

An Axiomatic Basis for Computer Programming on Relaxed Hardware Architectures: The AxSL Logics

ZONGYUAN LIU, Aarhus University, Denmark
ANGUS HAMMOND, University of Cambridge, UK
THIBAUT PÉRAMI, University of Cambridge, UK
PETER SEWELL, University of Cambridge, UK
LARS BIRKEDAL, Aarhus University, Denmark
JEAN PICHON-PHARABOD, Aarhus University, Denmark

Very relaxed concurrency memory models, like those of the Arm-A, RISC-V, and IBM Power hardware architectures, underpin much of computing but break a fundamental intuition about programs, namely that syntactic program order and the reads-from relation always both induce order in the execution. Instead, out-of-order execution is allowed except where prevented by certain pairwise dependencies, barriers, or other synchronisation. This means that there is no notion of the ‘current’ state of the program, making it challenging to design (and prove sound) syntax-directed, modular reasoning methods like Hoare logics, as usable resources cannot implicitly flow from one program point to the next.

We present AxSL, a family of separation logics for relaxed hardware memory models, and instantiate it on sequential consistency and on the Arm-A memory model. The Arm-A instance captures the fine-grained reasoning underpinning the low-overhead synchronisation idioms used by high-performance systems code. We mechanise AxSL in the Iris separation logic framework, illustrate it on key examples, and prove it sound with respect to the axiomatic memory model of Arm-A.

By instantiating AxSL on different memory models, we demonstrate the generality of our approach, and show that it is largely generic in the axiomatic model and in the instruction-set semantics, offering a potential way forward for compositional reasoning for other models, and for the combination of production concurrency models and full-scale ISAs.

CCS Concepts: • **Theory of computation** → **Separation logic**; • **Computer systems organization** → *Multicore architectures*.

Additional Key Words and Phrases: relaxed memory models, program logic, separation logic, Arm, Iris

1 Introduction

Systems code, such as operating system and hypervisor kernel code, is a prime target for software verification, being security-critical yet relatively small. However, it is highly concurrent, which raises two questions: What model to verify it above? And what verification theory to use? For example, the Arm-A architecture is used in essentially all mobile devices, and its base (“user”) relaxed concurrency model is now reasonably well-understood and stable [Arm Ltd. 2023, Ch.B2],[Alglave et al. 2021, 2014; Deacon 2016; Flur et al. 2017; Pulte et al. 2018]. However, there is little program verification theory or tooling that applies directly to Arm-A, nor to similarly relaxed architectures.

In this paper, we develop a family of separation logics that can be instantiated on relaxed hardware memory models, and yet is expressive, supporting local reasoning with higher-order ghost state and invariants, and mechanised in Rocq, using the Iris program logic framework [Jung et al. 2018, 2015]. We then instantiate this

Authors’ Contact Information: Zongyuan Liu, zy.liu@cs.au.dk, Aarhus University, Aarhus, Denmark; Angus Hammond, Angus.Hammond@cl.cam.ac.uk, University of Cambridge, Cambridge, UK; Thibaut Pérami, Thibaut.Perami@cl.cam.ac.uk, University of Cambridge, Cambridge, UK; Peter Sewell, Peter.Sewell@cl.cam.ac.uk, University of Cambridge, Cambridge, UK; Lars Birkedal, birkedal@cs.au.dk, Aarhus University, Aarhus, Denmark; Jean Pichon-Pharabod, jean.pichon@cs.au.dk, Aarhus University, Aarhus, Denmark.



This work is licensed under a [Creative Commons Attribution 4.0 International License](https://creativecommons.org/licenses/by/4.0/).

logic on the Arm-A “user” concurrency model, which is particularly challenging for program-logic reasoning because it (like RISC-V and IBM Power, but unlike x86) permits load-store reordering, as in the classic “load buffering” LB shape of Fig. 1. This means that the union of program order (po) and the reads-from relation (rf) is not guaranteed to be acyclic — but for compositional reasoning, one wants to attach assertions to particular program points, and program logics usually rely on the strength of program order captured by that acyclicity; they let resources implicitly flow in the proof context from one program point to the next. Previous program logics have either assumed $po \cup rf$ acyclic (which requires extra barriers), e.g. FSL++ [Doko 2021; Doko and Vafeiadis 2017], GPS [Turon et al. 2014], and iRC11 [Dang et al. 2020], or lack ghost state, e.g. FSL [Doko and Vafeiadis 2016], which makes the logic substantially less expressive and more awkward to use, or give extremely weak guarantees for non-synchronised reads, e.g. RSL [Vafeiadis and Narayan 2013]. Owicki-Gries [Owicki and Gries 1976] logics for C11 by Dalvandi et al. [2020, 2022] and Wright et al. [2021, 2023] support $po \cup rf$ cycles but do not feature full ghost state, nor thread-local modular reasoning. The Lace logic [Bornat et al. 2015] targeted relaxed architectural models but lacked a proof of soundness, and the Ogre and Pythia logic [Alglave and Cousot 2017] is a refinement of Owicki-Gries that is parameterised by (and sound for) a range of relaxed models, but (like Owicki-Gries) lacks thread-local reasoning.



Fig. 1. LB+pos

In contrast to those previous program logics, to allow sound usage of ghost resources even in the presence of LB, we prevent implicit flow of usable resources between program points along po, allowing it only when actual synchronisation is present — for example, for the Arm-A memory model, along the ordered-before ordering (ob). Different relaxed hardware architectures expose different combinations of ways to impose such synchronisation: address dependencies, release writes, etc. We allow explicit reasoning about those if need be, by exposing the structure of the axiomatic model, letting one reason about the low-cost ordering that the architecture under consideration guarantees from various forms of dependency (RSL, FSL, FSL++, GPS, and iGPS are all for C11 or RC11, without dependencies).

Stepping back, why would one want to reason directly above an architecture concurrency model? After all, high-level language concurrency models, e.g. C/C++11 [Batty et al. 2011; Boehm and Adve 2008] and the Linux kernel memory model, LKMM [Alglave et al. 2018; McKenney et al. 2020], were designed to obviate the need to program and reason about specific underlying architectures, with extensive work on the correctness of their compilation schemes [Batty et al. 2012; Lahav et al. 2017; Manerkar et al. 2016; Sarkar et al. 2012], and one would not envisage manual proof about large bodies of assembly code. There are three main reasons.

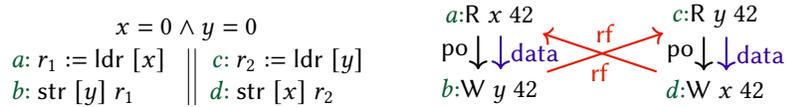


Fig. 2. LB+datas

First, those C/C++ language-level models are fundamentally flawed for highly relaxed code because of the out-of-thin-air problem [Becker 2011, §23.9p9] [Batty et al. 2015]: they allow arbitrary values to be created, e.g. for

the Fig. 2 LB+datas shape of relaxed atomic accesses and source-language data dependencies. Thin-air values are not believed to arise for conventional compilers and hardware, but it has proven challenging to define tractable semantics that exclude them while remaining sound w.r.t. conventional compiler and hardware optimisations – especially compiler dependency removal and hardware load-store reordering. The LKMM forbids thin-air outcomes by assuming some dependencies are respected, and in specific coding idioms they often are, but in general they can be removed by conventional compiler optimisations. There have been many attempts to solve this problem [Chakraborty and Vafeiadis 2019; Jeffrey and Riely 2016; Kang et al. 2017; Lee et al. 2020; Paviotti et al. 2020; Pichon-Pharabod and Sewell 2016], but so far none have been adopted – so we simply do not yet have any high-level language semantics suitable for reasoning about deployed highly relaxed code. In contrast, architecture concurrency models for Arm-A, x86, RISC-V, IBM Power, and others, are now well-established [Alglave et al. 2021, 2014; Arm Ltd. 2023; Deacon 2016; Flur et al. 2017; Owens et al. 2009; Pulte et al. 2018; Sarkar et al. 2011, 2009; Waterman and Asanović 2019], and do not suffer from the thin-air problem: these architectures guarantee respect for certain syntactic dependencies, ruling out thin-air. These architectural models thus give us a solid foundation that we can reason above.

Second, ultimately, the machine-code binary is what runs – and therefore one wants to verify down to the (concurrent) machine semantics, even if the bulk of one’s source-language verification is at the C level or above. There are several possible approaches to this: for example, one might have a source language with more restricted concurrency (without relaxed accesses), and then some verified compilation result down to the machine semantics [Cho et al. 2022; Tao et al. 2021]. But production systems-code in practice does use relaxed accesses for performance, and hence reasoning about them is an important problem. We thus aim here to first understand how to reason directly about the binary, where we have a good underlying model; future work can then use this as the basis for verified compilation or other verification approaches for higher-level code.

Third, systems code relies, in small but crucial parts, on assembly which is not C-language expressible – e.g. for particular barriers, and for management of systems features of the underlying architecture (instruction and data cache management [Simner et al. 2020], virtual memory [Simner et al. 2022][Arm Ltd. 2023, B2.3], exceptions, etc.). We do not cover systems semantics here, but our approach is designed to generalise to it.

