

ISA Semantics for ARMv8-A, RISC-V, and CHERI-MIPS

ANONYMOUS AUTHOR(S)

Architecture specifications notionally define the fundamental interface between hardware and software: the envelope of allowed behaviour for processor implementations, and the basic assumptions for software development and verification. But in practice, they are typically prose and pseudocode documents, not rigorous or executable artifacts, leaving software and verification on shaky ground.

In this paper, we present rigorous semantic models for the sequential behaviour of large parts of the mainstream ARMv8-A, RISC-V, and MIPS architectures, and the research CHERI-MIPS architecture, that are complete enough to boot operating systems, variously Linux, FreeBSD, or seL4. Our ARMv8-A models are automatically translated from authoritative ARM-internal definitions, and (in one variant) tested against the ARM Architecture Validation Suite.

We do this using a custom language for ISA semantics, Sail, with a lightweight dependent type system, that supports automatic generation of emulator code in C and OCaml, and automatic generation of proof-assistant definitions for Isabelle, HOL4, and (currently only for MIPS) Coq. We use the former for validation, and to assess specification coverage. To demonstrate the usability of the latter, we prove (in Isabelle) correctness of a purely functional characterisation of ARMv8-A address translation. We moreover integrate the RISC-V model into the RMEM tool for (user-mode) relaxed-memory concurrency exploration. We prove (on paper) the soundness of the core Sail type system.

We thereby take a big step towards making the architectural abstraction actually well-defined, establishing foundations for verification and reasoning.

1 INTRODUCTION

The architectural abstraction is a fundamental interface in computing: the architecture specification for each family of processors, ARMv8-A, AMD64, IBM POWER, Intel 64, MIPS, RISC-V, SPARC, etc., notionally defines the envelope of allowed behaviour for all hardware processor implementations of that family, providing the basic assumptions for portable software development. This decouples hardware and software implementation, as architectures are relatively stable over time, while processor implementations evolve rapidly.

In practice, industry architecture specifications have traditionally been prose documents, with decoding tables and (at best) pseudocode descriptions of instruction behaviour, while vendors have maintained internal “golden” reference models, often as large and highly confidential C++ codebases. The mainstream architectures have accumulated enormous complexity: 6300 and 4700 pages for recent ARMv8-A and Intel 64/IA-32 specification documents [ARM 2017; Intel Corporation 2017]. They comprise two main parts: the Instruction Set Architecture (ISA), describing the behaviour of each instruction in isolation, and cross-cutting aspects such as the concurrency model and interrupt behaviour. Understanding all these details is essential for achieving correct and robust behaviour of computer systems, but prose and pseudocode are simply not up to the task of precisely specifying them. These specification documents are moreover not executable as test oracles—they do not allow one to compute the set of all architecturally allowed behaviour of hardware tests, or to test software above the entire architectural envelope rather than just some specific implementation—and they do not support automatic test generation or test-suite specification coverage measurement.

Meanwhile, academic researchers in programming languages, semantics, analysis, and verification have increasingly aimed at mechanised reasoning about correctness down to the machine level, e.g. in the CakeML [Fox et al. 2017; Kumar et al. 2014; Tan et al. 2016], CerCo [Amadio et al. 2013], CompCert [Leroy 2009; Leroy et al. 2017], and CompCertTSO [Ševčík et al. 2013] verified compilers; the seL4 [Fox and Myreen 2010; Klein et al. 2014] and Hyper-V [Leinenbach and Santen 2009] verified hypervisors; the Verified Software Toolchain [Appel et al. 2017]; CertiKOS verified OS [Gu

et al. 2016]; Verasco verified static analysis [Jourdan et al. 2015]; RockSalt software fault isolation system [Morrisett et al. 2012]; Bedrock [Chlipala 2013]; PROSPER [Baumann et al. 2016; Guanciale et al. 2016]; machine-code program logics [Jensen et al. 2013; Kennedy et al. 2013; Myreen 2009]; and relaxed-memory semantics [Alglave et al. 2010, 2014; Flur et al. 2017; Gray et al. 2015; Pulte et al. 2018; Sarkar et al. 2011; Sewell et al. 2010]. Binary analysis tools such as Angr [Shoshitaishvili et al. 2016], BAP [Brumley et al. 2011], TSL [Lim and Reps 2013], and Valgrind [Nethercote and Seward 2007] also need architectural models, although typically less formally expressed.

On what semantics should such work be based? Recoiling, reasonably enough, from the scale of the full 6000+ page vendor architecture documents, and from the poorly specified complexities of the concurrency models and privileged “system-mode” aspects of the architectures (virtual memory, exceptions, interrupts, security domain transitions, etc.), many groups have hand-written formal models of modest ISA fragments. These typically cover just enough of the instruction set, and in just enough detail, for their purpose: usually only some aspects of the sequential behaviour of parts of the non-privileged “user-mode” ISA, and just for one proof assistant (Coq, HOL4, or Isabelle). Some are validated against actual hardware behaviour, to varying degrees, but none are tied to a vendor reference model. The multiplicity of models, each produced by a different group for their specific purpose, is inefficient and makes it hard to amortise any validation investment. A few go beyond user-mode fragments, including seL4, PROSPER, and the ACL2 X86isa model [Goel et al. 2017]; we return to these, and other related work, in §9. Emulators such as QEMU [qem 2017] and gem5 [gem 2017] effectively also develop models, often rather complete, but these are optimised for performance and hard to use for other purposes.

In this paper, we present rigorous semantic models for the sequential behaviour of large parts of the mainstream ARMv8-A, RISC-V, and MIPS architectures, and the research CHERI-MIPS architecture, that are complete enough to boot various operating systems: Linux above the ARMv8-A model, FreeBSD above MIPS and CHERI-MIPS, and seL4 and Linux above RISC-V. These are rather large semantics by usual academic standards: approximately 23 000 lines for ARMv8-A, and a few thousand for each of the others.

ARMv8-A is the ARM application-processor architecture, specifying the processors, designed by ARM and by their architecture partners, and produced by many vendors, that are ubiquitous in mobile devices. We build on a shift within ARM over recent years to specify ISA behaviour in an ARM-internal machine-processed language, ASL. We work with two versions: a recent public release of large parts of this for ARMv8.3 [Reid 2016, 2017; Reid et al. 2016], and a currently non-public more complete version thereof; our ARM models are automatically translated from these. We moreover validate the second by testing against the ARM-internal Architecture Validation Suite. These are thus substantially more complete, authoritative, and well-validated than previous models. For RISC-V and CHERI-MIPS, the situation is rather different: these are much simpler architectures, and they are in flux, currently being designed. Our models for these (and our MIPS model underlying CHERI-MIPS) are handwritten, feeding back into the architecture design process, and validated in part by comparison with previous simulator and formal models.

To be generally useful, our models should simultaneously:

- (1) be accessible to practising engineers who use existing vendor pseudocode descriptions;
- (2) be automatically translatable into executable sequential emulator code, with reasonable performance, to support validation of the models and software development above them;
- (3) be automatically translatable into idiomatic theorem prover definitions, to support formal mechanised reasoning about the architectures and about code above them — ideally for all the major provers, to enable use by each prover community;
- (4) provide bidirectional mappings between assembly syntax and binary opcodes;

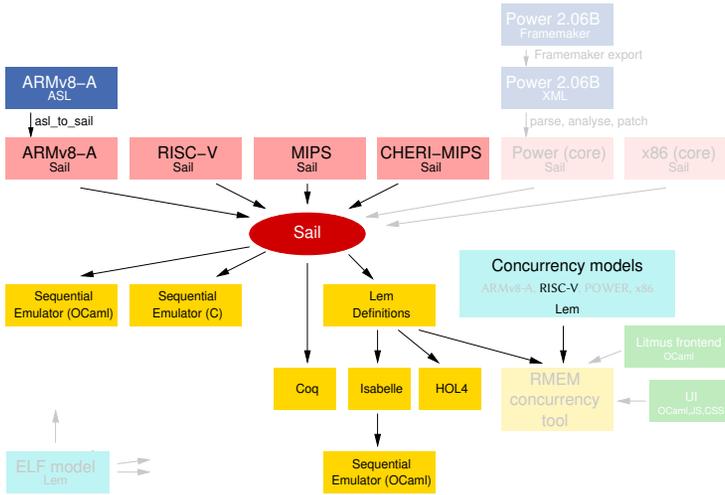


Fig. 1. Sail ISA semantics and (in yellow) the generated prover and emulator versions. The grey parts are previous concurrency and ISA models, user-mode only and not yet fully integrated into current Sail

- (5) provide the fine-grained execution information needed to integrate ISA semantics with the (user-mode) architectural relaxed-memory concurrency semantics previously developed;
- (6) be well-validated, to give confidence that they do capture the architectural intent and soundly describe hardware behaviour; and
- (7) be expressed in a well-engineered and robust infrastructure.

We achieve all this by designing a custom language for ISA semantics, Sail (§3), together with automatic translations as shown in Fig. 1: from the ARM-internal ASL language into Sail, from Sail to C and OCaml emulator code (§5), from Sail to Isabelle/HOL, HOL4, and Coq theorem-prover definitions (§4), and (currently only for RISC-V) from Sail to a representation used by the RMEM concurrency exploration tool [Pulte et al. 2018]. A common infrastructure for all our architecture models saves much duplication of effort.

Reconciling the above disparate goals is a delicate language-design problem. Sail has to be expressive enough to support each model idiomatically, especially the most-demanding ARMv8-A case, where the ASL source has accumulated features over time, including exceptions and complex (but not fully checked) dependent types for bitvector lengths. But to make Sail translatable into all the provers, especially the non-dependently-typed Isabelle/HOL and HOL4, and to fast C and OCaml code for emulation, it should be as *inexpressive* as possible, as we have to translate away any features that are not in the target language. We resolve this with a carefully designed lightweight dependent type system for checking vector bounds and integer ranges, inspired by Liquid types [Rondon et al. 2008], but which can be formalised in a simple, syntax-directed and single-pass style using a bidirectional approach [Dunfield and Krishnaswami 2013]. All constraints can be shown to exist within a decidable fragment, and are resolved using the Z3 SMT solver [De Moura and Bjørner 2008]. Our translations to Isabelle/HOL, HOL4, C, and OCaml rely on *monomorphising* these dependent types where they are not target-expressible, allowing us to use the existing well-developed machine word libraries for the first two, and efficient representations for the last two.

Otherwise, Sail is essentially a first-order functional/imperative language with a simple effect system, but with abstract register and memory accesses, for sequential and concurrent interpretations. Higher-order functions are unnecessary for our ISA models and would complicate the

translations to efficient C emulator code. Sail builds on earlier work [Flur et al. 2016; Gray et al. 2015] but has been substantially redesigned, especially the type system; the earlier work handled only modest user-mode ISA fragments for concurrency models, without the translation from ASL, prover definitions, fast emulation, or rather complete models we report on. We increase confidence in the Sail type system with a (paper) formalisation and soundness proof of a core MiniSail (§6).

We validate our models with the OS boots and ARM Architectural Validation Suite mentioned above, and with other test suites, using the executable OCaml and C versions produced by Sail (§7). This also lets us assess the *specification coverage* of such OS boot executions and test suites. We also validate the RISC-V model behaviour on concurrency litmus tests using RMEM.

We evaluate the usability of our generated theorem prover definitions by conducting an example proof in Isabelle/HOL about one of the most complex parts of the ARMv8-A specification, the translation from virtual to physical memory addresses. We prove correctness of a simple purely functional characterisation of address translation, under suitable preconditions (§8).

Considered as a specification or programming language, Sail is unusual in that it aims to support just a handful of specific programs – these and other architecture definitions of mainstream and research architectures – but the importance of those makes it necessary to do so well, and the specification scale and multiple demands listed above make that challenging.

Sail, along with our public ARMv8-A, RISC-V, MIPS, and CHERI-MIPS models, is publicly available under an open-source licence (the non-anonymous supplementary material has a github link), and with an OPAM package for Sail. The version of our ARMv8-A model derived from a non-public ARM source is currently not available, but we hope that will be possible in due course, and some of the legal infrastructure needed is in place. The anonymous supplementary material contains anonymised versions of the public models, the generated theorem-prover versions of them, our Isabelle proof scripts for §8, and our MiniSail definitions and paper proofs for §6.

Caveats and limitations. Our models cover considerably more than most formal ISA semantics of previous work, but they are still far from complete definitions of these architectures. For ARMv8-A, we translate only the AArch64 64-bit part of the architecture, not the AArch32 32-bit instructions. Including these should need only modest additional work. Our Coq generation has so far only been exercised for MIPS. Our assembly syntax support has only been exercised for RISC-V; for ARM it should be possible to generate this from ARM-supplied metadata, but that has not yet been done.

More substantially, we focus here on sequential behaviour. For RISC-V, our ISA model is integrated with the corresponding user-mode relaxed memory model, but we have not yet done that for ARM, and the relaxed-memory semantics of systems features (virtual memory, interrupts, etc.) is an open problem. Previous versions of Sail included models for modest fragments of the user-mode ISAs of IBM POWER [Gray et al. 2015], ARMv8 [Flur et al. 2016], and RISC-V and x86 (both previously unpublished); sufficient only for litmus tests and some user-mode concurrent algorithms. Those IBM POWER and x86 models have not yet been ported to the revised Sail of this paper, and that ARM model will be superseded by the one we present here when the above integration is done.

2 MODELS

The current status of our models and the generated definitions is summarised below.

