation,
 133 Testing, and Data Mining for Weak Memory. *ACM TOPLAS*, 36(2):7:1–7:74, July 2014.
- 134 **6** Alasdair Armstrong, Thomas Bauereiss, Brian Campbell, Alastair Reid, Kathryn E. Gray,
 135 Robert M. Norton, Prashanth Mundkur, Mark Wassell, Jon French, Christopher Pulte, Shaked
 136 Flur, Ian Stark, Neel Krishnaswami, and Peter Sewell. ISA semantics for armv8-a, risc-v, and
 137 CHERI-MIPS. *PACMPL*, 3(POPL):71:1–71:31, 2019. doi:10.1145/3290384.
- 138 **7** Stella Lau, Victor B. F. Gomes, Kayvan Memarian, Jean Pichon-Pharabod, and Peter Sewell.
 139 Cerberus-BMC: a principled reference semantics and exploration tool for concurrent and
 140 sequential C. In *CAV 2019*, July 2019. (to appear).
- 141 **8** Christopher Pulte, Shaked Flur, Will Deacon, Jon French, Susmit Sarkar, and Peter Sewell.
 142 Simplifying ARM Concurrency: Multicopy-atomic Axiomatic and Operational Models for
 143 ARMv8. In *POPL 2018*, July 2018. doi:10.1145/3158107.
- 144 **9** Alastair Reid. Who guards the guards? formal validation of the arm v8-m architecture
 145 specification. *Proc. ACM Program. Lang.*, 1(OOPSLA):88:1–88:24, October 2017. doi:
 146 10.1145/3133912.
- 147 **10** Robert N. M. Watson, Peter G. Neumann, Jonathan Woodruff, Michael Roe, Hesham Almatary,
 148 Jonathan Anderson, John Baldwin, David Chisnall, Brooks Davis, Nathaniel Wesley Filardo,
 149 Alexandre Joannou, Ben Laurie, A. Theodore Markettos, Simon W. Moore, Steven J. Murdoch,
 150 Kyndylan Nienhuis, Robert Norton, Alex Richardson, Peter Rugg, Peter Sewell, Stacey Son,
 151 and Hongyan Xia. Capability Hardware Enhanced RISC Instructions: CHERI instruction-set
 152 architecture (version 7). Technical report, Computer Laboratory, June 2019.