Contributions. We develop a family of separation logics for relaxed hardware architectures, AxSL, that is expressive, supporting reasoning with higher-order ghost state and invariants, and mechanised in Rocq, using the Iris program logic framework.

We then instantiate AxSL on two memory models: sequential consistency, and the Arm-A concurrency model. For both, we use an idealised instruction set architecture (ISA), but our approach is designed to generalise: our idealised ISA semantics and base logic are defined above the microinstructions of the Sail “outcome” interface [Gray et al. 2015][Pulte et al. 2018, §6.1][Pulte 2018, §2.3], so they should generalise straightforwardly to the full ISA of Arm-A or RISC-V. For the concurrency model, our approach is largely generic in the structure of the axiomatic model, so this work offers a path towards similar logics for other architecture axiomatic models (e.g. the RISC-V “user” model, which is similar to that of Arm-A), or, more speculatively, to extensions covering systems semantics, as has been developed for example for Arm [Alglave et al. 2024; Simner et al. 2025, 2022, 2020]. Moreover, both the Arm-A architecture reference manual and RISC-V specify their concurrency architecture in this axiomatic style [Arm Ltd. 2023, Ch.B2] and are actively maintained and occasionally changed, so (while semantics for new features have been developed in multiple styles), it is desirable to be able to track the reference-manual version with minimal effort.

Plan. We describe the program-logic and relaxed-memory context in §2. We explain the key ideas of our logic informally in §3. In §4, we describe the two languages we consider, and how we give their semantics in our novel ‘opax’ style that makes it possible to build expressive logics featuring higher-order ghost state above axiomatic memory models. We present one language for SC, combining a simplified assembly language with sequential

consistency; and one for Arm-A, combining a simplified assembly language featuring dependencies with the real LB-permitting Arm-A axiomatic concurrency model. To avoid adding overwhelming complexity to an already complex topic, we only consider the “user” Arm-A memory model of 2018 [Pulte et al. 2018], and not more recent extensions and changes: no mixed-size accesses, no instruction fetching, no virtual memory, and no pick dependencies, although these extensions are all in the shape that our approach supports. In §5, we describe the rules of our AxSL logic and exercise them on small, representative examples. We do this in three stages: first we present how to deal with axiomatic memory models ignoring relaxed memory, then we show how to structure the logic to deal with a relaxed memory model but in the simple setting of sequential consistency, and finally we deal with an actual relaxed memory model, namely that of Arm-A. In §6, we define the model of AxSL in Iris, following the same three stages, and present our non-standard definition of weakest precondition. In §7, we present our non-standard proof of adequacy of AxSL in Iris. In §8, we discuss some technical aspects and limitations of our work. We discuss related work in §9, and how our work can be used and extended further in §10. The Arm-A instance of AxSL, its soundness and the examples are formalised in Rocq using the Iris separation logic framework; the full development is available at <https://github.com/logsem/AxSL>.

Difference with the original paper. This article is an extended version of the original paper presented at POPL 2024 [Hammond et al. 2024]. In particular, it makes our contributions more accessible, especially to those who are less familiar with relaxed memory and the memory model of Arm-A, it elaborates the definitions to be self-contained, and it makes technical improvements to the proof technique. In detail:

- We introduce our ‘opax’ type of semantics using a simple language with a simpler memory model (namely sequential consistency) in §4.3.3.
- We explain the novel ideas of AxSL^{Arm} in that simpler setting, building two logics for that simple SC language: AxSL^{SC} and $\text{AxSL}^{\text{SCExt}}$. These two simpler logics serve as explanatory steps when building up the syntax and the semantic model of the AxSL^{Arm} .
 - We show how to define a first straightforward logic, AxSL^{SC} , on top of our novel ‘opax’ style of semantics in §5.2, and how to define a semantic model for it in §6.2.
 - We then show how to define a second, more elaborate logic, $\text{AxSL}^{\text{SCExt}}$, that uses the ideas that make AxSL^{Arm} work in the setting of relaxed memory, but still in the simple setting of sequential consistency in §5.3, and present its semantic model in §6.3.
- We give a self-contained presentation of the definitions of the model of AxSL^{Arm} , using precise definitions that were omitted in the original paper because of space constraints, in §6.4.
- We expand the explanation of various technical definitions and proofs, and add illustrations to make the technical material more accessible.
- We refine the model of AxSL^{Arm} , leading to some technical improvements that we elaborate on in §8.1, in particular a better proof of adequacy.

Moreover, by demonstrating the ideas of AxSL^{Arm} on a different (albeit simple) memory model, we have shown that our novel approach to defining semantics and program logics generalises.

2 Context: Program Logics and Relaxed Concurrency

Early work on program verification, in a sequential setting, could assume the existence of a simple program state of memory values, updated by each instruction, and program proof could be done by annotating a flowchart (as per Morris and Jones [1984]; Turing [1949] and Floyd [1967]), or syntactic program points (as per Naur [1966] and Hoare [1969]), with assertions on that state. In this setting, a fact about a part of the state untouched by some instruction remains true (and usable for program proof) after the instruction, though managing such framing had to be done manually. The first separation logics, of Reynolds [2002], and O’Hearn et al. [2001], refined this view with a separating conjunction, allowing assertions to express ownership of some part of such a state, with

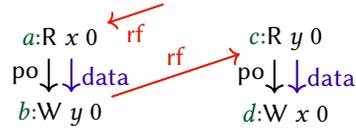


Fig. 3. An SC execution of LB+datas

an explicit frame rule. Simple concurrent separation logics, e.g. CSL [Brookes and O’Hearn 2016], are broadly similar except that ownership of parts of the state can be transferred at lock acquire and release points: facts about owned parts of the state remain true from one program point to the next, except where the state they mention is explicitly modified by the intervening instruction.

In a relaxed-memory concurrent setting, however, the notion of a program state acted on by all threads in some global interleaving is usually much more involved: threads do not execute in-order, and different threads can observe events in incompatible orders. To capture this, the underlying semantics have quite different forms to classical sequential or sequentially consistent concurrent semantics. Two styles of semantics for architectural relaxed-memory concurrency are common: abstract-microarchitectural operational models explain how the allowed observable behaviour arises from explicit speculative execution and event propagation, with roll-back when speculation turns out to violate some constraint, e.g. [Higham et al. 2007; Owens et al. 2009; Pulte et al. 2018; Sarkar et al. 2011], while axiomatic models define the allowed observable behaviour more concisely as predicates on candidate complete execution graphs, e.g. [Alglave et al. 2010, 2014; Gharachorloo 1995; Kohli et al. 1993], but do not straightforwardly support the incremental construction of valid executions. A third, “Promising”, style is, very roughly, intermediate between the two [Pulte et al. 2019]. All are challenging to work with, in different ways, as we discuss in §3.2.

We base the current work on axiomatic models. In these, a program gives rise to a large set of candidate complete execution graphs, each with a function from event IDs to events, a program order relation over event IDs (po) within each thread, and various other base relations. An axiomatic concurrency model imposes constraints on derived relations defined in terms of these. For example, in Arm-A, which we will use as our main case study, the model of Fig. 4 defines an *observed before* (ob) relation that captures synchronisation, and imposes constraints on those, in particular that ob is acyclic. The semantics of a program is the set of all candidate complete execution graphs that satisfy those properties and are consistent with the intra-instruction semantics. For example, the candidate execution for LB+pos in Fig. 1 is allowed by the Arm-A axiomatic model because the plain po relation, between reads and writes to different addresses, is not included in the ordered-before ob that is required to be acyclic (or in the internal or atomicity requirements). The candidate execution for LB+datas in Fig. 2 is forbidden in Arm-A because the intra-thread syntactic data dependencies create data edges, which are included in the Arm-A *locally ordered before* lob relation, and that and the inter-thread reads-from relation rfe are both contained in ob . In contrast, the candidate execution for LB+datas in which a reads from (rf) the initial state in Fig. 3 is allowed.