Architecture	Source	Size (LOS)	Boots	Generates		
ARMv8.3-A (public)	ARM ASL	23 000		C, OCaml	Isabelle, HOL4	RMEM
ARMv8.3-A (private)	ARM ASL	30 000	Linux	C, OCaml		
RISC-V	hand-written	5 000	seL4, Linux	OCaml	Isabelle, HOL4	
MIPS	hand-written	2 000	FreeBSD	C, OCaml	Isabelle, HOL4, Coq	
CHERI-MIPS	hand-written	4 000	FreeBSD	C, OCaml	Isabelle, HOL4	

```

197 union clause ast = LOAD : (bits(12), regbits, regbits)
198
199 mapping clause encdec = LOAD(imm, rs1, rd, is_unsigned, size, false, false)
200   <-> imm @ rs1 @ bool_bits(is_unsigned) @ size_bits(size) @ rd @ 0b0000011
201
202 function clause execute(LOAD(imm, rs1, rd, is_unsigned, width, aq, rl)) =
203   let vaddr : xlenbits = X(rs1) + EXTS(imm) in
204   if check_misaligned(vaddr, width)
205   then { handle_mem_exception(vaddr, E_Load_Addr_Align); false }
206   else match translateAddr(vaddr, Read, Data) {
207     TR_Failure(e) => { handle_mem_exception(vaddr, e); false },
208     TR_Address(addr) =>
209       match width {
210         BYTE => process_load(rd, vaddr, mem_read(addr, 1, aq, rl, false), is_unsigned),
211         HALF => process_load(rd, vaddr, mem_read(addr, 2, aq, rl, false), is_unsigned),
212         WORD => process_load(rd, vaddr, mem_read(addr, 4, aq, rl, false), is_unsigned),
213         DOUBLE => process_load(rd, vaddr, mem_read(addr, 8, aq, rl, false), is_unsigned)
214       }
215   }

```

Fig. 2. RISC-V load instruction in Sail

2.1 RISC-V

Most ISAs have been proprietary. In contrast, RISC-V is an open ISA, currently under development by a broad industrial and academic community, coordinated by the RISC-V Foundation. It is subdivided into a core and many separable features. We have handwritten a RISC-V ISA model based on recent versions of the prose RISC-V specifications [RIS 2017]. Our current model implements the 64-bit (RV64) version of the ISA: the *rv64imac* dialect (integer, multiply-divide, atomic, and compressed instructions), with user, machine, and supervisor modes, and the Sv39 address translation mode (3-level page tables covering 512GiB of virtual address space).

The model is partitioned into separate files for user-space definitions, machine- and supervisor-mode parts, the physical memory interface, virtual memory and address translation, instruction definitions, and the fetch-execute-interrupt loop. The main omissions are floating-point, PMP (Physical Memory Protection), modularisation for the “unified” 32-bit/64-bit model, and factoring to build machine/user and machine-only variants.

For example, Fig. 2 shows the Sail code defining the RISC-V LOAD instructions: a constructor of the *ast* Sail type, a clause of the *encdec* function (mapping between a 32-bit instruction word and the corresponding *ast* value containing the opcode fields), and a clause of the *execute* function expressing its dynamic semantics. The body of that is imperative code: *X(...)* refers to the RISC-V general-purpose registers, *mem_read* is a function that performs a read of physical memory, and *process_load* handles potential access exceptions. The boolean return value of the *execute* clause indicates whether the instruction retired successfully, and is used to update the *minstret* CSR register. The *aq* and *rl* flags are used to indicate the ordering constraints of the load to the memory model. Modulo minor syntactic variations, this should be readable by anyone familiar with typical industry ISA pseudocode descriptions.

To get a sense of what is required to make an ISA semantics complete enough to boot an OS, rather than a user-mode fragment, we describe some of what we have had to do. This model is parameterisable over various platform implementation choices that the ISA allows. In particular, it supports (i) trapping as well as non-trapping modes of accesses to misaligned data addresses, and

```

246 function clause execute (CIncoffsetImmediate(cd, cb, imm)) = {
247   checkCP2usable();
248   let cb_val = readCapReg(cb);
249   let imm64 : bits(64) = sign_extend(imm);
250   if register_inaccessible(cd) then
251     raise_c2_exception(CapEx_AccessSystemRegsViolation, cd)
252   else if register_inaccessible(cb) then
253     raise_c2_exception(CapEx_AccessSystemRegsViolation, cb)
254   else if (cb_val.tag) & (cb_val.sealed) then
255     raise_c2_exception(CapEx_SealViolation, cb)
256   else
257     let (success, newCap) = incCapOffset(cb_val, imm64) in
258     if success then
259       writeCapReg(cd, newCap)
260     else
261       writeCapReg(cd, int_to_cap(to_bits(64, getCapBase(cb_val)) + imm64))
262   }

```

Fig. 3. CHERI-MIPS capability increment-offset instruction in Sail

(ii) write updates as well as traps when a dirty-bit needs to be updated in a page-table entry during address translation. RISC-V also specifies various control and status registers (CSRs) as having bitfields with platform-defined behaviour on reads and writes, which allows a platform to choose legal values of a CSR bitfield, and how it handles writes to those fields. Our model supports these choices through user-specifiable *legaliser* functions that intercept read and write accesses to those CSRs that require such behaviour.

We have also endeavoured to keep other platform aspects explicitly separate from the Sail model. For example, the reservation state for Load-Reserved/Store-Conditional instructions is kept as part of the platform state, since the reservation state and progress guarantees provided are inherently platform-specific. This separation also simplifies reasoning about the RISC-V memory model.

The physical memory map for a platform is specified using the `extern` facility of the Sail language, which enables the ISA model itself to remain agnostic of the actual map, but allows the contexts of the various backend renderings of the model to provide these definitions. For example, the generated OCaml executable model is linked against modules that define the locations of valid physical memory regions, valid memory-mapped I/O regions, and the location of the timer and terminal devices. These modules also place the corresponding Device-Tree information generated from these values at the expected location in physical memory when the OCaml ISA emulator is initialised. The ISA model itself checks any physical address used for a data or instruction access against these before allowing the access or generating the appropriate memory fault exception.

Although not strictly part of the ISA specification, we have also implemented some aspects of simple memory-mapped devices in Sail (timer, terminal, device interrupt routing) as an exploration of the use of the Sail language to describe other components of a complete platform model.

Our development of the Sail model has led us to contribute improvements in the RISC-V prose specifications, e.g. in the description of page-faults expected during page-table walks, and fixes to bugs in the corresponding address translation code of the widely-used Spike reference simulator. It has also pointed out ambiguities in the specification of interrupt delegation, and cases of missing reservation yields in Spike.

2.2 MIPS and CHERI-MIPS

CHERI-MIPS [Watson et al. 2017, 2015; Woodruff et al. 2014] is an experimental research architecture that extends 64-bit MIPS with support for fine-grained memory protection and secure compartmentalisation. It provides hardware *capabilities*, compressed 128-bit values including a base virtual address, an offset, a bound, and permissions; and object capabilities that link code and data pointers. Additional tag memory, cleared by any non-capability writes, records whether each capability-sized and aligned unit of memory holds a valid capability. This and other features makes them unforgeable by software: each capability must be derived from a more-permissive one. One can either use capabilities in place of all pointers (“pure capability” code) or selectively (“hybrid”).

CHERI has used executable formal models of the architecture as a central design tool since 2014, largely in L3 [Fox and Myreen 2010], coupled with traditional prose and non-formal pseudocode in the ISA specification document. Executability of the formal model (at some 100s of KIPS) has been vital, both to provide a reference to test hardware implementations against, and as a platform for software development that is automatically in step with the frequent architecture changes. Isabelle definitions generated from L3 have been used for proofs about compressed capabilities and of security properties of the architecture as a whole. This has all provided invaluable experience for the design of Sail, and our Sail CHERI-MIPS model is now mature enough to replace both the earlier L3 model and the non-formal pseudocode; the latter using Sail-generated LaTeX.

Our MIPS Sail model is just over 2000 non-blank, non-comment lines of Sail code, including sufficient privileged architecture features to boot FreeBSD, but excluding floating point and other optional extensions. The CHERI-MIPS model extends the MIPS model with approximately 2000 more lines and includes support for either the original 256-bit capabilities or a compressed 128-bit format, with the instructions themselves being expressed in a manner that is agnostic to the exact capability format. This is important because CHERI is under continuous development and the capability format has changed many times. For example, Fig. 3 shows the Sail semantics for the CHERI `CIncOffsetImmediate` instruction, to increment the offset of a capability; it makes the various security checks (and the priority among them) explicit.

2.3 ARMv8-A

This is our most substantial example by far: ARMv8-A is a modern industry architecture, underlying almost all mobile devices. It was announced in 2011 and has been enhanced through to ARMv8.2-A (2016), ARMv8.3-A (2016), and ARMv8.4-A (2018). It includes both 64-bit (AArch64) and 32-bit (AArch32) instruction sets. ARM also define related microcontroller (-M) and real-time (-R) variants.

The ARM architecture specifications have long used a custom pseudocode metalanguage, ASL, to express instruction behaviour. ASL has evolved over time. It was initially purely a paper language, an important part of the manuals but not mechanically parsed, let alone type-checked or executed. Reid led an effort within ARM to improve this, so that “machine-readable, executable specifications can be automatically generated from the same materials used to generate ARM’s conventional architecture documentation” [Reid 2016, 2017; Reid et al. 2016]. This executable version of ASL is now used within ARM in documentation, hardware validation, and architecture design, alongside other modelling approaches.

In 2017 ARM released a machine-readable version of large parts of the ARMv8.2-A ASL, later updated to 8.3 and 8.4. This describes almost all of the sequential aspects of the architecture: instructions, floating point, page table walks, taking interrupts, taking synchronous exceptions such as page faults, taking asynchronous exceptions such as bus faults, user mode, system mode, hypervisor mode, secure mode, and debug mode. This provides a remarkable opportunity to rebase research on formal verification, analysis, and testing for ARM above largely complete (for sequential

code) models based on an authoritative vendor-supplied semantics. However, that public release does not include tools for executing or reasoning about the ASL code, and it is not in a form usable for mechanised proof or integration with relaxed-memory concurrency semantics.

Accordingly, we have co-designed Sail and an `asl_to_sail` translation tool that can translate these ASL specifications into Sail (itself open-source), and thence into multiple theorem-prover and emulator-code targets. We have done this both for that public 8.3 release and for an ARM-internal version of 8.3 that additionally includes semantics of the many hundreds of system registers, some of which are needed during an OS boot; we are exploring the possibilities for also releasing this.

The total size of the public v8.3 specification when translated into Sail is about 23 000 lines, including 1479 functions, and 245 registers. This includes all 64-bit instructions, which are expressed as 344 function clauses in Sail, each of which may correspond to multiple assembly mnemonics. So far we have focused on the AArch64 64-bit part of the architecture, and have not translated the (optional) AArch32 32-bit mode. For the non-public v8.3 specification, which additionally includes a full description of all the system registers, we opted to not translate the vector instructions (they add considerably to the size of the specification), as we were primarily interested in the system-level parts of that specification. However, even without vector instructions it contains approximately 30 000 lines of specification with 1279 functions and 501 registers, implementing a total of 390 instructions. In contrast to the simple RISC-V instruction shown in Fig. 2, a single ARM instruction may involve hundreds of auxiliary functions, e.g. for checks of the current exception level and suchlike. While booting Linux, we found that each instruction calls on average around 800 other auxiliary functions, and around 500 primitive operations.

In addition to translating the base specification, we have also added additional hand-written specification for timers, memory-mapped I/O (e.g. for a UART), and interrupt handling based on ARM's generic interrupt controller (GIC), which is sufficient to boot Linux using the model.

Our `asl_to_sail` tool is capable of translating the majority of ASL functions directly into Sail. Both Sail and ASL are first-order imperative languages, and most constructs can be translated in a straightforward manner. The main difficulty come from translating between the two type systems. Sail and ASL both have dependent types, but constructing well-typed Sail from ASL is sometimes non-trivial due to how the type systems differ. Sail's dependent type system and how it is translated from ASL is described more fully in §3. Roughly speaking our tool uses a mix of Sail's own type inference rules and some syntax-based heuristics to synthesise Sail types from ASL types. Some manual patching is needed, so `asl_to_sail` allows for interactive patching during translation. These patches are remembered and can be applied again automatically when the tool is re-run; We had to significantly re-engineer parts of the Sail language to support the kind of incremental parsing and type-checking required by `asl_to_sail`. Translating the non-public spec required 525 lines out of approximately 30 000 to be changed in some way, which represents patches to 143 top-level definitions out of 2158 that were translated. Most of these only require small tweaks and additional type annotations, with the median number of changed lines per patched top-level definition being 3. For the public specification, we also had to manually remove the mutual recursion from the translation table walk, as the ASL code is several hundred lines long and Isabelle has performance issues with such large mutually recursive functions. Fortunately, the maximum recursion depth in this case is only two.

Fig. 4 shows a sample instruction family automatically translated from ASL into Sail, the ADD / SUB (immediate) instructions. This illustrates several of the difficulties of working with the vendor definition: computed bitvector sizes and the use of imperatively updated local variables, initially **undefined**. `AddWithCarry` is an auxiliary pure function, defined in the ASL and also translated to sail, that does the required arithmetic over the mathematical integers and also computes the resulting flag values. Register accesses are indirected via other auxiliary functions, e.g. `aset_X`, which do

```

393 val aarch64_integer_arithmetic_addsub_immediate : forall ('datasize : Int).
394   (int, atom('datasize), bits('datasize), int, bool, bool) -> unit
395   effect {escape, rreg, undef, wreg}
396
397 function aarch64_integer_arithmetic_addsub_immediate ('d,datasize,imm,'n,setflags,sub_op)
398 = { assert(constraint('datasize in {8, 16, 32, 64, 128}), "datasize_constraint");
399   result : bits('datasize) = undefined;
400   operand1 : bits('datasize) = if n == 31 then aget_SP() else aget_X(n);
401   operand2 : bits('datasize) = imm;
402   nzcvcv : bits(4) = undefined;
403   carry_in : bits(1) = undefined;
404   if sub_op then {
405     operand2 = ~(operand2);
406     carry_in = 0b1
407   } else carry_in = 0b0;
408   (result, nzcvcv) = AddWithCarry(operand1, operand2, carry_in);
409   if setflags then (PSTATE.N @ PSTATE.Z @ PSTATE.C @ PSTATE.V) = nzcvcv else ();
410   if d == 31 & ~(setflags) then aset_SP(result) else aset_X(d, result)
411 }

```

Fig. 4. ARMv8-A ADD / SUB (immediate) Instruction in Sail, as translated from ASL

zero-extension if needed, select the appropriate register for the current exception level, check permissions, etc.