For some, relatively simple, forms of relaxed concurrency, one can adapt separation logic relatively straightforwardly (e.g., RSL, GPS). For example, rely/guarantee reasoning with acquire/release reads and writes lets one do thread-modular proofs, in which a thread might gain some resource at an acquire read, manipulate it freely, and then pass it on with a release write – with the resource still persisting from one program point to the next between those points (except where explicitly modified by this thread), as a read-acquire is ordered with all po -successors, a write-release with all po -predecessors, and no other access is allowed to race with non-release/acquire accesses.

```

(* Coherence-after *)
let ca = fr | co
(* Observed-by *)
let obs = rfe | fre | coe
(* Dependency-ordered-before *)
let dob = addr | data
| ctrl; [W]
| (ctrl | (addr; po)); [ISB]; po; [R]
| addr; po; [W]
| (ctrl | data); coi
| (addr | data); rfi
(* Atomic-ordered-before *)
let aob = rmw
| [range(rmw)]; rfi; [A | Q]
(* Barrier-ordered-before *)
let bob = po; [dmb.full]; po | [L]; po; [A]
| [R]; po; [dmb.ld]; po
| [A | Q]; po
| [W]; po; [dmb.st]; po; [W] | po; [L]
| po; [L]; coi
(* Locally ordered-before *)
let lob = dob | aob | bob
(* Ordered-before *)
let ob = (obs | lob)+
(* Internal visibility requirement *)
acyclic po-loc | ca | rf
(* External visibility requirement *)
irreflexive ob
(* Atomicity requirement *)
empty rmw & (fre; coe) as atomic

```

Fig. 4. Arm-A axiomatic model by Deacon [Pulte et al. 2018] (with lob separated out, following later Arm models [Alglave et al. 2021]), in herd’s cat syntax [Alglave et al. 2014] for relational algebra. Here $|$, $\&$, $;$, and $+$ are relational union, intersection, composition, and transitive closure; $[W]$, $[R]$, $[L]$, $[A]$ and $[Q]$ are the identity relations over all write, read, release, acquire and acquirePC events; $[ISB]$, $[dmb.full]$, $[dmb.st]$, $[dmb.ld]$ are the identity on those barrier events; $addr$, $data$, and $ctrl$ are the syntactic dependency-relation subsets of program order po ; $po-loc$ relates same-address memory accesses in po ; co is coherence over writes; the derived fr relates reads to coherence successors of the write they read from; rmw is the successful read/write-exclusive pairs; and the rf , co , and fr relations are subdivided into their “internal” (same-thread) and “external” (different-thread) parts, suffixed i and e respectively. The main “axiom” requires that ordered-before (ob) is irreflexive.

In these logics, the effect of reading from a shared variable on the thread’s logical state is accounted for thread-locally by relying on a protocol or invariant to abstract the possible actions of other threads. The protocol

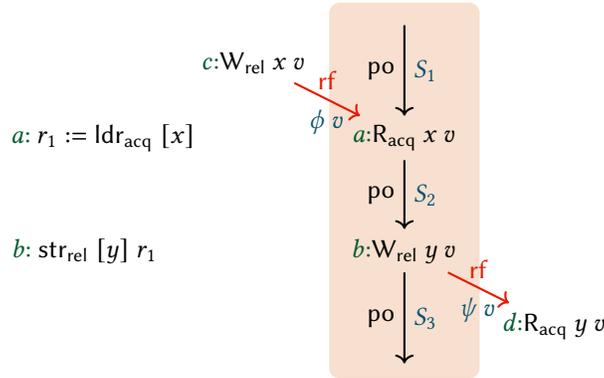


Fig. 5. Thread-local part of a candidate execution, annotated with the logical resources in blue flowing on edges. The thread program is on the left. Resources $S_{1,2,3}$ are those in hand at each program point, and the protocol specifies the resources ϕv and ψv passed along the release-acquire edges for x and y .

constrains what logical resources are transferred when accessing shared variables. In a candidate execution, one can see this as annotating the incoming (to reads) and outgoing (from writes) reads-from edges, for the part of the graph for each thread, with the resources that get transferred along them (Fig. 5).

In this view, as described for RSL, the events of the execution graph act following *flow implications*: “the annotation is locally valid around that action [when] basically the sum of the annotated heaps on the incoming edges should equal the sum of the annotated heaps on the outgoing edges, modulo the effect of [the] action”.

FSL generalises RSL to reason about C11’s release and acquire fences, but its assertions are still persistently freely usable along po , so they have to choose between soundness in the presence of load buffering (FSL), support for ghost state, at the cost of requiring $po \cup rf$ acyclic (FSL++), and non-thread-modular reasoning (Owicki-Gries). We describe these and related logics in more detail in §9.

3 Key Ideas

3.1 The First Problem: Relaxed Thread-local Ordering

The biggest challenge for reasoning about the more relaxed behaviour of mainstream (non-TSO) relaxed architectures, including Arm-A, RISC-V, and IBM Power, arises from the fact that they all permit out-of-order execution of program-ordered loads and stores, except where there is some dependency or barrier. This means that a resource gained on a load *cannot* be deemed to implicitly persist through to any program-order-later store where it might be passed on. For example, consider a thread consisting of two interleaved copies of the left thread of LB+datas (operating on disjoint addresses), as in Fig. 6. The data dependencies order a with c , and b with d , but that is all the ordering we get. In particular, *nothing* orders the last read b , before the first write c – in contrast to the release/acquire case.

The Arm-A axiomatic model’s locally-ordered-before (lob) relation specifies what thread-local ordering is respected, as introduced by barriers, synchronising accesses (store release, acquire reads, etc.), and register-to-register dependencies. All but the strongest barriers and synchronising accesses impose only a partial ordering and allow some intra-thread concurrency. In particular, some register dependencies merely impose a pairwise ordering of events; as such, these dependencies are particularly cheap, and are one of the motivations to directly write assembly for high-performance code, for example in the Linux kernel’s pervasive RCU library.

Our first key idea is that by attaching resources only to locally-ordered-before edges, rather than all of program order, we can make a sound logic even for relaxed architectures exhibiting load buffering and intra-thread concurrency. However, for practical and compositional reasoning, we want to annotate a program text, not the large set of its candidate executions. Moreover, to identify when ordering will arise from register dependencies to program-order-later events, it suffices to keep track of the *source(s)* of each register value. Concretely, a load a into register r of a value v , from some location following a protocol ϕ , will give us

$$(r \mapsto v@\{a\}) * (a \rightsquigarrow (\phi v))$$

Here our *register points-to* assertion $r \mapsto v@E$ keeps track of the set E of (thread-local) events that it stems from, along with r ’s current value v , while $a \rightsquigarrow (\phi v)$ records the resources gained (according to protocol ϕ) from the load, *tying* them to its event ID a . (The ϕ, ψ, \dots are per-location value-based protocols used for resource transfer, which we later generalise. Resource transfer using protocols is sound by requiring all threads’ memory operations to follow the protocol: stores must satisfy the protocol resources so that loads can obtain them.)

These register points-to and tied resources then do flow down to later program points (except where transferred away), but, crucially, they can only be used for an event b when a is locally-ordered-before b , e.g. where b is a program-order and data-dependent write after a , which might consume some or all of the resources in passing them to another thread. It is tempting to try to combine these two assertions into one, bypassing the indirection,

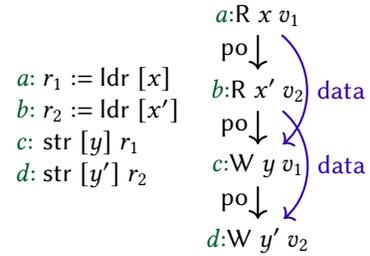


Fig. 6. Intra-thread concurrency

as $r \mapsto v \& (\phi v)$, but this breaks down for all but the simplest use cases: moving the contents of a register into another must distribute the resources, or use indirection as we do via events. Lace logic [Bornat et al. 2015] had a somewhat similar mechanism, but more explicitly in terms of edges than sources, which fits their setting where they dictate ordering (à la Crary and Sullivan [Crary and Sullivan 2015]) better, but is less convenient for ours, where ordering emerges in program order. In general, of course, there may be many dynamic instances – and hence memory events – arising from each static instruction; that can also be dealt with within the logic, by existential quantification and counters for event IDs [Alglave and Cousot 2017; Lamport 1977].

Crucially, we allow any Iris proposition to be tied to an event. This includes any piece of ghost state gs , embedded into an Iris proposition as \boxed{gs} . Ghost state is very flexible [Dinsdale-Young et al. 2013, 2010; Jung et al. 2018, 2015; Svendsen and Birkedal 2014], and, as usual in Iris, we use it both (1) to track the physical state piecemeal, so that we for example can talk about the state of a single register, which is done, as usual in Iris, by enforcing in the definition of weakest precondition that it keeps in sync with physical state introduced in §3.2 and §4.3; and (2) to track the logical state of a (concurrent) program, for example with an exclusive permission to commit to a value, where owning such a permission to commit refutes observing another thread having done something that required having committed to a value (as in §5.6). In a sense, ghost state instruments the physical state of the operational semantics, but unlike physical state, ghost state can be updated freely by a *view shift* $P \Rightarrow Q$, as long as the update is frame-preserving, meaning that it does not contradict other pieces of ghost state; the view shift can be viewed as a generalised implication. Finally, our $a \mapsto P$ assertion is itself defined using ghost state, using the fact that Iris ghost state is higher-order, in the sense that it is mutually defined with Iris propositions.

```

1   $\{r_1 \mapsto \_ * R_0\} // S_1$ 
2   $a: r_1 := \text{ldr } [x]$ 
3   $\{\exists v. r_1 \mapsto v@{a} * (a \mapsto (\phi v)) * R_0\} // S_2$ 
4   $b: \text{str } [y] r_1 \quad (R_0 * \phi v) \Rightarrow (\psi v)$ 
5   $\{r_1 \mapsto v@{a}\} // S_3$ 

```

Fig. 7. Proof sketch for a plain-access (non-release/acquire) version of Fig. 5. The **flow implication** is on line 4.