3 THE SAIL LANGUAGE

Sail as a language has to be sufficiently expressive to idiomatically express real ISAs, but no more expressive than necessary, otherwise translations to idiomatic prover definitions and fast emulator code would be more challenging, and readability by practising engineers would suffer.

The language is statically checked, with type inference and checking both to detect specification errors and to aid the generation of target code. We also have an interactive Sail interpreter, which can be used for debugging via breakpoints and interactively stepping through the evaluation of functions, and also provides a useful reference semantics for the language.

Following existing industry ISA pseudocode (both paper languages and ASL), Sail is essentially a first-order imperative language. Avoiding higher-order functions simplifies translation into C for efficient emulator code, simplifies proof about the ISA definitions, and avoids readability difficulties for the many engineers who are not familiar with functional languages. Instruction semantics are intrinsically effectful: instructions read and write registers and memory. In the sequential world, one might imagine that each instruction atomically updates a global machine state. In a realistic relaxed-memory concurrent setting, that is no longer the case, as one has to deal with finer-grain interactions between instructions. Perhaps surprisingly, though, at least for user-mode code it has so far been possible to treat the intra-instruction semantics sequentially, albeit with care to sequence specific register and memory operations correctly (and excluding ARM load-pair) [Flur et al. 2017; Gray et al. 2015; Pulte et al. 2018; Sarkar et al. 2011]. Whether this will remain true for systems-mode concurrency is unknown, but for the moment Sail does not require or support any intra-instruction concurrency.

Instructions refer to a global collection of the architected registers. Some ISA specifications, including ARMv8-A, also rely on imperatively updatable local variables, but general references are not used. Sail supports passing references to registers, which is occasionally useful when trying

442 to stick closely to the appearance of some industry ISAs, but we usually find it preferable to pass
443 numeric (integer-range) register indices instead.

444 Most computation is over bitvectors, integer ranges, or integers, but user-defined enumerations
445 are also needed, as are labelled records and (non-recursive) sums. Sail includes a built-in polymorphic
446 list type. We also support user-defined type-polymorphic functions, and sum types can also be type-
447 polymorphic, so one can implement a standard ‘option’ datatype as in most functional languages,
448 but there is relatively little use of type polymorphism in our ISA models. However, we do need
449 dependent types for bitvector lengths, integer range sizes, and operations on these, as such types
450 can have arbitrary numeric constraints attached to them. This is the most technically challenging
451 aspect of the language design, discussed in the next subsection. Operations on subvectors, and
452 registers and record types with named sub-vector-bitfields, are also needed, including complex
453 l-values for updates to specific parts of complex register state.

454 Architecture specifications commonly leave some bits loosely specified in specific contexts, or
455 have broader loosely specified behaviour. Sail supports the former (**undefined**), with our backends
456 providing various semantics as needed for different purposes. ARM also contains unpredictable
457 behaviour, but this is modelled directly in the specification using ordinary functions that specify
458 how the unpredictable behaviour should be handled.

459 The language includes both loops and recursion, as these are needed in the examples, e.g. in
460 the ARMv8-A address translation code and the `BigEndianReverse` function, that reverses the bytes
461 of a 16/32/64/128-bit bitvector. The Sail for each instruction should be terminating, but Sail itself
462 does not check that; instead it is left to the theorem-prover targets. The termination arguments are
463 usually very simple, e.g. the address translation table walk has at most 4 iterations, and manual
464 termination proofs are rarely needed because most loops are inherently terminating foreach loops.

465 Translating ASL into Sail led us to make changes to the language, to better express the ARMv8-A
466 ASL code. For example, we originally did not plan to include exceptions in Sail, but ASL includes
467 exceptions and exception handling, and uses them to implement some key aspects of the architecture,
468 so we needed to add these to Sail to generate clean definitions (we translate these away in the
469 various targets, as appropriate). We also had to add support for arbitrary-precision rational numbers,
470 as ASL specifies several floating point operations by converting the binary floating point values to
471 rationals, performing arbitrary-precision rational arithmetic, and then rounding back to floating
472 point values with the appropriate precision. ASL also assumes that various components in the
473 model are configurable at run-time, so we had to add support for special ‘configuration registers’
474 to be set by command line flags when the model is used. Such command line flags had to be made
475 compatible with ARM’s tools, so we could run our model with the ARM-internal AVS test-suite.

476 The language includes pattern matching, used especially for bitvector-concatenation pattern
477 matching in decode functions, and for tuples.

478 We support various convenience features tuned for ISA specification. They are typically large
479 and flat, so Sail supports splitting functions and type definitions into multiple clauses which can be
480 *scattered* throughout the file, interleaved with other definitions Fig. 2 shows those clauses for a
481 load instruction from RISC-V formalised in Sail, grouped together in the way they would be in an
482 ISA manual; they could be followed by the clauses for another instruction, perhaps in a different
483 file. Some ISAs, including ARMv8-A, rely on syntactic sugar to define pseudoregisters, that can be
484 used either within lvalues or expressions, with semantics defined by user-defined functions; we
485 support this with an overloading mechanism, much as ASL and L3 do. We include mechanisms for
486 specifying bi-directional mappings between binary opcodes and assembly syntax, discussed below.

487 Good concrete syntax design is important for accessibility. Initially we aimed to exactly match
488 the various industry ISA pseudocode languages, idiosyncratic as they are, and to use a C-like syntax
489 for types and type annotations (e.g. `int x = ...`). Experience showed that neither were sustainable,
490

491 and so we redesigned the Sail syntax more cleanly, but in a way that should still be readable by a
492 broad community of hardware, software, and tool developers, who may be unfamiliar with modern
493 functional languages.

494 Targeting multiple provers —currently Isabelle/HOL, HOL4, and Coq— forces us to be careful that
495 all language features can be translated into usable theorem prover definitions for each, taking their
496 different logical foundations into account. In general, we want to make use of our type system to
497 generate nicer prover definitions where possible. As detailed in §4, we are currently able to generate
498 Coq that preserves most of the liquid types from the Sail specification, whereas for Isabelle and
499 HOL4 we perform a specialised partial monomorphisation that retains useful typing information
500 where possible and tries to avoid duplicating code (as a naive full-monomorphisation pass would
501 do).

502 Any proof based upon an ISA specification is dependent on the specification being correct, but
503 an executable ISA specification is a large and complex program in its own right, as is Sail itself.
504 We prioritise clarity over emulation performance when expressing the specifications, and we have
505 devoted considerable effort to testing Sail, e.g. to ensure that the libraries of bitvector operations do
506 the same in all targets. We only provide arbitrary precision integers, integer ranges, and rationals
507 in Sail—this costs us some performance but guarantees that the specification cannot contain the
508 kind of integer overflow and underflow issues that commonly affect programs written in languages
509 like C. We have implemented Sail so that every intermediate rewriting step from the original Sail
510 source to our theorem prover definitions can be type-checked.

511 512 513 3.1 Dependent types for bitvector lengths and integer ranges

514 Bitvector indexing and manipulation is ubiquitous throughout ISA specifications, including bitvector
515 concatenation and taking sub-bitvector slices, as is indexing into arrays, e.g. indexing from a 5-bit
516 opcode field into an array of 32 general-purpose registers. In a simple idealised ISA the sizes of
517 these bitvectors might all be constants, but in more realistic cases, especially in ARMv8-A, they are
518 very often parameterised or computed. For example, ‘size’ arguments in functions are often small
519 powers of two, like 16, 32, or 64, and instructions often come in variants for multiple sizes. It is also
520 extremely common for such arguments to be linked to others and the return type in dependent
521 ways, such as one argument giving the length of another argument in bytes. Expressions used for
522 indexing often involve nontrivial integer ranges. Sometimes the context determines a bitvector
523 size, e.g. for the result type of a zero- or sign-extend operation.

524 ASL necessarily supports all this, but it does not statically check bitvector accesses. In contrast,
525 Sail is designed to statically check these things wherever we can, without needing the specification
526 to fall back onto bit-list representations. We do so for many reasons: to statically catch many
527 specification errors; to enable specifications to more directly express their intent; to make it
528 possible to generate theorem prover definitions in which the correctness of bitvector accesses and
529 suchlike are guaranteed by the prover type system, rather than needing additional proof; and to
530 simplify the generation of fast emulator code, using fixed-width bitvectors instead of bit-lists.

531 Accordingly, Sail supports a form of *lightweight dependent types* for statically checking vector
532 bounds and integer ranges. To ensure our type system remains as lightweight and engineer-friendly
533 as possible, we use a system inspired by Rondon et al’s liquid types [Rondon et al. 2008], which
534 uses the Z3 SMT solver to automatically solve vector bounds and integer range constraints. In our
535 experience, liquid types are ideal for an ISA description language, as they easily express the often
536 relatively simple numeric constraints that occur when bounds-checking vector accesses or the
537 use of integer ranges, without imposing much burden on the user. Often we only need types with
538 appropriate constraints to be declared as a top-level type signatures, and all the types within the
539

540 function can be automatically inferred, and all types and constraints in the function body can be
 541 automatically inferred and discharged.

542 As mentioned, it is extremely common to want to represent an integer value that is either 16, 32,
 543 or 64. This would be represented in Rondon et al’s notation as: $\{i : \text{int} \mid i = 16 \vee i = 32 \vee i = 64\}$
 544 Our syntax differs slightly from this for historical reasons (previous versions of Sail had a type-
 545 system more similar to dependent ML [Xi 2007]), and in Sail such a type would be specified as
 546 $\{i. 'i \text{ in } \{16, 32\ 64\}, \text{int}('i)\}$. This allows us to write commonly used types succinctly, e.g.
 547 $\text{bits}(8 * 'n)$ for a bitvector of n bytes, but such types can be converted into liquid types notation
 548 such as $\{m : \text{bits} \mid \text{length}(m) = 8 * n\}$ in this case, as described in §6.

549 Rondon et al’s inference algorithm operates in steps: First they perform Hindley-Milner type
 550 inference, before using syntax directed liquid typing rules to generate *liquid constraints*, which are
 551 solved in a final third step. In Sail we use a syntax-directed bidirectional type-system (along the
 552 lines of [Dunfield and Krishnaswami 2013]), so we can generate and solve the constraints as part of
 553 the ordinary typing-rules in a single type checking pass. While this means we do not have full type
 554 inference, in practice we mostly require top-level type declarations, with types within function
 555 bodies being automatically inferred.

```

556
557
558 val LSL_C : forall ('N : Int), 'N >= 1.
559   (bits('N), int) -> (bits('N), bits(1)) effect {escape}
560
561 function LSL_C (x, shift) = {
562   assert(shift > 0);
563   let shift as 'S : range(1, 'N) = if shift > 'N then 'N else shift;
564   let extended_x : bits('S + 'N) = x @ Zeros(shift);
565   let result : bits('N) = slice(extended_x, 0, 'N);
566   let carry_out : bits(1) = [extended_x['N]];
567   return (result, carry_out)
568 }

```

568 Fig. 5. Fully annotated left shift with carry function

570 Fig. 5 shows an example of how dependent types for bitvectors are often used in Sail. The assert
 571 guarantees that the shift variable is greater than zero, and the next **let** statement forces shift to be
 572 in the (inclusive) range 1 to 'N, which the type checker will prove based on the assert and the type
 573 constraint 'N >= 1. In order to refer to the value of shift in type signatures later in this function,
 574 we give it a name as a type variable 'S. The next line extends the input bitvector x with a number
 575 of zeros equal to shift, resulting in a bitvector of length 'S + 'N. Then we take a slice from index 0
 576 of length 'N. Here the type system will prove that 'N <= 'S + 'N to show that the slice does not
 577 violate the bounds of extended_x. The next line accesses the carry bit. Here the type system relies
 578 on the fact that 'S must be greater than 0 to establish that 'N is a valid index into extended_x. In
 579 practice most of the manual type signatures in Fig. 5 are not required, and the body of the function
 580 can be written as below.

```

581 let shift : range(1, 'N) = if shift > 'N then 'N else shift;
582 let extended_x = x @ Zeros(shift);
583 let result = slice(extended_x, 0, 'N);
584 let carry_out = [extended_x['N]];
585

```

586 **Translating from ASL to Sail: Dependent Types** As mentioned in §2.3 there are differences
 587 in the type systems between ASL and Sail that make generating type-correct Sail a significantly
 588

```

589
590
591 bits(N) FPTThree(bit sign)
592   assert N IN {16,32,64};
593   constant integer E =
594     (if N == 16 then 5
595      elif N == 32 then 8
596      else 11);
597   constant integer F = N - (E + 1);
598   exp = '1':Zeros(E-1);
599   frac = '1':Zeros(F-1);
600   return sign : exp : frac;

```

Fig. 6. Original FPTThree ASL

```

val FPTThree : forall 'N. bits(1) -> bits('N)
effect {escape}
function FPTThree sign = {
  assert('N == 16 | 'N == 32 | 'N == 64);
  let E : {|5, 8, 11|} =
    if 'N == 16 then 5
    else if 'N == 32 then 8
    else 11;
  let F = 'N - (E + 1);
  let exp = 0b1 @ Zeros(E - 1);
  let frac = 0b1 @ Zeros(F - 1);
  sign @ exp @ frac }

```

Fig. 7. FPTThree function translated into Sail

601
602

603 harder task than just generating syntactically-correct Sail. ASL's typesystem is a compromise
604 between expressivity and the ability to detect errors: like Sail, it provides dependent types for
605 bitvector sizes and statically checks every function call but, unlike Sail, uses of bitvector indexing
606 are not statically checked. During the conversion of ARM's pseudocode to ASL, ARM's architects
607 requested a more flexible type system, and some form of flow-sensitive typing was considered but
608 then rejected because it was not clear how to get good error messages, how to explain to users
609 what could and could not be typechecked and how to avoid path explosion. Automatic translation
610 of ASL to Sail is therefore not just practically useful to Sail users but also useful to ARM, since it
611 demonstrates that ASL could also adopt flow-typing and gain the benefits of more expressive types
612 and stronger checking. ARM's internal ASL steering committee is currently exploring this option.

613 To illustrate the translation of these dependent types, consider the Sail function FPTThree in Fig. 7
614 translated from the ASL in Fig. 6. It constructs the floating point value 3.0, as either a 16, 32, or 64-bit
615 vector. As can be seen, the length of both the exponent (E) and mantissa (F), are calculated based on
616 the length of the returned bitvector, given by the type variable 'N. Sail will check that the length
617 of sign @ exp @ frac is equivalent to 'N. In Fig. 7, the length of the exponent E has the integer set
618 type {|5, 8, 11|}. Currently only this type signature must be present for this function to type check,
619 while the other type signatures can be inferred automatically (in practice our translation tool will
620 add type signatures wherever it can, but we have omitted them here for brevity).

621 Unlike ASL, Sail has flow-sensitive typing, so the assert statement will guarantee to the type
622 checker that 'N is either 16, 32 or 64 in the body of the function. Typically in our hand-written Sail
623 models, one would put such a constraint on 'N in the type signature, as **val** FPTThree : forall 'N in
624 {|16,32,64|. bits(1)-> bits('N), rather than an assert statement, but in ASL such information
625 is often encoded in runtime assertions. Rather than trying to lift this information into the type
626 signatures, we have generally found that sticking closely with idioms found in ASL, and ensuring
627 that such idioms also work well in Sail (e.g. by adjusting our rules for flow-sensitive typing) has
628 been the best way to easily translate large amounts of ASL into Sail without a large amount of
629 manual effort. Despite this we do make some stylistic improvements when translating ASL code
630 where possible, such as turning some mutable variables in ASL into immutable let-bindings, e.g. exp
631 and frac in Fig. 6. We also have to add an *escape* effect to the function in Sail, as the assert could
632 fail and exit the function. Sail has a basic effect system that keeps track of whether functions
633 read and write registers, and how they interact with memory, as well as other effects such as the
634 aforementioned escape for non-local control flow.

635 Currently we have slightly relaxed Sail's strict bounds checking behaviour for the translated ASL.
636 Sail is able to fully check 2695 bounds checking problems encountered in the ARM specification,

637

```

638 enum rop = { RISCV_ADD, RISCV_SUB, ... }
639 union clause ast = RTYPE : (regbits, regbits, regbits, rop)
640 mapping rtype_mnem : rop <-> string = { RISCV_ADD <-> "add", RISCV_SUB <-> "sub", ... }
641 mapping clause assembly =
642   RTYPE(rs2, rs1, rd, op) <->
643     rtype_mnem(op) ^ spc() ^ reg_name(rd) ^ sep() ^ reg_name(rs1) ^ sep() ^ reg_name(rs2)

```

Fig. 8. Parts of the Sail assembly syntax for RISC-V RTYPE binary operation instructions.

with 48 that are currently not automatically solvable. While we could resolve this by simply adding assertions in the specification where these problems occur using `asL to_sail`'s patching mechanism, we instead plan to improve `asL to_sail`'s ability to infer tight ranges on integer variables, which should help in these cases and also improve code generation.

3.2 Mappings and string pattern-matching

So far we have described the aspects of Sail needed to specify the decoding of binary instructions and their dynamic semantics. When working with an ISA specification, one often also needs the ability to define the assembly language syntax, and the pretty printing and parsing (disassembly and encoding/assembly) functions between it and binary instructions.

Sail *mappings* allow the definition of both sides of a bidirectional function at once, for example a parser and pretty-printer. This is similar to existing work on bidirectional programming, e.g. Boomerang [Bohannon et al. 2008], but much more lightweight. Mappings can be simply defined as a set of pattern-matching clauses, where the right-hand-side of the pattern-match is in itself a pattern, or as pairs of functions, allowing for more complex behaviour such as string conversion to and from integers. The type system allows mappings to be called as if they were functions, with the inferred result type determining in which direction the mapping runs. (This gives rise to the restriction that the types on either side of a mapping must be different.) In the implementation, mappings are expanded into two conventional pattern-matches. Mappings interact usefully with string pattern-matching. We allow string concatenation to be used as an operator in pattern-matches, and attempt a simple left-to-right exact matching (compiled into successively nested guarded pattern-matches). To date, we have handwritten mappings for RISC-V, as in Fig. 8; it should also be possible to generate mappings for ARMv8-A from ARM-supplied metadata.

4 GENERATION OF THEOREM PROVER DEFINITIONS

One of our main goals is to provide theorem prover models upon which verification projects that need detailed ISA specifications can build. For this purpose, we implement automatic translations from Sail code to definitions for different popular theorem provers; we target Isabelle/HOL, HOL4, and Coq (we currently have complete Coq translation only for MIPS). Most of the translation pipeline from Sail to those targets is shared, transforming features such as pattern guards and scattered definitions into forms supported by the targets. Some parts of this pipeline are also shared with generation of emulator code, in §5.

For Isabelle/HOL and HOL4, we first translate to Lem as an intermediate language, using Lem to generate the prover definitions [Mulligan et al. 2014]. Since the RMEM concurrency models are specified in Lem, this translation is also used for the integration of Sail ISA models into RMEM [Pulte et al. 2018]. For Coq, we generate Coq definitions directly from Sail to make better use of Coq's type system, in particular to preserve dependent types for bitvector lengths, which are not supported by Lem or the other provers. We describe how we deal with those dependent types for Isabelle/HOL and HOL4 in §4.1. We explain further details of the translation, in particular the monadic treatment

687 of effects, in §4.2. Our translations are generally intended to handle all of Sail, but there are areas
 688 where we currently require additional restrictions, which are all compatible with our models. For
 689 example, in monomorphisation we currently only support case splits on the types used in practice.

690

691

4.1 Bitvector Length Monomorphisation

692

693

694

695

696

697

698

As described above, Sail’s type system can track the sizes of bitvectors with a reasonably rich
 suite of type-level arithmetic operations, backed by constraint solving. This is convenient for
 expressing data-dependent bitvector sizes, such as the data size used in the instruction shown in
 Fig. 4. However, Isabelle/HOL and HOL4 only permit very simple expressions at the type level;
 essentially just constants and variables. To translate into these, we have added a bitvector library
 to the intermediate Lem language, and perform a partial monomorphisation of models to fit them
 into these less expressive type systems.

699

700

701

702

703

704

705

706

The approach is similar to one previously used by ARM during translation to Verilog for model
 checking [Reid et al. 2016, §4], where additional case splits are added to ensure that all bitvector
 sizes will be constant, and constant propagation reveals exactly what those sizes are. Our goals
 are slightly different, however. We want to retain the original model structure as far as possible,
 in particular avoiding the duplication of functions due to specialisation. Fortunately, Isabelle and
 HOL4 support non-dependent size parametricity, representing sizes as type variables. For example,
 in the ARMv8-A model a case split for the data size can sometimes be introduced in the decoder,
 and the more complex execution function left parametric in the size.

707

708

709

710

711

712

713

The location of the case splits to be introduced is determined by an automated interprocedural
 dependency analysis. Case splits on bitvector and enumeration variables are simple to introduce,
 but for integer variables we consult the Sail typing to find the set of possible values. The constant
 propagation is also mildly interprocedural so that trivial helper functions can be eliminated. When
 a case split refines the type of an argument or a result, e.g., from `bits('n)` to `bits(8)`, etc. by a case
 split on 'n, we introduce a cast using a primitive zero-extension operation, which will change the
 type but not the value.

714

715

716

717

To reduce the amount of code duplication we perform a transformation on type signatures before
 monomorphisation. This lifts complex sizes out of types in function signatures, allowing them to
 be treated as type parameters. For example, a simple memory loading function might have the
 signature

718

719

720

721

722

```
val load : forall 'n, 'n >= 0. (bits(64), bits(8 * 'n)) -> bits(64)
```

suggesting that `8 * 'n` must be monomorphised in the body of `load` because it cannot be represented
 in Lem’s type system. Instead, we rewrite it to the equivalent signature

723

724

725

726

727

728

729

730

```
val load : forall 'n 'm, 'n >= 0 & 'm = 8 * 'n. (bits(64), bits('m)) -> bits(64)
```

making the size a proper type parameter, which can be expressed in Lem.

729

730

731

732

733

734

735

For some combinations of variable-size bitvector operations it is preferable to rewrite them in
 terms of shifting and masking on a suitably large fixed-size bitvector. For example, comparing two
 slices of bitvectors `v[x .. y] == w[x .. y]` can be replaced by masking `v` and `w` and comparing,
 without needing to monomorphise `y-x`. We have a small library of combined operations like this,
 and a set of rewrites to use them.

730

4.2 Monadic Translation of Effects

731

732

733

734

735

The translation of imperative, effectful Sail code into monadic code for the generation of prover
 definitions is largely standard, rewriting into a sequence of monadic and let-bindings similar to
 A-normal form [Flanagan et al. 1993], but where the criterion is that arguments to functions must
 be pure. For example, the effectful first operand of `add_vec` in Fig. 9 has been pulled out into a

```

736 fun execute_LOAD :: "12 word=>5 word=>5 word=>bool=>word_width=>bool=>bool=>bool M" where
737 "execute_LOAD imm rs1 rd is_unsigned width aq rl = (
738   rX (regbits_to_regno rs1) >>= (λ w__0.
739     let (vaddr :: xlenbits) = add_vec w__0 ((EXTS 64 imm)) in
740     if check_misaligned vaddr width then
741       handle_mem_exception vaddr E_Load_Addr_Align >> return False
742     else
743       translateAddr vaddr Read Data >>= (λ w__1 :: TR_Result.
744         (case w__1 of
745           TR_Failure (e) => handle_mem_exception vaddr e >> return False
746         | TR_Address (addr) =>
747           (case width of
748             BYTE => mem_read addr 1 aq rl False >>= (λ w__2 :: 8 word MemoryOpResult.
749               process_load rd vaddr w__2 is_unsigned)
750           | HALF => mem_read addr 2 aq rl False >>= (λ w__4 :: 16 word MemoryOpResult.
751               process_load rd vaddr w__4 is_unsigned)
752           | WORD => ...
753           | DOUBLE => ...)))))"

```

Fig. 9. RISC-V load instruction translated into Isabelle

754
755

756 monadic bind. Local mutable variable updates are translated to pure let-bindings, where local
 757 blocks that update variables, e.g. loop bodies and the branches of if-expressions, are rewritten
 758 to return the updated values so that they can be picked up by the surrounding context, while
 759 respecting their scoping. This avoids generating and handling per-function local state spaces, and
 760 the need for a polymorphic state that is difficult to support in the non-dependently typed backends.
 761 Early return statements in functions are translated in terms of the Sail exception mechanism,
 762 by throwing the return value and wrapping the function body in a try-catch-block, where early
 763 returns and proper exceptions are distinguished using a sum type. The translation assumes a left-
 764 to-right evaluation order of effectful function arguments. Boolean conjunction and disjunction are
 765 special-cased, however, to give them a short-circuiting semantics. This is required for the ARMv8-A
 766 specification, which includes expressions such as `UsingAArch32() && AArch32.ExecutingLSMInstr()`,
 767 where an assertion fails in the right-hand function if the left-hand function does not return `true`.

768 Our translation targets two monads with different purposes. The first is a state monad with
 769 nondeterminism and exceptions. It is suitable for reasoning in a sequential setting, assuming that
 770 effectful expressions are executed without interruptions and with exclusive access to the state.
 771 Nondeterminism is needed for aspects that the architecture loosely specifies, and for features such
 772 as load reserve/store conditional instructions that can succeed or fail. The second is a monad for a
 773 concurrent semantics, where a standard state monad interpretation of the Sail code is insufficient.
 774 In particular, in the relaxed memory models of ARMv8 and RISC-V, instructions observably execute
 775 out-of-order, speculatively, and non-atomically, and so the semantics needs to expose the instruc-
 776 tions' effects at a finer granularity. For example, a store instruction waits until all program-order
 777 preceding memory accesses have resolved their address before it can propagate, and so it can
 778 observe intermediate states in the execution of those preceding instructions. To support integrating
 779 Sail with these concurrency models, we use a free monad of an effect datatype. It is implemented
 780 in terms of a monad type as below, parameterised by the return value type 'a, the register value
 781 type 'r, and the exception type 'e (such a monad is often implemented using a generic functor
 782 `Free`, e.g. in Haskell, but since this is not supported by the type system of Isabelle, we merged it
 783 with the concrete effects into a single type).

784

```

785 type monad 'r 'a 'e =
786   | Done of 'a
787   | Read_mem of read_kind * address * nat * (list mem_byte -> monad 'r 'a 'e)
788   | Write_ea of write_kind * address * nat * monad 'r 'a 'e
789   | Write_memv of list mem_byte * (bool -> monad 'r 'a 'e)
790   ...
791   | Barrier of barrier_kind * monad 'r 'a 'e
792   | Read_reg of register_name * ('r -> monad 'r 'a 'e)
793   | Write_reg of register_name * 'r * monad 'r 'a 'e
794   | Undefined of (bool -> monad 'r 'a 'e)
795   | Fail of string (* Assertion failure with error message *)
796   | Exception of 'e (* Exception thrown *)

```

797 A value of this type is either Done *a*, representing a finished computation with a pure value
798 of type 'a, or an *effect request*: each of the other constructors represents an effect, typically to-
799 gether with some parameters specifying the particular request, and a continuation. For example,
800 Read_reg "PC" *k* is a request to the execution context to read the PC register and pass its value
801 into the continuation *k*. Another example is Undefined *k*, which requests a Boolean value from
802 the execution context, e.g. to make a nondeterministic choice or to resolve an undefined bit to
803 a concrete value. The definition of the monad leaves the meaning of these instruction effects
804 open—the monad’s bind operator simply “nests” the requests—and the monad instead delegates
805 handling the effects to an *effect interpreter* outside the instruction semantics definition. To support
806 the integration with a concurrency model that executes these instruction definitions out-of-order,
807 the monad type has effects for all concurrency-relevant events of the instruction’s execution: for
808 example, the Barrier effect announces memory barriers, register reads and writes are explicit
809 requests (Read_reg and Write_reg) to enable handling the fine-grained memory ordering resulting
810 from dataflow dependencies, and the writing of memory is split into the announcement of the
811 write address, Write_ea, and the writing of the value, Write_memv, so program-order succeeding
812 instructions can be informed about the address as early as it is known.