Using these assertions, we can write concrete proofs for synchronisation involving thread-local dependencies, as sketched in Fig. 7, without relying on the program-order strength of release-acquire reasoning we illustrated in Fig. 5. (This proof sketch is more complicated than needed for LB+datas, which has a very simple proof just asserting all writes write 0, but it generalises to variations of LB, as we show in §5.) The initial logical state of the thread on line 1, S_1 , includes a register points-to for r_1 containing unknown, irrelevant data, which we write with an underscore: $r_1 \mapsto _$, and some potential extra logical state R_0 . The load a reads some value v , so now we have $r_1 \mapsto v@{a}$ and $a \mapsto (\phi v)$. When performing the store b on line 4, the proof rule requires us to establish the corresponding flow implication. Because the store has a data dependency on a , we get to use not only the ambient R_0 , but also the ϕv tied to a , to establish (because this is a store) the protocol for v for y , ψv . The flow implication for b that the proof rule requires us to establish is thus $(R_0 * \phi v) \Rightarrow (\psi v)$.

If the data dependency between a and b is removed (so the store, e.g. now of a constant, can execute early, and hence the relaxed LB behaviour of Fig. 1, where both reads read a non-zero value, is allowed by Arm-A), then the proof does not go through anymore, as desired, because the flow implication for b no longer has ϕv_1 available. This illustrates how our assertions allow us to soundly use ghost state to reason about relaxed architectures exhibiting load buffering and intra-thread concurrency.

Framing. We are separating resources flowing from different sources to different targets by necessity. Relatedly, one of the points of separation logic is allow separate resources to flow side-by-side, for convenience (specifically, for modularity). In the example of Fig. 6, reasoning about c is not allowed to use the resource from b , only from a — thanks to framing, it does not need to mention the resource from b either (similarly, reasoning about d is not allowed to use the resource from a , and does not need to mention them either). We support framing and splitting tied resources, so that given an instruction that merely needs $a \rightsquigarrow P$, for example store b in Fig. 8, we can split $a \rightsquigarrow (P * Q)$ into $(a \rightsquigarrow P) * (a \rightsquigarrow Q)$ (Line 2–3) and frame the latter off (Line 3–4 and 7–8); we explain this mechanism further in §6.4.3.

```

1  a: r := ldr [x]
2  {r ⊢ v@{a} * a ⇝ (P * Q)}
3  {r ⊢ v@{a} * (a⇝P) * (a⇝Q)}
4  {r ⊢ v@{a} * a ⇝ P}
5  b: str [y] r // uses P
6  {r ⊢ v@{a} * a ⇝ ⊤}
7  {r ⊢ v@{a}}
8  {r ⊢ v@{a} * a ⇝ Q}
9  c: str [z] r // uses Q
10 {r ⊢ v@{a} * a ⇝ ⊤}

```

Fig. 8. Splitting and framing tied resources

3.2 The Second Problem: Operationalising the Relaxed Arm-A Model

The next challenge is that of selecting – or developing – a version of the Arm-A concurrency architecture to underlie the soundness proof for our logic. A priori, one might use existing abstract-microarchitectural operational [Pulte et al. 2018], axiomatic [Pulte et al. 2018], or Promising-Arm [Pulte et al. 2019] models, which are proved equivalent (for the features covered by all). We would like to express the logic as an instantiation of Iris [Jung et al. 2018, 2015], an expressive separation logic framework, to get the benefits of its higher-order ghost state, guarded recursion, and existing mechanisation. That requires the underlying semantics to be phrased as a small-step operational semantics, for the logical setup for higher-order ghost state to apply, and it should work ‘enough’ along program order for the soundness proof of our syntax-directed proof rules to be tractable.

The abstract microarchitectural operational model is explanatory, based on hardware intuition, and it is operational, but it is in this respect too close to hardware, with explicit out-of-order execution (speculative execution is generally incongruent for syntax-directed program-logic reasoning); it also splits memory reads and writes into multiple fine-grained events. The axiomatic model as normally presented is not straightforwardly operational: phrased as acyclicity requirements on the ob and certain other derived relations of a whole-program complete candidate execution, expressed using relational algebra with fixpoints over basic relations po , $data$, rf , etc. One might imagine constructing an operational model from the axiomatic model by fiat, with a state that is a set of events, and steps that add an arbitrary event and recheck the axiomatic-model validity predicate, but for Arm-A (and similarly RISC-V and IBM Power), because $po \cup rf$ is not acyclic, this cannot straightforwardly follow program order, as reads would have to sometimes read from events that have not yet been introduced. One might follow ob , but that would be at odds with the structure of the soundness proof. Or one might permit such reads to read new symbolic values, and propagate those through the instruction semantics, but that adds substantial complexity.

Instead, we develop a novel operationalisation of the axiomatic memory model in a mixed operational-axiomatic style (§4), our **second key idea**. This *opax* semantics is sufficiently close to a small-step operational semantics that it is not too difficult to instantiate the Iris logical framework to it, it works enough along program order for the soundness proof of our syntax-directed proof rules to be tractable, and it remains manifestly equivalent to the reference axiomatic model.

Executions in our *opax* semantics are with respect to an ambient complete candidate execution graph that satisfies the axiomatic model validity predicate (but unconstrained by the thread-local ISA semantics), which is picked non-deterministically at the start. The semantics executes threads individually: there is no substantive interleaving or interaction directly between threads, only between single threads and the ambient memory

graph. The instructions of each thread execute in order, keeping only thread-local state – the next memory event identifier, the contents of registers, sources of control dependencies, etc. Each instruction acts as an assertion about the existence of a corresponding memory event at a particular position in the ambient execution graph, and the thread is stuck if an appropriate memory event does not exist in the graph (in which case that specific execution is stuck in the opax semantics, but of course the assembly program itself does not get stuck). This explicitly manipulates a non-thread-local graph; but in the logic, we manage to hide this non-thread-locality in normal cases (as we show in §5).

Candidate graphs in which one or more thread(s) get stuck are simply ignored. This is unusual: getting stuck is not an error; it indicates rather that this particular graph is not consistent with the thread-local semantics of instructions. This was inspired by a related approach taken for the Islaris logic [Sammler et al. 2022] for reasoning about sequential Arm-A machine-code, which faced a similar challenge in that rather different context¹. In a sense, this opax model is merely permuting the order of the usual construction of the axiomatic model: it starts by guessing a valid execution graph (that is, an execution graph that follows the constraints), and then checks that each thread’s contribution in the graph does indeed correspond to an execution of the thread.

One could instead try to work directly over Promising-Arm [Pulte et al. 2019], which is also operational enough for our current purposes in the above senses. In Promising-Arm, apart from promises of future writes (which can all be done at the start of execution, and which also inspire our up-front nondeterministic choice of graph) each thread executes in program order; threads interact through a linear history of writes, keeping track of certain integer *timestamps* (indices into the history of writes), which constrain how instructions can interact with the history. Timestamps keep track of lower bounds on the sources of register values (and some whole-thread bounds for barriers), abstracting the set of source events for each. These integer timestamps might be technically easier to work with than graphs, but we found the explicit nodes with explicit edges of the axiomatic model helpful in developing our model of assertions. In a sense, our opax semantics is a reformulation of an axiomatic model made to look more like a promising model, but with the advantage that changes to the axiomatic model apply directly.

Note that, while our opax semantics technically qualifies as an operational semantics, it does not meet most usual expectations of such. In particular, there is no reasonable sense in which it is executable.

By putting an axiomatic model in the required shape, we need a non-standard definition of weakest precondition (§6.4) and a non-standard proof of adequacy (§7) (even more so as our threads are executed independently), but we still benefit from more fundamental Iris features like higher-order ghost state, which one would not want to reconstruct.

Developing a separation logic directly over an axiomatic memory model has previously been done either using a non-standard semantics of assertions (e.g. RSL, FSL, and GPS, which, as noted by Kaiser et al. [Kaiser et al. 2017, §1.2], requires significant effort), or by defining an equivalent, operational model (e.g. iGPS [Kaiser et al. 2017] and ORC11 [Dang et al. 2020]), which is challenging when the model allows very relaxed behaviour.

3.3 The Third Problem: Structuring the Adequacy Proof

Finally, given a proof in AxSL^{Arm} using our new assertions, we then need an adequacy theorem (§7), which, given a family of thread-local proofs in our logic, gives a statement about the soundness of a whole program in the meta-logic w.r.t. the Arm-A semantics. To prove such an adequacy theorem, we need to address a tension between our program-logic proof, which is syntax-directed, and therefore in program order, and synchronisation, which is along rf – even though there can be cycles in $\text{po} \cup \text{rf}$, which prevents doing an induction on it. **Our third key idea** is that, to solve this tension, we can split the proof of adequacy in two phases: first along po , and

¹To reason above the full Arm-A ISA semantics without being overwhelmed with irrelevant detail, Islaris simplified the semantics of each instruction with respect to chosen assumptions, e.g. about Arm-A system register values and alignment facts, using Isla SMT-assisted symbolic execution [Armstrong et al. 2021]. The resulting symbolic traces contain asserts on some paths, which (when they fail to hold) discard those paths from the instruction semantics – which the Islaris instantiation of Iris exploits.

then along ob . The first phase, along po , uses thread-local resources to establish, for each thread, and for each memory event of that thread, that the flow implication for that event holds. The second phase, along ob , stitches the flow implications together. For example, this second phase walks through the thread of Fig. 6 twice, for the two disjoint components of ob : once from a to c , and once from b to d (with both orderings being possible).