813 4.3 Target-specific Differences in the Translation

814 Most of the translation pipeline is shared between the different provers, e.g. the rewriting of
815 bitvector patterns to guarded patterns, and then the rewriting of those to a combination of if-
816 expressions and unguarded pattern matches using an algorithm similar to that of [Spector-Zabusky
817 et al. 2018, §3.4].¹ There are some differences, however, mainly due to the differences in the type
818 systems of the provers.

820 **Isabelle** The prover definitions generated from a Sail model should ideally be parametric in
821 the monad, but this is not supported by Isabelle’s type system. Hence, when generating Isabelle
822 definitions, we use the prompt monad, and provide a lifting to the state monad that enables
823 reasoning in terms of the latter, if desired (cf. §8).

824 **HOL4** When generating HOL4 definitions, we use only the state monad, since HOL4’s datatype
825 package does not currently support the prompt monad’s type (it has a recursion on the right of a
826 function arrow).

828 **Coq** The dependent type system in Coq enables us to give a much more direct translation of Sail’s
829 rich type information than would be possible with Lem’s rudimentary Coq output support. The

830 ¹The main differences are that we use a different grouping strategy for clauses (overlapping instead of mutually exclusive
831 groups, since bitvector pattern rewriting can lead to many consecutive, overlapping patterns), and that we keep fall-through
832 branches in place instead of pulling them out into let-bindings, since that could interfere with both effects and flow-typing.

```

834 Definition FPThree (N_tv : Z) (sign : mword 1) : M (mword N_tv) :=
835   assert_exp' ((Z.eqb N_tv 16) \\/ (Z.eqb N_tv 32) \\/ (Z.eqb N_tv 64)) "" >>= fun _ =>
836   let '(existT _ E _) :=
837     (if sumbool_of_bool ((Z.eqb N_tv 16)) then build_ex 5
838     else if sumbool_of_bool ((Z.eqb N_tv 32)) then build_ex 8
839     else build_ex 11)
840     : {n : Z & ArithFact (In n [5; 8; 11])} in
841   let F := Z.sub N_tv (Z.add E 1) in
842   let exp := concat_vec (vec_of_bits [B1] : mword 1) (Zeros__0 (Z.sub E 1)) in
843   let frac := concat_vec (vec_of_bits [B1] : mword 1) (Zeros__0 (Z.sub F 1)) in
844   returnm ((autocast (concat_vec sign (concat_vec exp frac))) : mword N_tv).

```

Fig. 10. FPThree function translated into Coq

845

846

847

848

849

850

851

852

853

854

855

856

857

858

859

860

861

862

863

864

865

866

867

868

869

870

871

872

873

874

875

876

877

878

879

880

881

882

main difference in our Coq translation compared to our other backends is that the type-level sizes and constraints are fully retained in the generated Coq definitions. In particular, Sail's existential types are translated to dependent pairs in Coq. This can be seen in Fig. 10, the translation of Fig. 7, where a dependent pair is built for E to show that it is in the set $\{5, 8, 11\}$. However, it would be extremely challenging to reuse proofs about the constraints from the SMT solver used during Sail type checking, so instead we use a Coq typeclass wrapper to trigger a constraint solving tactic. In Fig. 10 this is done by the `build_ex` function. The core of the tactic is Coq's implementation of the Omega Presburger arithmetic decision procedure [Pugh 1991], with additional preprocessing to transform information from the context into a useful form and to evaluate constant powers of 2. The solver can also be extended by adding facts to a Coq hints database.

There is an important difference between the type checking in Sail and Coq: Sail uses the SMT solver to assist with type equivalence and subtyping checks automatically, whereas Coq only uses its built-in notion of reduction. This is often inadequate; for example, even $1 * z$ does not reduce to z for Coq's integer type, Z , whereas Sail considers the types `bits(1 * z)` and `bits(z)` to be interchangeable. Our Coq backend detects differences like this and inserts a cast function. The cast function has a constraint that the two integer expressions are equal (which is automatically inferred from the context by Coq), and triggers the constraint solving tactic during type checking. For example, in the last line of `FPThree` the cast uses a proof of $1 + (1 + (E - 1) + (1 + (F - 1))) = N_tv$.

The Coq backend is still under development. For example, it is still restricted to fairly simple uses of Sail's existential types. Nonetheless, it is already sufficient to produce a full Coq translation of our MIPS model.

5 GENERATION OF EMULATORS

It is also important to support emulation, with enough performance for validation purposes. We implement a simple direct mapping from Sail into OCaml, as well as more involved optimised compilation path to C. The simple OCaml translation is primarily used as a validation tool for the more involved C translation, and for prototyping.

Our C generation involves several steps. First we use the same type-preserving rewrites we use when generating theorem prover code, to eliminate some features and syntactic sugar found in the full Sail language, then we map into an A-normal form representation which is very similar to the MiniSail language described in §6. This is then translated to a lower-level intermediate representation, before we generate C code. Our intermediate representation is not particularly tied to C, so we could easily switch to e.g. LLVM IR if desired at some point in the future. There are

883 three main optimisations that we perform during this compilation process that greatly speed up
 884 the resulting code.

885 First, bitvectors that are statically known to be 64-bits in length or less are mapped to 64-bit
 886 unsigned integers in C, whereas variable length bitvectors or those that are larger than 64 bits are
 887 mapped onto arbitrary length bitvectors implemented using GMP integers. Furthermore, some
 888 bitvector types larger than 64-bits are mapped onto multiple 64-bit integers if necessary—this was
 889 key to get good performance for MIPS address translation, which features 117-bit wide bitfields for
 890 TLB entries.

891 Secondly, we use our liquid types and constraints to detect integer types that are bounded to
 892 fit within an 64-bit signed integer type. For example, the `int_of_accessLevel` function below is
 893 returning an integer constrained to be either 0, 1, or 2. As such rather than using an arbitrary
 894 precision GMP integer, we return a fixed-width integer type. In general we can optimise any such
 895 integer types, provided Z3 can prove they fit within the bounds of a 64-bit signed integer.

```
896 function int_of_AccessLevel(level : AccessLevel) -> {|0,1,2|} =  

  897   match level { User => 0, Supervisor => 1, Kernel => 2 }
```

899 This optimisation turns out to be important, because small often-used functions like the above
 900 can be quite costly if they are forced to use arbitrary-precision integers. The `int_of_accessLevel`
 901 function accounted for nearly 5% of the time taken booting FreeBSD on our MIPS model before we
 902 implemented this optimisation.

903 Thirdly, we note that the vast majority of functions in ISA specifications are non-recursive, with
 904 the ARM specification containing only a single recursive function (for endianness reversal) and
 905 one small group of mutually recursive functions (for nested page table walks). For all non-recursive
 906 functions we are able to statically preallocate any space they need on the heap for arbitrary-precision
 907 bitvectors and integers to avoid calling `malloc` and `free`.

908 In total these optimisations gave us around a 13x performance increase in the performance of the
 909 generated emulator code. For ARM, we went from approximately 4000 instructions per second (IPS)
 910 to around 53 000 IPS. For MIPS, which is significantly simpler, we were able to achieve performance
 911 of between 500 000 and 1.5 million IPS (this difference being caused by the number of memory
 912 accesses), with an average of about 850 000 IPS. By contrast, when compiling Sail into OCaml we
 913 can execute instructions at around 1800 IPS for ARM, and when using our interactive interpreter
 914 we can execute ARM instructions at about 30 IPS.

915 6 FORMAL TYPE SYSTEM

916 ISA specifications are long-lived, and any formal work above them will involve substantial invest-
 917 ment, so we want Sail to be robust and stable. Unfortunately, the initial version of Sail (like most
 918 architecture description languages) was implementation-defined: Sail was whatever the implemen-
 919 tation would accept. This made evolving the system—even bug-fixing—a fraught process, and
 920 the fact that Sail’s type inference relied on a complex custom constraint solver meant that were
 921 many bugs to fix. Solving this problem required a degree of care, since we wanted to improve the
 922 language design “in place”. To guide the evolution of Sail, we introduced a kernel calculus, MiniSail.

924 The two key properties we prove of MiniSail are (a) type safety, and (b) decidability of type
 925 checking. As a first-order language, the dynamic semantics of Sail are largely straightforward, but
 926 type safety is not entirely obvious, because Sail’s support for type-dependency means that safety
 927 can rely upon control- and data-dependent properties. For software engineering reasons, we wanted
 928 to move away from a hand-rolled constraint solver to using an off-the-shelf SMT solver such as Z3.
 929 This both simplifies our implementation, and increases its reliability, as a widely-used solver like
 930 Z3 will be tested much more thoroughly than a hand-rolled solver. However, a danger is that we
 931

would merely replace one ad-hoc set of heuristics (embodied in our solver) with Z3's – a different set not even under own control. Luckily, SMT solvers come with a very clear guarantee: any query in the quantifier-free fragment is decidable, and in practice those we generate are efficient. So our decidability proof fundamentally exists to ensure that Sail's type system only generates pure SMT queries, ensuring that the specification of Sail is independent of the details of the solver.

Fig. 11 presents a subset of MiniSail's grammar and a table of judgements, and Fig. 12 presents a selection of the typing rules. The grammar in Fig. 11 defines a language in a slight variant of A-normal form. Programs are defined from expressions (which include arithmetic, variables, and function applications), and statements (which include let-binding, conditionals, and the assignment and declaration of mutable variables). Note that we distinguish bindings of immutable variables **let** $x = e$ **in** s from the declaration of mutable variables **var** $u : \tau = v$ **in** s . Types τ are set-comprehension-style: they are the elements of a first-order base type, together with a boolean constraint restricting which elements of the base type are in τ .

As described in §3, MiniSail is a bidirectional type system, with a type synthesis mode $\Pi; \Gamma; \Delta \vdash e \Rightarrow \tau$ for expressions, and a type checking mode $\Pi; \Gamma; \Delta \vdash s \Leftarrow \tau$ for statements. The presence of three contexts arises from the fact that MiniSail is a first-order, imperative language. Function definitions are separated into a context Π , and Γ and Δ control the scoping of immutable bindings and mutable variables, respectively.

Value	$v ::= x \mid n \mid \mathbf{T} \mid \mathbf{F}$	Expression $e ::= v \mid v \oplus v \mid f \ v \mid u$
Statement	$s ::= v \mid \mathbf{let} \ x = e \ \mathbf{in} \ s \mid \mathbf{var} \ u : \tau = v \ \mathbf{in} \ s \mid u := v \mid \mathbf{if} \ v \ \mathbf{then} \ s_1 \ \mathbf{else} \ s_2$	
Constraint	$\phi ::= \top \mid e_1 = e_2 \mid e_1 \leq e_2 \mid \phi_1 \wedge \phi_2 \mid \phi_1 \vee \phi_2 \mid \neg \phi$	
Base Type	$b ::= \mathbf{int} \mid \mathbf{bool} \mid \mathbf{unit}$	Type $\tau ::= \{z : b \mid \phi\}$

Π	Function definition context	Γ	Immutable variable context	Δ	Mutable variable context
$\Pi; \Gamma \vdash v \Rightarrow \tau$	Type synthesis values		$\Pi; \Gamma \vdash v \Leftarrow \tau$	Type checking values	
$\Pi; \Gamma; \Delta \vdash e \Rightarrow \tau$	Type synthesis expressions		$\Pi; \Gamma; \Delta \vdash s \Leftarrow \tau$	Type checking statements	
$\Pi; \Gamma \vdash \tau_1 \lesssim \tau_2$	Subtyping		$\Pi; \Gamma \models \phi$	Validity	

Fig. 11. MiniSail Grammar Fragment and Judgements

$\frac{}{\Pi; \Gamma \vdash n \Rightarrow \{z : \mathbf{int} \mid z = n\}} \mathbf{1}$	$\frac{}{\Pi; \Gamma \vdash \mathbf{T} \Rightarrow \{z : \mathbf{bool} \mid z = \mathbf{T}\}} \mathbf{2}$	$\frac{}{\Pi; \Gamma \vdash \mathbf{F} \Rightarrow \{z : \mathbf{bool} \mid z = \mathbf{F}\}} \mathbf{3}$
$\frac{x : b[\phi] \in \Gamma}{\Pi; \Gamma \vdash x \Rightarrow \{z : b \mid z = x\}} \mathbf{4}$	$\frac{\Pi; \Gamma \vdash v_1 \Rightarrow \{z_1 : \mathbf{int} \mid \phi_1\} \quad \Pi; \Gamma \vdash v_2 \Rightarrow \{z_2 : \mathbf{int} \mid \phi_2\}}{\Pi; \Gamma; \Delta \vdash v_1 + v_2 \Rightarrow \{z_3 : \mathbf{int} \mid z_3 = v_1 + v_2\}} \mathbf{5}$	$\frac{\mathbf{val} \ f : (x : b[\phi]) \rightarrow \tau \in \Pi \quad \Pi; \Gamma \vdash v \Leftarrow \{z : b \mid \phi\}}{\Pi; \Gamma; \Delta \vdash f \ v \Rightarrow \tau[v/x]} \mathbf{6}$
$\frac{u : \tau \in \Delta}{\Pi; \Gamma; \Delta \vdash u \Rightarrow \tau} \mathbf{7}$	$\frac{\Pi; \Gamma \vdash v \Leftarrow \tau}{\Pi; \Gamma; \Delta \vdash v \Leftarrow \tau} \mathbf{8}$	$\frac{\Pi; \Gamma; \Delta \vdash e \Rightarrow \{z : b \mid \phi\} \quad \Pi; \Gamma; x : b[\phi[x/z]]; \Delta \vdash s \Leftarrow \tau}{\Pi; \Gamma; \Delta \vdash \mathbf{let} \ x = e \ \mathbf{in} \ s \Leftarrow \tau} \mathbf{9}$
$\frac{\Pi; \Gamma \vdash v \Leftarrow \tau \quad \Pi; \Gamma; \Delta, u : \tau \vdash s \Leftarrow \tau_2}{\Pi; \Gamma; \Delta \vdash \mathbf{var} \ u : \tau := v \ \mathbf{in} \ s \Leftarrow \tau_2} \mathbf{10}$	$\frac{\Pi; \Gamma \vdash v \Rightarrow \{x : \mathbf{bool} \mid \phi_1\} \quad \Pi; \Gamma; \Delta \vdash s_1 \Leftarrow \{z_1 : b \mid (v = \mathbf{T} \wedge (\phi_1[v/x])) \Rightarrow (\phi[z_1/z])\} \quad \Pi; \Gamma; \Delta \vdash s_2 \Leftarrow \{z_2 : b \mid (v = \mathbf{F} \wedge (\phi_1[v/x])) \Rightarrow (\phi[z_2/z])\}}{\Pi; \Gamma; \Delta \vdash \mathbf{if} \ v \ \mathbf{then} \ s_1 \ \mathbf{else} \ s_2 \Leftarrow \{z : b \mid \phi\}} \mathbf{11}$	
$\frac{u : \tau \in \Delta \quad \Pi; \Gamma \vdash v \Leftarrow \tau}{\Pi; \Gamma; \Delta \vdash u := v \Leftarrow \{z : \mathbf{unit} \mid \top\}} \mathbf{12}$	$\frac{\Pi; \Gamma \vdash v \Rightarrow \{z_2 : b \mid \phi_2\} \quad \Pi; \Gamma \vdash \{z_2 : b \mid \phi_2\} \lesssim \{z_1 : b \mid \phi_1\}}{\Pi; \Gamma \vdash v \Leftarrow \{z_1 : b \mid \phi_1\}} \mathbf{13}$	$\frac{\Pi; \Gamma, z_1 : b[\phi_1] \models \phi_2[z_1/z_1]}{\Pi; \Gamma \vdash \{z_1 : b \mid \phi_1\} \lesssim \{z_2 : b \mid \phi_2\}} \mathbf{14}$

Fig. 12. Selected MiniSail Typing Rules

981 The key technical idea ensuring decidability is as follows: whenever we have an expression in the
 982 SMT fragment, we record an exact constraint. Otherwise, we merely propagate any SMT constraints
 983 by attaching them to variables, using A-normal form to ensure that there is always a name to attach
 984 a constraint to. Rules 1-5 in Fig. 12 give typing rules for expressions. Since each of these terms
 985 is in the SMT fragment, we can generate an exact equality constraint for them. However, rules
 986 6 and 7 (for function applications and mutable variables) are for terms not in the SMT fragment,
 987 and so we use the type as an approximation, with mutable variables just looking up the type in
 988 the environment Δ and function applications returning the result type with the argument value
 989 substituted in. Rules 8-12 play the same game with statements. The rules for variables (9 and 10)
 990 state no equalities between expressions and variables, since the expressions include forms (e.g.,
 991 function calls) outside of the SMT fragment. On the other hand, rule 11 for if's scrutinises a value
 992 and can flow the value into the branches.

993 Finally, the constraint discipline pays off in rules 13 and 14, where the subtyping relation is used.
 994 One type is a subtype of another just when they have a common base type and the first type's
 995 constraint implies the second, under the assumption of all the constraints in the context. Since no
 996 rule ever introduces a quantifier, we only generate entailments strictly in the SMT fragment.

997 MiniSail's design is heavily inspired by the observation in Liquid Types [Rondon et al. 2008]
 998 that if logical constraints are determined by the actual arguments to a function, there is no need to
 999 introduce existential constraint variables. However, we do not need a prepass deriving a simply-
 1000 typed skeleton. Our bidirectional [Dunfield and Krishnaswami 2013] algorithm is completely
 1001 syntax-directed, with subtyping checks (the only source of SMT queries) occurring at (syntactically
 1002 evident) checking/synthesis boundaries.

1003 The original Sail implementation had a Hindley-Milner-style typechecker, mated to a custom
 1004 arithmetic constraint solver. This codebase was complicated and could not handle many of the
 1005 constraints generated from the the ARM specification. Sail is now bidirectional, mostly replacing
 1006 unification with constraint solving. The transition is ongoing: unification still plays a role in the
 1007 implementation of function calls, and we still allow the declaration of non-argument-constrained
 1008 quantifiers. Still, performance has dramatically improved: checking a fragment of the ARM spec
 1009 has gone from 10-15 minutes to under 3 seconds.

1010 The operational semantics of the full language (including tuples and sums) is standard, and can
 1011 be found in the supplementary material. The full type safety proof is also in the appendix: the proof
 1012 is long (due to the presence of dependency) but not fundamentally difficult.

1014 7 VALIDATION

1015 We validate our models by using the generated emulators to run test suites and boot various
 1016 operating systems. This also serves to validate the translation from Sail to the various backends.
 1017 Ideally, one might want to have mechanised proofs about the correctness of the translation w.r.t.
 1018 a deep embedding of Sail in each of the provers. However, the effort for this would have been
 1019 prohibitive, especially while the Sail language itself was still evolving. Instead, we follow a testing
 1020 approach here too. We have used Isabelle's code generation feature to extract an OCaml emulator
 1021 from the Isabelle model of CHERI-MIPS, which successfully executes the CHERI test suite, albeit
 1022 slowly. This gives us end-to-end validation for the nontrivial translation pipeline from Sail via Lem
 1023 to Isabelle, including bitvector length monomorphisation and translation of effects (§4).

1025 **ARM** We validated our ARM models first by booting Linux on the non-public v8.3 version with
 1026 system register support (used for the timer, handling interrupts, and controlling the availability of
 1027 architectural features). This does not directly validate the public version of our ARM model, but
 1028 as the two are generated in much the same way from the same ultimate sources, it does provide
 1029

1030 significant confidence. We were able to boot older versions of the Linux kernel, in particular Linux
1031 4.4 (2016). For more recent versions of the kernel, we observed issues with context switching
1032 when run above our model. Linux has changed how context switching is handled to unmap kernel
1033 memory, and it seems that a page fault that is supposed to happen at a certain point does not occur.
1034 This could be due to a bug in the address translation code of our model, in a systems feature, or in
1035 our tooling. However, the problem seems to be subtle enough to only be triggered in some versions
1036 of Linux, and we have not yet fully diagnosed it.

1037 ARM's *Architecture Validation Suite* (AVS) is an extensive set of architectural compliance tests
1038 that are used as part of the signoff criteria for ARM-compatible processors. These tests are usually
1039 run on systems composed of processors, RAM and a verification device that can be used to monitor
1040 the processor's behaviour (e.g. memory accesses and their attributes) and to generate stimulus
1041 (e.g. patterns of interrupts). ARM currently runs these tests on an extension of the public ASL
1042 specification that adds a particular set of configuration choices for the implementation-defined
1043 behaviour, and an ASL specification of the verification device [Reid 2016]. ARM does not currently
1044 publicly release the tests, or the configuration or specification of the test device, but we were able
1045 to use them to test and debug our translation of the ASL specification.

1046 The tests cover many aspects of the AArch64 architecture including all usermode behaviour
1047 (i.e., integer, float and SIMD instructions), and system behaviour (i.e., bigendian support, switching
1048 between 32-bit mode and 64-bit mode, memory protection, exceptions/interrupts, privilege levels,
1049 security and virtualisation). The tests for usermode behaviour make up 31% of the tests (this is
1050 roughly proportional to the fraction of ARM's specification documents and their ASL specification
1051 that describes usermode behaviour). Many of the tests for system behaviour explore obscure corners
1052 of the architecture, such as whether a memory access should be cacheable or non-cacheable if an
1053 operating system marks the page as cacheable but a hypervisor (in which the operating system is
1054 running) marks the same page as non-cacheable, or what exception should be signalled if a bus
1055 fault occurs during a page table walk. (These are just two of thousands of scenarios that are tested.)
1056 The tests consist of over 30 000 test programs and the tests run for billions of instructions.

1057 Our current translation to Sail and our C model generation do not handle certain features of
1058 ARM's specification including the AArch32 instruction set, SIMD instructions, multiprocessor
1059 support and a small number of instructions added in the v8.3 model, and so we restricted our
1060 attention to 15 400 tests that do not rely on these features. Of those 15 400 tests, currently 24 (0.15%)
1061 pass on the ASL model but not on the Sail model. 12 of those are floating point failures, due to
1062 the square root primitive operation returning a rational number; 8 are exception handling failures
1063 due to a particular unallocated exception being misthrown; and the remaining 4 are memory
1064 management failures involving marking page table entries dirty. We are working on fixing these
1065 issues.

1066 **RISC-V** We validated the RISC-V model with the seL4 and Linux boots and against the Spike
1067 reference simulator (the current platform model of our RISC-V OCaml emulator matches that
1068 of Spike). The OCaml emulator is run regularly against the tests in the `riscv-tests` test-suite
1069 repository, and passes all tests for integer and compressed instructions for the user, supervisor
1070 and machine modes (currently amounting to 181 tests). An official compliance test-suite is under
1071 construction by the RISC-V Compliance Working Group, but it has yet to create tests for the 64-bit
1072 architecture. We also compare the trace outputs of the Sail model and a version of Spike modified
1073 to provide additional execution traces, and to have a more regular I/O and timer interrupt dispatch
1074 schedule. Our comparison tool checks that the two simulators execute matching instructions,
1075 integer register writes, CSR reads and writes, LR/SC reservation state modifications, and outputs to
1076 certain device ports. We have ensured that these traces match on all but one of the above tests. The
1077

sole exception is the test for the breakpoint instruction, where the Sail model passes the test but the execution trace differs due to the absence of a debug module.

MIPS and CHERI-MIPS To validate these models we ran the CHERI test suite (which also tests MIPS ISA features) and booted FreeBSD-MIPS with a minimal system model consisting of just a write-only UART for console output. Using Sail’s C backend and gcc 5.4 on an Intel Core i7-4770K desktop CPU clocked at 3.50GHz the boot reached a shell prompt after about 90 million instructions in less than 2 minutes, averaging about 850 000 instructions per second

Coverage An executable-as-test-oracle architectural model makes it possible to assess the specification coverage of tests. We did this for the MIPS and CHERI-MIPS models, simply using the gcov coverage tool on the compiled C. Booting FreeBSD on the MIPS model touched 84.8% of the lines of generated C. Most unexecuted lines were due to instructions that were not used (e.g. debugging, cache management, fused multiply&add) and exception cases that were not hit, such as reserved instructions. The MIPS-only subset of the CHERI test suite covered 97.8% of the MIPS model, with the uncovered code due to missing tests for MIPS features such as unusual TLB page sizes and supervisor mode that are not used by FreeBSD. Coverage for the CHERI model was 94.8%. This found a recently introduced instruction that had no tests and highlighted many exception paths that need more testing.

RMEM concurrency integration We integrated our RISC-V ISA model with the RMEM concurrency exploration tool [Pulte et al. 2018], allowing exploration of its relaxed-memory multi-threaded behaviour. For validation, we compared its behaviour on the library of 7251 litmus tests used to develop the RISC-V memory model [RIS 2017, App. A]. They concur on all except 4, due to a discrepancy between the RISC-V memory model and the Spike single-threaded reference simulator: the former allows store-conditionals to fail early before reading any registers, while the latter does not. We currently forbid this, to match traces with Spike. In addition, for emulator performance reasons, the sequential ISA model uses a definition of the JALR instruction that does not allow the write-before-read behaviour of the concurrent specification.

8 MECHANISED PROOF

To evaluate the usability of the generated theorem prover definitions, we proved a nontrivial property of the ARMv8-A specification in Isabelle/HOL. We focus on address translation from virtual to physical memory addresses. This is a critical part of the architecture specification; playing an important role in separating user-space processes from each other and from the operating system. ARMv8-A address translation is also an informative benchmark of the usability of our theorem prover definitions, as it is one of the most complex parts of the most detailed specification we have. The translation table walk function alone consists of over 500 lines of Sail code, not counting various helper functions. It includes a loop for the table walk, does the construction of the physical address from variable-length bitvector slices, reads and writes memory, and exhibits nondeterminism. The latter arises from underspecification that can be refined by implementations. For example, there is a validity check of page table entries that an implementation may choose to perform (potentially faulting) or to ignore. This is “implementation defined” behaviour in the ASL and translated to a nondeterministic choice in our model. Another source of nondeterminism is undefined values. Address translation returns a record containing the output address and other fields such as permission bits. If one of those fields does not make sense in a given situation, such as the device type field for non-device memory, the ASL code sets it to an “unknown” value or leaves it uninitialised. Again, this is translated to a nondeterministic choice of a value in Sail.

1128 Details like these are typically abstracted away in verification projects involving an ISA semantics.
 1129 This may be essential for reasoning about the ISA semantics in a scalable way, but the underlying
 1130 assumptions should be made explicit. Proving soundness of an abstraction against our model allows
 1131 —and requires— us to do this, in terms of the model. As our example, we therefore defined a purely
 1132 functional characterisation of ARMv8-A address translation in a user-mode setting. Our function
 1133 `read_tables` extracts from memory a snapshot of the translation tables (up to four hierarchical
 1134 levels deep) starting at a given base address, while `walk_tables` is a partial function that takes a
 1135 table snapshot and an input address and looks up the corresponding descriptor. The partial function
 1136 `translate_address` calls those two, checks the permission bits, and, if all checks succeed, constructs
 1137 a result record containing an address descriptor with the output address and its attributes, and
 1138 potentially a descriptor update, if hardware updating of access and dirty bits is enabled. The function
 1139 `update_descriptor` writes back the updated descriptor, if necessary.

1140 This characterisation of address translation is quite detailed, but we do make some simplifying
 1141 assumptions. We assume a setting in 64-bit user mode and not in a “secure” state, which is an
 1142 isolation feature of the ARM architecture. We also assume that no virtualisation is active, so only
 1143 one and not two stages of address translation. Moreover, we assume that hardware updating of
 1144 descriptor flags is enabled (the Linux kernel uses this in its default configuration). Without it,
 1145 translating an address within a page or block without the access flag set results in a translation
 1146 fault. Finally, we assume that the MMU is enabled and debug events are disabled. We formalise
 1147 these assumptions as state predicates. For example, the predicate `HwUpdatesFlags(s)` requires that
 1148 bits 39 and 40 of the `TCR_EL1` system register are set. We omit the definitions of these predicates
 1149 and functions here and refer to the supplementary material.