4 The Languages

As the focus of this paper is on real-world concurrency rather than realistic instruction set architectures, we consider a simplified assembly language, TinyArm, in which to write simple Arm-A concurrency-model programs. However, akin to itrees [Xia et al. 2020], we give its semantics by elaborating it into a new mechanisation of the outcome interface type of Sail [Gray et al. 2015][Pulte et al. 2018, §6.1][Pulte 2018, §2.3] (§4.1) by Pérami et al. [2026] (which is not a contribution of this paper), translating instructions into sequences of their semantic “microinstructions”: primitive register and memory accesses, and Arm-A fences.

These are what our logic actually reasons about. Using our basic rules for the interface events, we then give high-level rules for our toy instructions. We expect the logic to extend naturally to the full Sail semantics for a large fragment of the Arm-A instruction-set architecture (ISA), using either the Sail-generated Rocq definitions for the ISA, or (as in Islaris [Sammler et al. 2022]) the output of the Isla symbolic evaluator for Sail [Armstrong et al. 2021], both of which express the intricate real semantics of instructions in terms of that same outcome interface type.

Our simplified language is shown in Fig. 9. Loads and stores are parameterised by an *ordering strength*, os , either plain, release/acquire, or weak-acquire, and a *variety*, vr : non-exclusive or exclusive. The output register of a store is used only for the success/fail value of a store exclusive; a dummy register is used for other stores.

Besides TinyArm, Fig. 9 also shows the syntax of TinySc, an even simpler language in which we write concurrent programs for a sequentially consistent (SC) memory model. We use this compact language to demonstrate the core idea of opax which is the formal foundation that two logics in §5 build upon. TinySc is syntactically a sublanguage of TinyArm, where we elide the os and vr and omit the barriers. Therefore, in the rest of this section, we present their semantics by detailing one and then merely explaining how the other relates to it.

4.1 The Elaboration Semantics of Instructions into the Sail Outcome Interface

The Sail outcome interface defines the intra-instruction semantics for each instruction, independently from the behaviour of registers and memory. It does this in terms of abstract microinstructions, formally a free monad of effects on *outcomes*, with constructors `RegRead`, `RegWrite`, `MemRead`, `MemWrite`, etc. Each of these takes the appropriate arguments, and the free monad constructor `Next` pairs it with a continuation which takes any register or memory read result and gives the subsequent intra-instruction semantics.

We show how to elaborate TinyArm instruction using the interface, especially how to handle Arm’s architectural dependencies, and then how to reuse the elaboration for TinySc instructions by merely ignoring those Arm specifics.

4.1.1 The Elaboration of TinyArm Instructions. Given an ordering strength os , a variety (exclusive or non-exclusive) vr , and address x , and dependencies $d \in \text{Dep} \triangleq P(\text{Reg}) \times P(N)$ (composed of register dependencies r , and intra-instruction event dependencies m), `MemRead os vr x d` has type (roughly) `Outcome Word`, wrapped in the instruction monad `IMon A` which has constructors `NextT : Outcome T → (T → IMon A) → IMon A` and `Ret : A → IMon A`. Rather than give an exhaustive definition, we sketch how a few special cases of our instructions elaborate into (a meta-level Rocq program over) this Sail outcome interface in Fig. 10.

A plain, non-exclusive load `r := ldr [taddr]`, into register r from address t_{addr} , elaborates into a plain (and non-exclusive) memory read from that address with address dependencies given by the auxiliary function $\mathbb{D}[-]$ (in our simplified language, dependencies can be computed from the syntax), the value of which is bound to v ,

$i_{\text{TinyArm}} ::=$	instructions	
nop		$r \in \text{Reg} \triangleq \{r_0, r_1, \dots\}$
$r := t$	register assignment	$v, x \in \text{Word} \triangleq 0..2^{64} - 1$
br x	branch to address x	$op ::= + \mid - \mid \times$
bne $t \ x$	conditional branch	$t ::= v \mid r \mid t_1 \ op \ t_2$
$r := \text{ldr}_{os, vr} [t_{\text{addr}}]$	memory load	$os ::= \text{plain} \mid \text{relacq} \mid \text{weakacq}$
$r := \text{str}_{os, vr} [t_{\text{addr}}] \ t_{\text{data}}$	memory store	$vr ::= \text{nexcl} \mid \text{excl}$
dmb sy dmb st dmb ld isb	Arm-A fences	
$r := \text{ldr} [t_{\text{addr}}]$	\triangleq	$r := \text{ldr}_{\text{plain}, \text{nexcl}} [t_{\text{addr}}]$
$r := \text{ldar} [t_{\text{addr}}]$	\triangleq	$r := \text{ldr}_{\text{relacq}, \text{nexcl}} [t_{\text{addr}}]$
$r := \text{ldr}_x [t_{\text{addr}}]$	\triangleq	$r := \text{ldr}_{\text{plain}, \text{excl}} [t_{\text{addr}}]$
$\text{stlr} [t_{\text{addr}}] \ t_{\text{data}}$	\triangleq	$r := \text{str}_{\text{relacq}, \text{nexcl}} [t_{\text{addr}}] \ t_{\text{data}}$
$r := \text{str}_x [t_{\text{addr}}] \ t_{\text{data}}$	\triangleq	$r := \text{str}_{\text{plain}, \text{excl}} [t_{\text{addr}}] \ t_{\text{data}}$

$$i_{\text{TinySc}} ::= \text{nop} \mid r := t \mid \text{br } x \mid \text{bne } t \ x \mid r := \text{ldr} [t_{\text{addr}}] \mid r := \text{str} [t_{\text{addr}}] \ t_{\text{data}}$$

Fig. 9. Instructions and syntactic sugar of TinyArm and TinySc

followed by a register write to r of v with dependencies $\langle \emptyset, \{0\} \rangle$ meaning that the register write has no register dependency and a dependency on the 0th MemRead of the instruction. See the corresponding reduction rule [A-REG-WRITE](#) of our semantics in §4.3 for how the real dependencies are computed from this bookkeeping type Dep. Finally, the elaboration is followed by a program counter increment (which we write with a bind $\gg=$ to be systematic).

A non-exclusive store release $\text{stlr} [t_{\text{addr}}] \ t_{\text{data}}$, elaborates into the sequence of the elaboration of t_{addr} to obtain the address x , and of the elaboration of t_{data} to obtain the value v (both defined using the auxiliary function $\mathbb{T}[-]$), followed by a release (and non-exclusive) memory write to x of v , with its address and data dependencies (as given by the auxiliary function $\mathbb{D}[-]$), and again a PC increment (the dummy register is not mentioned).

We use the Sail interface BranchAnnounce outcome to capture dependencies of branches, which we elaborate into evaluation of their condition, followed by, depending on whether the condition holds (using the conditional of the meta-language), either a write of the given address x to the program counter, or a normal program counter increment.

4.1.2 The Elaboration of TinySc Instructions. We reuse the elaborations of plain, non-exclusive load and store of TinyArm as the elaborations of load and store of TinySc respectively, and identical elaborations for other shared instructions. The elaboration is further simplified by ignoring register dependencies (defining $\mathbb{D}[-]$ to $\lambda_. \emptyset$), since register dependencies do not matter for SC. For the sake of simplicity, we omit the constant arguments of microinstructions, and for example only write MemRead x for MemRead plain nexcl $x \ \emptyset$, when we present the microinstructions for TinySc.