1150 We have proved the following soundness result about our characterisation w.r.t. the original
 1151 function `AArch64_TranslateAddress` defined in the model, where $\llbracket \cdot \rrbracket$ denotes the lifting from free
 1152 to state monad mentioned in §4.2, the relation \approx_{det} denotes equivalence of the deterministic parts
 1153 of address descriptors, ignoring undefined parts, and `Value` indicates a successful outcome of an
 1154 expression in the state monad, as opposed to an exception denoted using `Ex` (where in this case,
 1155 the preconditions guarantee that there is no exception).

1156 THEOREM 8.1. *If*

- 1157 • $\text{InUserMode}(s) \wedge \text{NonSecure}(s) \wedge \text{MMUEnabled_EL01}(s) \wedge \text{VirtDisabled}(s) \wedge$
 1158 $\text{HwUpdatesFlags}(s) \wedge \text{UsingAArch64}(s) \wedge \text{DebugDisabled}(s)$ *and*
- 1159 • $\text{translate_address}(vaddr, acctype, iswrite, aligned, size, s) = r$

1160 *then*

1161 $\forall (\text{Value}(r'), s') \in \llbracket \text{AArch64_TranslateAddress}(vaddr, acctype, iswrite, aligned, size) \rrbracket (s).$

1162 $r' \approx_{det} \text{addrdesc}(r) \wedge s' = \text{update_descriptor}(r, acctype, iswrite, s)$

1163 The assumption that the partial function `translate_address` successfully returns a value implies
 1164 that all checks have passed and all table entries related to the input address are valid. If one of
 1165 those checks fails, then the original address translation function returns a record detailing which
 1166 kind of fault occurred; we do not currently model faulting behaviour in our characterisation.

1167 This means that Theorem 8.1 may not shed light on any potential address translation bug related
 1168 to the Linux booting issue of §7, as that would involve a page fault. However, our proof did uncover
 1169 a missing endianness reversal and several potential uses of uninitialised variables in the original
 1170 ASL code, which have been reported to and confirmed by ARM.

1171 Our Isabelle proof is with respect to the sequential Sail model in the state-nondeterminism-
 1172 exception monad. We manually stated and proved a loop invariant for the translation table walk,
 1173 and Hoare triples about various helper functions. This helps reduce the complexity of the main
 1174 and Hoare triples about various helper functions. This helps reduce the complexity of the main
 1175 and Hoare triples about various helper functions. This helps reduce the complexity of the main
 1176 and Hoare triples about various helper functions. This helps reduce the complexity of the main

1177 proof, which uses an automatic proof method that iteratively applies the basic proof rules of the
1178 Hoare logic and the helper lemmas to derive a precondition for a given postcondition. The Isabelle
1179 proof scripts can be found in the supplementary material.

1180 9 RELATED WORK

1182 There is extensive work on low-level verification using ISA specifications, as well as language design
1183 for ISA description languages, e.g. [Dias and Ramsey 2010; Misra and Dutt 2008]. As mentioned
1184 in the introduction, there exist many smaller partial formal ISA models, usually created for very
1185 specific purposes. Here we mostly focus on work that involves larger specifications including some
1186 system-level features.

1187 The ACL2 X86isa model [Goel et al. 2017], is a hand-written specification of the (64-bit) IA-32e
1188 mode of the x86 architecture. It contains a very comprehensive specification of user-mode parts of
1189 the architecture, as well as system-level features including paging, segmentation, and a system call
1190 interface. Their model has been extensively validated via co-simulation with actual x86 processors.
1191 This work represents the most complete public x86 specification to date. Our work differs mainly
1192 in targeting different architectures, providing for multiple LCF-family theorem provers (Isabelle,
1193 HOL4, and Coq) rather than ACL2, using a dependently typed metalanguage and (validated but not
1194 proved) translations from it rather than working entirely within ACL2, and in translating from the
1195 vendor-supplied ARMv8-A specification. Our models have sufficient system-feature coverage to
1196 boot operating systems, though that is particularly challenging for x86.

1197 L3 [Fox 2012, 2015; Fox et al. 2017] is a well-developed ISA specification language, which like Sail,
1198 supports multiple prover targets (HOL4 and Isabelle/HOL), and has existing models for numerous
1199 architectures. L3 was a key inspiration in the design of Sail, which differs principally in its more
1200 sophisticated type-system (better able to express and check the dependent features found in ASL),
1201 its integration with concurrency models, and features to better support direct translation of ASL
1202 pseudocode, such as exception handling.

1203 seL4 [Klein et al. 2014] uses a specification of the ARMv7 architecture [Fox and Myreen 2010] to
1204 verify binary correctness of all seL4 functions. However, this binary verification is not done for
1205 certain machine-interface functions that interact with system-level parts of the architecture, which
1206 were originally assumed correct as part of the main seL4 proof. The CertiKOS project [Gu et al.
1207 2016] presents another verified operating system, which defines a machine-model for x86 [Gu et al.
1208 2015] in Coq extended with support for devices and interrupts [Chen et al. 2016]. This machine
1209 model is based on the 32-bit x86 subset specified in CompCert [Leroy et al. 2017].

1210 Syeda and Klein [Syeda and Klein 2018] formalise an ARMv7 style memory management unit
1211 (MMU) in Isabelle/HOL, with a translation lookaside buffer and multiple levels of page tables.
1212 They are able to reason about system-level code in the presence of a TLB, including operating
1213 system context-switching. Joloboff et al [Joloboff et al. 2015; Shi 2013] develop a verified instruction
1214 set simulator using Coq for the ARMv6 architecture. They compile C code implementing each
1215 instruction using CompCert, before proving equivalence between the CompCert instruction set
1216 semantics and a model of ARMv6 extracted to Coq from the ARM architecture reference manual
1217 PDF. With ARM's release of a machine readable specification [Reid 2017], which we have used,
1218 such an extraction process is no longer necessary.

1219 The PROSPER project [Baumann et al. 2016; Guanciale et al. 2016] has extended L3 models of
1220 ARMv8 [Fox 2015] with system features sufficient to verify a virtualisation platform including
1221 secure boot and a hypervisor. This specification is based on hand-translating the required parts
1222 from the ARM architecture reference manuals. In contrast, by basing our ARMv8 model on ASL,
1223 we are able to more easily keep track of the constant revisions to the architecture, as well as cover
1224 more obscure corner cases in the architecture with improved confidence.

1225

REFERENCES

- 1226 2017. The gem5 Simulator. <http://gem5.org>.
- 1227 2017. QEMU: the FAST! processor emulator. <https://www.qemu.org/>.
- 1228 2017. The RISC-V Instruction Set Manual. Volume I: User-Level ISA; Volume II: Privileged Architecture. <https://riscv.org/specifications/>. 236 pages.
- 1229 J. Alglave, L. Maranget, S. Sarkar, and P. Sewell. 2010. Fences in Weak Memory Models. In *Proceedings of CAV 2010: the 22nd International Conference on Computer Aided Verification, LNCS 6174*. https://doi.org/10.1007/978-3-642-14295-6_25
- 1230 Jade Alglave, Luc Maranget, and Michael Tautschnig. 2014. Herding Cats: Modelling, Simulation, Testing, and Data Mining for Weak Memory. *ACM TOPLAS* 36, 2, Article 7 (July 2014), 74 pages. <https://doi.org/10.1145/2627752>
- 1231 Roberto M. Amadio, Nicholas Ayache, François Bobot, Jaap Boender, Brian Campbell, Ilias Garnier, Antoine Madet, James McKinna, Dominic P. Mulligan, Mauro Piccolo, Randy Pollack, Yann Régis-Gianas, Claudio Sacerdoti Coen, Ian Stark, and Paolo Tranquilli. 2013. Certified Complexity (CerCo). In *Foundational and Practical Aspects of Resource Analysis - Third International Workshop, FOPARA 2013, Bertinoro, Italy, August 29-31, 2013, Revised Selected Papers*. 1–18. https://doi.org/10.1007/978-3-319-12466-7_1
- 1232 Andrew W. Appel, Lennart Beringer, Robert Dockins, Josiah Dodds, Aquinas Hobor, Gordon Stewart, and Qinxiang Cao. 2017. Verified Software Toolchain. <http://vst.cs.princeton.edu/download/>.
- 1233 ARM. 2017. ARM Architecture Reference Manual. ARMv8, for ARMv8-A architecture profile. v8.2 Beta. 6354 pages.
- 1234 Christoph Baumann, Mats Näslund, Christian Gehrman, Oliver Schwarz, and Hans Thorsen. 2016. A high assurance virtualization platform for ARMv8. In *European Conference on Networks and Communications, EuCNC 2016, Athens, Greece, June 27-30, 2016*. 210–214. <https://doi.org/10.1109/EuCNC.2016.7561034>
- 1235 Aaron Bohannon, J. Nathan Foster, Benjamin C. Pierce, Alexandre Pilkiewicz, and Alan Schmitt. 2008. Boomerang: Resourceful Lenses for String Data. In *Proceedings of the 35th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL '08)*. ACM, New York, NY, USA, 407–419. <https://doi.org/10.1145/1328438.1328487>
- 1236 David Brumley, Ivan Jager, Thanassis Avgerinos, and Edward J Schwartz. 2011. BAP: A binary analysis platform. In *International Conference on Computer Aided Verification*. Springer, 463–469.
- 1237 Hao Chen, Xiongnan (Newman) Wu, Zhong Shao, Joshua Lockerman, and Ronghui Gu. 2016. Toward Compositional Verification of Interruptible OS Kernels and Device Drivers. In *Proceedings of the 37th ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI '16)*. ACM, New York, NY, USA, 431–447. <https://doi.org/10.1145/2908080.2908101>
- 1238 Adam Chlipala. 2013. The Bedrock structured programming system: combining generative metaprogramming and Hoare logic in an extensible program verifier. In *ACM SIGPLAN International Conference on Functional Programming, ICFP'13, Boston, MA, USA - September 25 - 27, 2013*. 391–402. <https://doi.org/10.1145/2500365.2500592>
- 1239 Leonardo De Moura and Nikolaj Bjørner. 2008. Z3: An Efficient SMT Solver. In *Proceedings of the Theory and Practice of Software, 14th International Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS'08/ETAPS'08)*. Springer-Verlag, Berlin, Heidelberg, 337–340. <http://dl.acm.org/citation.cfm?id=1792734.1792766>
- 1240 João Dias and Norman Ramsey. 2010. Automatically Generating Instruction Selectors Using Declarative Machine Descriptions. In *Proceedings of the 37th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL '10)*. ACM, New York, NY, USA, 403–416. <https://doi.org/10.1145/1706299.1706346>
- 1241 Joshua Dunfield and Neelakantan R. Krishnaswami. 2013. Complete and Easy Bidirectional Typechecking for Higher-Rank Polymorphism. In *International Conference on Functional Programming (ICFP)*. [arXiv:1306.6032\[cs.PL\]](https://arxiv.org/abs/1306.6032).
- 1242 Cormac Flanagan, Amr Sabry, Bruce F. Duba, and Matthias Felleisen. 1993. The Essence of Compiling with Continuations. In *Proceedings of the ACM SIGPLAN'93 Conference on Programming Language Design and Implementation (PLDI), Albuquerque, New Mexico, USA, June 23-25, 1993*, Robert Cartwright (Ed.). ACM, 237–247. <https://doi.org/10.1145/155090.155113>
- 1243 Shaked Flur, Kathryn E. Gray, Christopher Pulte, Susmit Sarkar, Ali Sezgin, Luc Maranget, Will Deacon, and Peter Sewell. 2016. Modelling the ARMv8 Architecture, Operationally: Concurrency and ISA. In *Proceedings of POPL: the 43rd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*. <https://doi.org/10.1145/2837614.2837615>
- 1244 Shaked Flur, Susmit Sarkar, Christopher Pulte, Kyndylan Nienhuis, Luc Maranget, Kathryn E. Gray, Ali Sezgin, Mark Batty, and Peter Sewell. 2017. Mixed-size Concurrency: ARM, POWER, C/C++11, and SC. In *The 44th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, Paris, France*. 429–442. <https://doi.org/10.1145/3009837.3009839>
- 1245 Anthony C. J. Fox. 2012. Directions in ISA Specification. In *Interactive Theorem Proving - Third International Conference, ITP 2012, Princeton, NJ, USA, August 13-15, 2012. Proceedings*. 338–344. https://doi.org/10.1007/978-3-642-32347-8_23
- 1246 Anthony C. J. Fox. 2015. Improved Tool Support for Machine-Code Decompilation in HOL4. In *Interactive Theorem Proving - 6th International Conference, ITP 2015, Nanjing, China, August 24-27, 2015, Proceedings*, Christian Urban and Xingyuan Zhang (Eds.), Vol. 9236. Springer, 187–202. https://doi.org/10.1007/978-3-319-22102-1_12
- 1247 Anthony C. J. Fox and Magnus O. Myreen. 2010. A Trustworthy Monadic Formalization of the ARMv7 Instruction Set Architecture. In *Interactive Theorem Proving, First International Conference, ITP 2010, Edinburgh, UK, July 11-14, 2010. Proceedings*. 243–258. https://doi.org/10.1007/978-3-642-14052-5_18

- 1275 Anthony C. J. Fox, Magnus O. Myreen, Yong Kiam Tan, and Ramana Kumar. 2017. Verified compilation of CakeML to
1276 multiple machine-code targets. In *Proceedings of the 6th ACM SIGPLAN Conference on Certified Programs and Proofs, CPP*
1277 *2017, Paris, France, January 16–17, 2017*. 125–137. <https://doi.org/10.1145/3018610.3018621>
- 1278 Shilpi Goel, Warren A. Hunt Jr., and Matt Kaufmann. 2017. Engineering a Formal, Executable x86 ISA Simulator for Software
1279 Verification. In *Provably Correct Systems*. 173–209. https://doi.org/10.1007/978-3-319-48628-4_8
- 1280 Kathryn E. Gray, Gabriel Kerneis, Dominic Mulligan, Christopher Pulte, Susmit Sarkar, and Peter Sewell. 2015. An integrated
1281 concurrency and core-ISA architectural envelope definition, and test oracle, for IBM POWER multiprocessors. In
1282 *Proc. MICRO-48, the 48th Annual IEEE/ACM International Symposium on Microarchitecture*. <https://doi.org/10.1145/2830772.2830775>
- 1283 Ronghui Gu, Jérémie Koenig, Tahina Ramananandro, Zhong Shao, Xiongnan (Newman) Wu, Shu-Chun Weng, Haozhong
1284 Zhang, and Yu Guo. 2015. Deep Specifications and Certified Abstraction Layers. In *Proceedings of the 42nd Annual ACM*
1285 *SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL '15)*. ACM, New York, NY, USA, 595–608.
<https://doi.org/10.1145/2676726.2676975>
- 1286 Ronghui Gu, Zhong Shao, Hao Chen, Xiongnan (Newman) Wu, Jieung Kim, Vilhelm Sjöberg, and David Costanzo. 2016.
1287 CertiKOS: An Extensible Architecture for Building Certified Concurrent OS Kernels. In *12th USENIX Symposium on*
1288 *Operating Systems Design and Implementation, OSDI 2016, Savannah, GA, USA, November 2–4, 2016*. 653–669. <https://www.usenix.org/conference/osdi16/technical-sessions/presentation/gu>
- 1289 Roberto Guanciale, Hamed Nemati, Mads Dam, and Christoph Baumann. 2016. Provably secure memory isolation for Linux
1290 on ARM. *Journal of Computer Security* 24, 6 (2016), 793–837. <https://doi.org/10.3233/JCS-160558>
- 1291 Intel Corporation. 2017. Intel 64 and IA-32 Architectures Software Developer’s Manual. Combined Volumes: 1, 2A, 2B, 2C,
1292 2D, 3A, 3B, 3C, 3D, and 4. <https://software.intel.com/en-us/articles/intel-sdm>. 325462-063US. 4744 pages.
- 1293 Jonas B. Jensen, Nick Benton, and Andrew Kennedy. 2013. High-level Separation Logic for Low-level Code. In *Proceedings of*
1294 *the 40th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL '13)*. ACM, New York,
NY, USA, 301–314. <https://doi.org/10.1145/2429069.2429105>
- 1295 Vania Joloboff, Jean-François Monin, and Xiaomu Shi. 2015. Towards Verified Faithful Simulation. In *Dependable Software*
1296 *Engineering: Theories, Tools, and Applications - First International Symposium, SETTA 2015, Nanjing, China, November 4–6,*
1297 *2015, Proceedings (LNCS)*, Xuandong Li, Zhiming Liu, and Wang Yi (Eds.). Springer, 105–119. https://doi.org/10.1007/978-3-319-25942-0_7
- 1298 Jacques-Henri Jourdan, Vincent Laporte, Sandrine Blazy, Xavier Leroy, and David Pichardie. 2015. A Formally-Verified
1299 C Static Analyzer. In *Proceedings of the 42nd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming*
1300 *Languages, POPL 2015, Mumbai, India, January 15–17, 2015*. 247–259. <https://doi.org/10.1145/2676726.2676966>
- 1301 Andrew Kennedy, Nick Benton, Jonas Braband Jensen, and Pierre-Évariste Dagand. 2013. Coq: the world’s best macro
1302 assembler?. In *15th International Symposium on Principles and Practice of Declarative Programming, PDP '13, Madrid,*
1303 *Spain, September 16–18, 2013*. 13–24. <https://doi.org/10.1145/2505879.2505897>
- 1304 Gerwin Klein, June Andronick, Kevin Elphinstone, Toby Murray, Thomas Sewell, Rafal Kolanski, and Gernot Heiser.
1305 2014. Comprehensive Formal Verification of an OS Microkernel. *ACM TOCS* 32, 1 (Feb. 2014), 2:1–2:70. <https://doi.org/10.1145/2560537>
- 1306 Ramana Kumar, Magnus O. Myreen, Michael Norrish, and Scott Owens. 2014. CakeML: A Verified Implementation of ML.
1307 In *Proceedings of the 41st ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL '14)*. ACM,
1308 New York, NY, USA, 179–191. <https://doi.org/10.1145/2535838.2535841>
- 1309 Dirk Leinenbach and Thomas Santen. 2009. Verifying the Microsoft Hyper-V Hypervisor with VCC. In *FM 2009: Formal*
1310 *Methods, Second World Congress, Eindhoven, The Netherlands, November 2–6, 2009. Proceedings*. 806–809. https://doi.org/10.1007/978-3-642-05089-3_51
- 1311 Xavier Leroy. 2009. A formally verified compiler back-end. *J. Automated Reasoning* 43, 4 (2009), 363–446. <https://doi.org/10.1007/s10817-009-9155-4>
- 1312 Xavier Leroy et al. 2017. CompCert 3.1. <http://compcert.inria.fr/>.
- 1313 Junghee Lim and Thomas W. Reps. 2013. TSL: A System for Generating Abstract Interpreters and its Application to
1314 Machine-Code Analysis. *ACM TOPLAS* 35, 1 (2013), 4:1–4:59. <https://doi.org/10.1145/2450136.2450139>
- 1315 Prabhat Misra and Nikil Dutt (Eds.). 2008. *Processor Description Languages*. Morgan Kaufmann.
- 1316 Greg Morrisett, Gang Tan, Joseph Tassarotti, Jean-Baptiste Tristan, and Edward Gan. 2012. RockSalt: better, faster, stronger
1317 SFI for the x86. In *ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI '12, Beijing,*
1318 *China - June 11 - 16, 2012*. 395–404. <https://doi.org/10.1145/2254064.2254111>
- 1319 Dominic P. Mulligan, Scott Owens, Kathryn E. Gray, Tom Ridge, and Peter Sewell. 2014. Lem: reusable engineering of real-
1320 world semantics. In *Proceedings of ICFP 2014: the 19th ACM SIGPLAN International Conference on Functional Programming*.
1321 175–188. <https://doi.org/10.1145/2628136.2628143>
- 1322 Magnus Oskar Myreen. 2009. *Formal verification of machine-code programs*. Ph.D. Dissertation. University of Cambridge,
1323 UK. <http://ethos.bl.uk/OrderDetails.do?uin=uk.bl.ethos.611450>

- 1324 Nicholas Nethercote and Julian Seward. 2007. Valgrind: A Framework for Heavyweight Dynamic Binary Instrumentation.
 1325 In *Proceedings of the 28th ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI '07)*.
 1326 ACM, New York, NY, USA, 89–100. <https://doi.org/10.1145/1250734.1250746>
- 1327 William Pugh. 1991. The Omega test: a fast and practical integer programming algorithm for dependence analysis. In
 1328 *Proceedings Supercomputing '91, Albuquerque, NM, USA, November 18-22, 1991*, Joanne L. Martin (Ed.). ACM, 4–13.
 1329 <https://doi.org/10.1145/125826.125848>
- 1330 Christopher Pulte, Shaked Flur, Will Deacon, Jon French, Susmit Sarkar, and Peter Sewell. 2018. Simplifying ARM Concur-
 1331 rency: Multicopy-atomic Axiomatic and Operational Models for ARMv8.
- 1332 Alastair Reid. 2016. Trustworthy Specifications of ARM v8-A and v8-M System Level Architecture. In *FMCAD 2016*. 161–168.
 1333 <https://alastairreid.github.io/papers/fmcad2016-trustworthy.pdf>
- 1334 Alastair Reid. 2017. ARM Releases Machine Readable Architecture Specification. [https://alastairreid.github.io/
 1335 ARM-v8a-xml-release/](https://alastairreid.github.io/ARM-v8a-xml-release/).
- 1336 Alastair Reid, Rick Chen, Anastasios Deligiannis, David Gilday, David Hoyes, Will Keen, Ashan Pathirane, Owen Shepherd,
 1337 Peter Vrabel, and Ali Zaidi. 2016. End-to-End Verification of Processors with ISA-Formal. In *Computer Aided Verification
 1338 - 28th International Conference, CAV 2016, Toronto, ON, Canada, July 17-23, 2016, Proceedings, Part II (Lecture Notes in
 1339 Computer Science)*, Swarat Chaudhuri and Azadeh Farzan (Eds.), Vol. 9780. Springer, 42–58. [https://doi.org/10.1007/
 1340 978-3-319-41540-6_3](https://doi.org/10.1007/978-3-319-41540-6_3)
- 1341 Patrick M. Rondon, Ming Kawaguci, and Ranjit Jhala. 2008. Liquid Types. In *Proceedings of the 29th ACM SIGPLAN
 1342 Conference on Programming Language Design and Implementation (PLDI '08)*. ACM, New York, NY, USA, 159–169.
 1343 <https://doi.org/10.1145/1375581.1375602>
- 1344 Susmit Sarkar, Peter Sewell, Jade Alglave, Luc Maranget, and Derek Williams. 2011. Understanding POWER Multiprocessors.
 1345 In *Proceedings of PLDI 2011: the 32nd ACM SIGPLAN conference on Programming Language Design and Implementation*.
 1346 175–186. <https://doi.org/10.1145/1993498.1993520>
- 1347 Peter Sewell, Susmit Sarkar, Scott Owens, Francesco Zappa Nardelli, and Magnus O. Myreen. 2010. x86-TSO: A Rigorous
 1348 and Usable Programmer’s Model for x86 Multiprocessors. *Comm. of the ACM* 53, 7 (July 2010), 89–97. [https://doi.org/10.
 1349 1145/1785414.1785443](https://doi.org/10.1145/1785414.1785443)
- 1350 Xiaomu Shi. 2013. *Certification of an Instruction Set Simulator*. Theses. Université de Grenoble. [https://tel.archives-ouvertes.
 1351 fr/tel-00937524](https://tel.archives-ouvertes.fr/tel-00937524)
- 1352 Yan Shoshitaishvili, Ruoyu Wang, Christopher Salls, Nick Stephens, Mario Polino, Audrey Dutcher, John Grosen, Siji Feng,
 1353 Christophe Hauser, Christopher Kruegel, and Giovanni Vigna. 2016. SoK: (State of) The Art of War: Offensive Techniques
 1354 in Binary Analysis. In *IEEE Symposium on Security and Privacy*.
- 1355 Antal Spector-Zabusky, Joachim Breitner, Christine Rizkallah, and Stephanie Weirich. 2018. Total Haskell is Reasonable
 1356 Coq. In *Proceedings of the 7th ACM SIGPLAN International Conference on Certified Programs and Proofs (CPP 2018)*. ACM,
 1357 New York, NY, USA, 14–27. <https://doi.org/10.1145/3167092>
- 1358 Hira Taqdees Syeda and Gerwin Klein. 2018. Program Verification in the Presence of Cached Address Translation. In *ITP
 1359 2018*, Jeremy Avigad and Assia Mahboubi (Eds.). Springer, 542–559.
- 1360 Yong Kiam Tan, Magnus O. Myreen, Ramana Kumar, Anthony C. J. Fox, Scott Owens, and Michael Norrish. 2016. A new
 1361 verified compiler backend for CakeML. In *Proceedings of the 21st ACM SIGPLAN International Conference on Functional
 1362 Programming, ICFP 2016, Nara, Japan, September 18-22, 2016*. 60–73. <https://doi.org/10.1145/2951913.2951924>
- 1363 Jaroslav Ševčík, Viktor Vafeiadis, Francesco Zappa Nardelli, Suresh Jagannathan, and Peter Sewell. 2013. CompCertTSO: A
 1364 Verified Compiler for Relaxed-Memory Concurrency. *J. ACM* 60, 3, Article 22 (June 2013), 50 pages. [https://doi.org/10.
 1365 1145/2487241.2487248](https://doi.org/10.1145/2487241.2487248)
- 1366 Robert N. M. Watson, Peter G. Neumann, Jonathan Woodruff, Michael Roe, Jonathan Anderson, John Baldwin, David
 1367 Chisnall, Brooks Davis, Alexandre Joannou, Ben Laurie, Simon W. Moore, Steven J. Murdoch, Robert Norton, Stacey Son,
 1368 and Hongyan Xia. 2017. *Capability Hardware Enhanced RISC Instructions: ChERI Instruction-Set Architecture (Version
 1369 6)*. Technical Report UCAM-CL-TR-907. University of Cambridge, Computer Laboratory. [http://www.cl.cam.ac.uk/
 1370 techreports/UCAM-CL-TR-907.pdf](http://www.cl.cam.ac.uk/techreports/UCAM-CL-TR-907.pdf)
- 1371 Robert N. M. Watson, Jonathan Woodruff, Peter G. Neumann, Simon W. Moore, Jonathan Anderson, David Chisnall, Nirav H.
 1372 Dave, Brooks Davis, Khilan Gudka, Ben Laurie, Steven J. Murdoch, Robert Norton, Michael Roe, Stacey D. Son, and Munraj
 1373 Vadera. 2015. ChERI: A Hybrid Capability-System Architecture for Scalable Software Compartmentalization. In *2015 IEEE
 1374 Symposium on Security and Privacy, SP 2015, San Jose, CA, USA, May 17-21, 2015*. 20–37. <https://doi.org/10.1109/SP.2015.9>
- 1375 Jonathan Woodruff, Robert N.M. Watson, David Chisnall, Simon W. Moore, Jonathan Anderson, Brooks Davis, Ben Laurie,
 1376 Peter G. Neumann, Robert Norton, and Michael Roe. 2014. The ChERI capability model: revisiting RISC in an age of risk.
 1377 In *ISCA '14: Proceeding of the 41st annual international symposium on Computer architecture*. IEEE Press, Piscataway, NJ,
 1378 USA, 457–468. <https://doi.org/10.1145/2678373.2665740>
- 1379 Hongwei Xi. 2007. Dependent ML An approach to practical programming with dependent types. *J. Funct. Program.* 17, 2
 1380 (2007), 215–286. <https://doi.org/10.1017/S0956796806006216>