4.2 The Conventional Axiomatic Concurrency Model Semantics

A program working over the Sail outcome interface can then be connected to different types of memory models: operational, axiomatic, or promising. For an axiomatic memory model, this is usually done by recursively computing the set of thread-local instruction-semantics pre-executions of (the control-flow unfoldings of) each

$$\begin{aligned}
 \llbracket r := \text{ldr } [t_{\text{addr}}] \rrbracket &\triangleq \mathbb{T}\llbracket t_{\text{addr}} \rrbracket \gg= \lambda x. \text{Next} (\text{MemRead plain nexcl } x \ (\mathbb{D}\llbracket t_{\text{addr}} \rrbracket)) \\
 &\quad (\lambda v. \text{Next} (\text{RegWrite } r \ v \ \langle \emptyset, \{0\} \rangle) (\lambda(). \text{Ret } ())) \gg \text{IncPC} \\
 \llbracket \text{stlr } [t_{\text{addr}}] \ t_{\text{data}} \rrbracket &\triangleq \mathbb{T}\llbracket t_{\text{addr}} \rrbracket \gg= \lambda x. \mathbb{T}\llbracket t_{\text{data}} \rrbracket \gg= \lambda v. \\
 &\quad \text{Next} (\text{MemWrite rel nexcl } x \ v \ (\mathbb{D}\llbracket t_{\text{addr}} \rrbracket) \ (\mathbb{D}\llbracket t_{\text{data}} \rrbracket)) (\lambda(). \text{Ret } ()) \gg \text{IncPC} \\
 \llbracket \text{bne } t \ x \rrbracket &\triangleq \mathbb{T}\llbracket t \rrbracket \gg= \lambda v. \text{Next} (\text{BranchAnnounce } x \ \mathbb{D}\llbracket t \rrbracket) \\
 &\quad \left(\lambda(). \begin{array}{l} \text{if } v = 0 \text{ then IncPC} \\ \text{else Next} (\text{RegWrite pc } x \ \langle \emptyset, \emptyset \rangle) (\lambda(). \text{Ret } ()) \end{array} \right) \\
 \text{IncPC} &\triangleq \text{Next} (\text{RegRead pc}) (\lambda v. \text{Next} (\text{RegWrite pc } (v + 4) \ \langle \emptyset, \emptyset \rangle) (\lambda(). ())) \\
 \\
 \mathbb{T}\llbracket v \rrbracket &\triangleq \text{Ret } v & \mathbb{D}\llbracket v \rrbracket &\triangleq \emptyset \\
 \mathbb{T}\llbracket r \rrbracket &\triangleq \text{Next} (\text{RegRead } r) \text{Ret} & \mathbb{D}\llbracket r \rrbracket &\triangleq \{r\} \\
 \mathbb{T}\llbracket t_1 \text{ op } t_2 \rrbracket &\triangleq \mathbb{T}\llbracket t_1 \rrbracket \gg= \lambda v_1. \mathbb{T}\llbracket t_2 \rrbracket \gg= \lambda v_2. \text{Ret } (v_1 \circ [op] v_2) & \mathbb{D}\llbracket t_1 \text{ op } t_2 \rrbracket &\triangleq \mathbb{D}\llbracket t_1 \rrbracket \cup \mathbb{D}\llbracket t_2 \rrbracket
 \end{aligned}$$

Fig. 10. A few cases of the elaboration of TinyArm into the outcome interface (eliding some details), where `ldr` is the syntactic sugar of `ldrplain,nexcl` and `stlr` is the syntactic sugar of `stlrrel,nexcl`. The computation of intra-instruction dependencies is **highlighted**.

thread, allowing arbitrary concrete values for register and memory reads, and then taking the Cartesian product of these sets, which ensures that all the pre-executions are consistent with (the instruction semantics of) the program. For each such pre-execution, one enumerates the set of candidate executions, decorating the pre-execution with `rf` and `co` relations (unconstrained except for some well-formedness properties). Finally, one filters those with the axiomatic-model validity predicate.

4.3 Our Opax Concurrency Model Semantics

As discussed in §3.2, it is not clear how to define a syntax-directed program logic over this axiomatic style of semantics. Hence, we reformulate the combination of axiomatic model and instruction semantics in our novel *opax* semantics, mixing *operational* and *axiomatic* styles. We first present the language-agnostic shape of this new style of semantics, and then give concrete *opax* semantics definitions for TinySc and TinyArm. Similar to how the instruction elaboration of TinyArm extends the instruction elaboration of TinySc with bookkeeping, the instantiation of TinyArm as an *opax* semantics extends the instantiation of TinySc with bookkeeping local states for register dependencies, etc.

4.3.1 Opax Candidate Executions. Unlike the conventional definition of candidate execution, an *opax candidate execution* has a different meaning. It consists of the usual events and relations, but we swap the validity check with the program consistency check. That is, an *opax candidate execution* has to be well-formed and satisfies the axiomatic-model validity predicate, but is not assumed to be consistent with a program. We define it using the outcome interface. Formally, an *opax candidate execution* (graph) X comprises a collection of events in the form of a function `lab` from event IDs (of type *Eid*) to events, and various relations between events of type $Eid \times Eid$, where an event is of type $\text{Outcome } T \times T$, that is, a pair of an outcome request and its response. The validity predicate usually consists of acyclicity requirements on compound relations (defined using base relations), and differs from model to model. The well-formedness condition for relations is standard, except for `po`, for which the constraint depends on the implementation of *Eid* (we elaborate on this soon). The semantics of a program is then defined as the set of *opax candidate executions* consistent with the program. We give two formal instantiations following this new notion below.

Opax Candidate executions of SC. An opax candidate execution for SC is defined as

$$X \triangleq \langle \text{lab, po, rf, co, fr, sc} \rangle$$

with reads of arbitrary values and writes of values, and with arbitrary reads-from (rf), coherence (co), and fr (equivalent to rf^{-1} ; co) relations between them. The validity requirement is the acyclicity of $\text{sc} \triangleq \text{po} \cup \text{rf} \cup \text{co} \cup \text{fr}$, where $e \text{ sc } e'$ means event e happens before e' .

Opax Candidate executions of Arm-A. An opax candidate execution for Arm-A is defined as

$$X \triangleq \langle \text{lab, po, rf, co, fr, ctrl, addr, data, rmw, \dots (compound relations)} \rangle$$

with new ctrl, data, and addr dependencies, and rmw base relation for exclusives. The compound relations and validity requirements are defined in Fig. 4.

4.3.2 The Shape of the Opax Semantics. The opax semantics works in two phases, to filter out execution graphs that are both consistent with respect to the concurrency model, and conformed with the program semantics. In the first phase, execution of a whole program starts by guessing a complete opax candidate execution graph X for the program. This execution graph is well-formed and consistent, but is otherwise unconstrained, and in particular is for now unrelated to the program itself. The guessing action does not have an operational step, which is simply reflected by a universally quantification.

In the second phase, we validate if the guessed candidate execution graph from the previous phase conform with the program by executing each thread independently to completion operationally:

$$\begin{array}{c} \text{VALIDATION} \\ \frac{\textcircled{1} s_{\text{init}} c_1 \xrightarrow{1, X, I}_h^* \textcircled{2} \text{Done } \langle _ \rangle \quad \dots \quad (s_{\text{init}} c_n) \xrightarrow{n, X, I}_h^* \text{Done } \langle _ \rangle}{\textcircled{3} \langle (c_1 \parallel \dots \parallel c_n), I, X \rangle \rightarrow \checkmark} \end{array}$$

We start this phase with a global configuration $\textcircled{3}$, where c_1, \dots, c_n are the initial program counter values of the different threads, and X is the guessed graph. For simplicity, we assume a fixed instruction memory I , a map from addresses to instructions (elaborated into micro-instructions). Instruction fetching could be accurately modelled in the memory model, as per Simner et al. [Simner et al. 2020], but this would lead to significant complexity, and only be relevant for programs that modify their code.

The global configuration passes the validation if it passes the local check of each thread, which is implemented by our thread-local opax semantics. Thread transitions $s \xrightarrow{tid, X, I}_h s'$ are indexed by the thread ID, execution graph, and instruction memory, and are deterministic. Each thread state s is either Ctd C (“continued”), which represents an ongoing thread execution, or Done T , which represents a completed thread execution. Here C is a tuple $\langle p, T \rangle$ where p is the remaining microinstruction program in the Sail outcome interface for the current instruction, and T is the thread state which is dependent on the language and the concurrency model.

For each thread tid , we start the local validation with its initial thread state $\textcircled{1} (s_{\text{init}} c_{tid})$, with pc set to c_{tid} and microinstruction program $\text{Ret } ()$ (before the first instruction has started). The execution terminates with the terminating state $\textcircled{2}$ when it has finished execution of the current microinstruction program and the program counter points outside of instruction memory. The completion of the execution means that the fragment of the guessed graph X for the thread conforms well with the thread’s program. We perform this process for all threads to obtain the global conformity.

Remarks. It is worth-noting that the only source of non-determinism is the guessing step; all following thread-local reductions are deterministic. A similar pattern appears in operational semantics with a quantified scheduler, where picking the scheduler is non-deterministic, and the interleaving is determined by the scheduler. It is also worth-pointing-out that a stuck thread is not an error state: it rather indicates that the guessed graph does not

correspond to this program. Rule **VALIDATION** of our semantics ignores these wrongly guessed graphs, leaving only the execution graphs of the program.

In the rest of this subsection, we give the instantiations of the opax semantics for the two languages. We start with the semantics of TinySc, which is compact and makes it easy to demonstrate the core idea, and then show the semantics of TinyArm, explaining how to handle the complexities that come with Arm-A: access kinds and dependencies.

4.3.3 Opax Semantics for TinySc. For TinySc, we instantiate thread state T with a tuple $\langle regs, IT \rangle$ to track a register state $regs$ - a finite map from register names to values - and an intra-instruction state IT which only comprises a thread local event counter $cntr$ (We will extend it with more components later for TinyArm).

We sketch selected rules of the thread operational semantics in Fig. 11.

A thread executes by executing the current microinstruction program until it ends, at which point all microinstructions of the current instruction must have been executed. The rule **S-RELOAD** checks that with ⑥, and then fetches the microinstruction program of the next instruction at the address in register pc by looking it up in I (⑤). The rule **S-TERM** terminates the execution, when pc is outside the instruction memory (⑧), at which point there must not be further events by this thread in the graph (checked by ⑨).

Each other rule in Fig. 11 validates one event of the guess graph with respect to the current microinstruction ins of Next ins K . The event identifier of the current microinstruction $e = \langle tid, IT.cntr \rangle$ comprises a thread identifier (zero being reserved for the ‘initial’ thread that contains all the initial writes), and the event counter $IT.cntr$ which is an ordered pair comprising an instruction counter $icntr$ and an intra-instruction event counter $ecntr$.

A MemRead microinstruction can execute (rule **S-MEM-READ**) only when there is a corresponding memory read event in the execution graph (①); otherwise, this instruction (and thus this thread) is stuck. This event has to have the appropriate po edges to it. To check po edges, the intra-instruction counter $ecntr$ of $IT.cntr$ gets incremented with next-e after executing every microinstruction (③); the instruction counter $icntr$ gets incremented with next-i (⑦), which additionally resets $ecntr$ when finishing up an instruction (in **S-RELOAD**). po edges are special, in the sense that a po edge between two non-initial events can be checked for by determining whether their identifiers have same thread id and their local event counter values are lexicographically ordered:

$$\frac{e_1 = \langle tid, \langle icntr, ecntr \rangle \rangle \quad e_2 = \langle tid, \langle icntr', ecntr' \rangle \rangle \quad (icntr < icntr') \vee (icntr = icntr' \wedge ecntr < ecntr')}{e_1 X.po e_2}$$

This is part of the well-formedness condition of the execution graph. This indicates that we do not need to explicitly check if there is a po edge between two events in the graph – we only need to increment the counters correctly.

A RegWrite microinstruction can similarly execute (rule **S-REG-WRITE**) only when there is a corresponding register write event in the execution graph. The graph register write event needs to agree with the thread-local register state $regs$ on the value of the write (④). (This register event is not used in the axiomatic memory models for either SC or Arm-A, because the former does not need it and the latter instead uses primitive dependency relations, but the interface includes it to support operational models and other axiomatic models.)

4.3.4 Opax Semantics for TinyArm. We define the opax semantics for TinyArm by extending the TinySc semantics with instrumentation to track Arm-A dependencies and access kinds. Concretely, we first extend the thread state T with the set $srcs_{ctrl}$ of sources of control dependencies. Then, we augment every register in $regs$ with dependency information $srcs_d$ alongside its value. Finally, we add the list of identifiers of intra-instruction read events seen so far mrd as a new field of IT . The thread state T is now defined as $\langle regs; srcs_{ctrl}; IT \rangle$ where $regs \in RegName \rightarrow \{v : Val; srcs_d : \mathcal{P}(Eid)\}$, $srcs_{ctrl} \in \mathcal{P}(Eid)$, and $IT \in \{cntr : \{icntr : Nat; ecntr : Nat\}; mrd : list(Eid)\}$. In §6.4.4, we will further extend T with the previous exclusive read event $srcs_{rmw}$ to support exclusives.

$$\begin{array}{c}
\text{S-MEM-READ} \\
\frac{\textcircled{1} X.\text{lab}(e) = R \ x \ v \quad \textcircled{2} e = \langle tid, IT.cntnr \rangle}{\text{Ctd} \langle \text{Next} (\text{MemRead } x) \ K, \langle \text{regs}, IT \rangle \rangle \xrightarrow{tid, X, I}_h \text{Ctd} \langle K \ v, \langle \text{regs}, \textcircled{3} \text{next-e}(IT) \rangle \rangle} \\
\text{S-MEM-WRITE} \\
\frac{X.\text{lab}(e) = W \ x \ v \quad e = \langle tid, IT.cntnr \rangle}{\text{Ctd} \langle \text{Next} (\text{MemWrite } x \ v) \ K, \langle \text{regs}, IT \rangle \rangle \xrightarrow{tid, X, I}_h \text{Ctd} \langle K \ (), \langle \text{regs}, \text{next-e}(IT) \rangle \rangle} \\
\text{S-REG-WRITE} \\
\frac{X.\text{lab}(e) = \text{RegW } r \ v \quad e = \langle tid, IT.cntnr \rangle}{\text{Ctd} \langle \text{Next} (\text{RegWrite } r \ v) \ K, \langle \text{regs}, IT \rangle \rangle \xrightarrow{tid, X, I}_h \text{Ctd} \langle K \ (), \langle \textcircled{4} \text{regs}[r \mapsto v], \text{next-e}(IT) \rangle \rangle} \\
\text{S-REG-READ} \\
\frac{X.\text{lab}(e) = \text{RegR } r \ v \quad e = \langle tid, IT.cntnr \rangle \quad \text{regs}(r) = v}{\text{Ctd} \langle \text{Next} (\text{RegRead } r) \ K, \langle \text{regs}, IT \rangle \rangle \xrightarrow{tid, X, I}_h \text{Ctd} \langle K \ (), \langle \text{regs}, \text{next-e}(IT) \rangle \rangle} \\
\text{S-RELOAD} \\
\frac{\text{regs}(\text{pc}) = x \quad \textcircled{5} I(x) = p \quad \textcircled{6} \text{instr-done}(X, tid, IT.cntnr)}{\text{Ctd} \langle \text{Ret} \ (), \langle \text{regs}, ?R, IT \rangle \rangle \xrightarrow{tid, X, I}_h \text{Ctd} \langle p, \langle \text{regs}, \textcircled{7} \text{next-i}(IT) \rangle \rangle} \\
\text{S-TERM} \\
\frac{\text{regs}(\text{pc}) = x \quad \textcircled{8} x \notin \text{dom}(I) \quad \textcircled{9} \text{prog-done}(X, tid, IT.cntnr)}{\text{Ctd} \langle \text{Ret} \ (), \langle \text{regs}, IT \rangle \rangle \xrightarrow{tid, X, I}_h \text{Done} \langle \text{regs}, IT \rangle} \\
\text{next-e}(\langle icntnr, ecntnr \rangle) \triangleq \langle icntnr, ecntnr + 1 \rangle \\
\text{next-i}(\langle icntnr, ecntnr \rangle) \triangleq \langle icntnr + 1, 0 \rangle \\
\text{instr-done}(X, tid, \langle icntnr, ecntnr \rangle) \triangleq \forall e \in \text{dom}(X.\text{lab}). \text{same-instr}(e, tid, cntnr.icntnr) \rightarrow e.cntnr.ecntnr < ecntnr \\
\text{prog-done}(X, tid, \langle icntnr, ecntnr \rangle) \triangleq \forall e \in \text{dom}(X.\text{lab}). \text{same-thd}(e, tid) \rightarrow \\
(e.cntnr.icntnr < icntnr) \vee (e.cntnr.icntnr = icntnr \wedge e.cntnr.ecntnr \leq ecntnr)
\end{array}$$

Fig. 11. Selected reduction rules of opax semantics for TinySc with auxiliary functions.

The two selected rules in Fig. 12 illustrate how these extensions check dependency edges. A non-exclusive MemRead microinstruction now has to have the appropriate addr and ctrl edges to it (rule **A-MEM-READ-NEXCL**), as checked using the thread state (the set of data dependencies d_{addr} and the control dependency sources $\text{srcs}_{\text{ctrl}}$ respectively). Unlike how we treat po, here we have to check if those dependency edges exist in the graph explicitly ($\textcircled{1}$ and $\textcircled{2}$). The event ID e is appended to the intra-instruction memory read list $IT.mrd$ by auxiliary function intra-read-app ($\textcircled{3}$), so that later register microinstructions can obtain e by providing a position in the list to check a dependency from e if the register value is computed from the read value v (we demonstrate this with the elaboration of load in Fig. 10 in the following paragraph).

A RegWrite microinstruction (rule **A-REG-WRITE**) now also updates the dependency of the register d_{reg} for the local registers in regs . The dependency d_{reg} is a set of event identifiers computed from two sources tracked with

d : the union of the dependencies of every register in $d.r$ (④), and intra-instruction event dependencies that are memory reads whose indices are in $d.m$ (⑤). For instance, in the elaboration of load in Fig. 10 where d is instantiated to $\langle \emptyset, \{0\} \rangle$, d_{reg} is computed to be $\{e\}$ when $IT.mrd$ is $[e]$ (that is, we take the 0th event from the list), where e is the event ID of the memory read event preceding the register write. Therefore, we conclude that the data of the register v comes from e , and update $regs$ accordingly.

$$\begin{array}{c}
 \text{A-MEM-READ-NEXCL} \\
 \hline
 X.\text{lab}(e) = R_{os, \text{nexcl}} x v \quad e = \langle tid, IT.\text{cntr} \rangle \quad \textcircled{1} \{ \langle e_d, e \rangle \mid e_d \in \text{srcs}_{ctrl} \} = to(e, X.\text{ctrl}) \\
 \textcircled{2} \{ \langle e_d, e \rangle \mid r \in d_{addr}.r \wedge \text{regs}(r) = \langle _ , \text{srcs}_d \rangle \wedge e_d \in \text{srcs}_d \} = to(e, X.\text{addr}) \quad \textcircled{3} \text{intra-read-app}(IT, e) = IT' \\
 \hline
 \text{Ctd} \langle \text{Next} (\text{MemRead } os \text{ nexcl } x d_{addr}) K, \langle \text{regs}, \text{srcs}_{ctrl}, IT \rangle \rangle \xrightarrow{tid, X, I}_h \text{Ctd} \langle K v, \langle \text{regs}, \text{srcs}_{ctrl}, \text{next-e}(IT') \rangle \rangle \\
 \\
 \text{A-REG-WRITE} \\
 \hline
 X.\text{lab}(e) = \text{RegW } r v \quad e = \langle tid, IT.\text{cntr} \rangle \quad d_{reg} = \textcircled{4} \left(\bigcup_{\{ \text{srcs}_d \mid r \in d.r \wedge \text{regs}(r) = \langle _ , \text{srcs}_d \rangle \}} \text{srcs}_d \right) \cup \textcircled{5} \left(\bigcup_{i \in d.m} \{ IT.mrd[i] \} \right) \\
 \hline
 \text{Ctd} \langle \text{Next} (\text{RegWrite } r v d) K, \langle \text{regs}, \text{srcs}_{ctrl}, IT \rangle \rangle \xrightarrow{tid, X, I}_h \text{Ctd} \langle K (), \langle \text{regs}[r \mapsto \langle v, d_{reg} \rangle], \text{srcs}_{ctrl}, \text{next-e}(IT) \rangle \rangle
 \end{array}$$

Fig. 12. Selected reduction rules of our operationalised semantics. We write $to(e, \mathcal{R})$ for the set of edges of type \mathcal{R} with target e . The instrumentation to deal with Arm dependencies is **highlighted**.

4.3.5 Stuckness. The guessing step and the fact that executions can get stuck mean that this model is not executable as such, but this is not problematic for a logic. First, we discard stuck executions by assuming unstuckness in the definition of weakest preconditions. Second, the guessing does not appear in the definition of weakest preconditions, which takes the guessed graph as a parameter; instead, the guessing is handled by a quantification in the adequacy theorem, when the proofs of the individual threads are combined.

4.3.6 Infinite Executions. Handling infinite executions in memory models exhibiting load buffering is currently an open problem. The problem manifests in axiomatic models in the form of an infinite regress, where an event is justified by a program-order-later event, itself justified by another program-order-later event, ad infinitum, without an eventual grounding, but because this is not a cycle, most axiomatic models do not reject this kind of execution. The same underlying problem appears in the promising and operational models of Arm-A under a different guise. We do not attempt to tackle this problem, and our opax semantics sticks to the axiomatic model as-is.

5 The Logics

The goal of this section is to build up to AxSL^{Arm} . We do this incrementally, introducing two intermediate logics to explain the different building blocks of AxSL^{Arm} .

(1) We start by tackling only the challenge of defining a logic on top of an opax semantics, and illustrate it with our first logic: AxSL^{SC} . We start from a simple setting: sequentially consistent (SC) concurrency: the question here is how to deal, in a logic like Iris, with a fixed execution graph representing the shared memory. This is merely the first step: AxSL^{SC} is built on the right foundations (namely, an opax semantics) to scale to relaxed concurrency, but it still bakes in too much ordering in its structure.

(2) We then move on to describe our novel style of assertions compatible with relaxed memory, and illustrate it with our second logic: $\text{AxSL}^{\text{SCExt}}$, a logic with both foundations and an assertion style compatible with relaxed concurrency. For simplicity, $\text{AxSL}^{\text{SCExt}}$ stays in the context of sequential consistency, but follows the

fine-grained resource management style sketched in §2. In particular, $\text{AxSL}^{\text{SCExt}}$ employs ‘tied-to’ assertions and flow implications in order to be compatible with relaxed concurrency.

(3) Finally, we present AxSL^{Arm} , our logic for the relaxed memory of Arm. The final challenge is to deal with the subtleties of such a memory model: syntactic dependencies, external vs. internal reads, exclusives, etc.

For each of the three logics, we present a selection of its proof rules followed by examples. The proof rules (and the Hoare triples used in them) of the three logics are defined at two abstraction levels: the underlying rules are for microinstructions, and are proven sound against the semantics of Hoare triples described in §6.4, while the high-level rules explicitly used in proofs of programs are for the surface instructions of Fig. 9, derived from the former by reasoning about instruction elaboration (§4.1). Before diving into the three logics, we first discuss the resource transfer mechanism that they employ.

5.1 Resource Transfer with Protocols

Concurrent separation logic (CSL) uses invariants to share and transfer resources. However, invariants are unsound in the context of relaxed concurrency [Dang et al. 2020; Turon et al. 2014]. Intuitively, this is because, with relaxed concurrency, threads may have different views on the shared memory, thus owning potentially inconsistent resources describing those views, while classic CSL invariants require the transferred resources to be consistent across all threads.

Furthermore, even though invariants have been shown to work well with heap reasoning, it is unknown how they work in conjunction with graph reasoning. Recall that in CSL, to share memory resources (e.g. some points-tos), one usually allocates some invariant with them, and then distributes the duplicates of the invariant to the threads. The threads then may obtain the ownership of the resources (temporarily) by opening the invariant for loading and storing. In the graph-based approach that we present in this section, one is not required to own any shared memory assertions to access the shared memory; instead, one makes *assertions* about the graph representing memory. It is not clear to us how this can be made to work with invariants, in particular how to make connections between the newly-gained graph assertions and the resources shared by invariants.

Instead, in AxSL, we use a notion of protocol for resource transfer, inspired by previous relaxed memory logics including RSL and GPS, whose relation to AxSL we discuss in more detail in §9. Our protocols Φ are a variant of rely/guarantee-style protocols, and enable thread-local reasoning by expressing the intended resource transfers between threads across an entire program. More specifically, for each location x , $\Phi(x)$ specifies what resources are transferred between threads by accesses to location x : writes to x are required to send resources mandated by $\Phi(x)$, and in exchange, reads from x are allowed to assume that they receive the resources mandated by $\Phi(x)$, as usual. These protocols are compatible with the invariants we use for exclusives in §5.7. (For simplicity, we only consider static protocols with a single state in the formalisation.) Our Hoare triples are then parametrised by this fixed protocol, and ensure that write operations respect the protocol by requiring that protocol resources are provided, which allows reads to rely on these guarantees by assuming the same resources. Triples can only be composed between threads when they use the same protocol, enforcing whole-program agreement.

Technically, a protocol Φ of type $Loc \rightarrow Val \rightarrow Eid \rightarrow iProp$ takes as arguments the location x , a value v , and an event ID e . The event ID e is associated with the write event that fulfils the protocol. This event ID argument e is used for explicit reasoning about the execution graph, as illustrated in the message passing example below. For example, it makes it possible to state “there exists a write to a certain address of a certain value that is lob-before e ”. For simpler cases, this last argument can be elided, as we have so far.

5.2 Dealing with Opax Semantics: The AxSL^{SC} Logic

AxSL^{SC} is a thin logical layer above our opax semantics of TinySc. It essentially exposes all the details of the opax semantics to the users of the logic, which gives the logic an unique shape and new reasoning principles

compared to classic operational-based CSLs. The distinction primarily comes from the diverging representations of the shared memory of the underlying semantics. Since the opax semantics guesses a fixed execution graph upfront, in AxSL^{SC}, we can only deal with persistent knowledge on graph events and memory orders, rather than the usual points-to assertions that represent the state of a shared and dynamic heap at individual locations (We show how to close this gap by recovering points-tos in AxSL^{SC} in §B). To understand how AxSL^{SC} works, in particular its graph-based compositional reasoning, we show some of its proof rules, and then verify a message passing example.

$$\begin{array}{l}
 \text{SC-HT-MICRO-MEMREAD} \\
 \left\{ \textcircled{1} \text{PoPred}(e_{\text{po}}) * \forall e, v, e_w. \left(\textcircled{2} \text{GraphFactsR}(e, x, v, e_w, e_{\text{po}}) * \textcircled{3} \Phi(x, v, e_w) \right) \right\} \\
 \quad \Rightarrow \textcircled{4} P(e, x, v, e_w, e_{\text{po}}) * \Phi(x, v, e_w) \\
 \text{MemRead } x \\
 \left\{ v. \exists e, e_w. \textcircled{5} \text{PoPred}(e) * \text{GraphFactsR}(e, x, v, e_w, e_{\text{po}}) * P(e, x, v, e_w, e_{\text{po}}) \right\}_{tid, \Phi} \\
 \\
 \text{SC-HT-MICRO-MEMWRITE} \\
 \left\{ \textcircled{1} \text{PoPred}(e_{\text{po}}) * \forall e. \textcircled{2} \text{GraphFactsW}(e, x, v, e_{\text{po}}) \Rightarrow \textcircled{3} \Phi(x, v, e) \right\} \\
 \text{MemWrite } x \ v \\
 \left\{ (). \exists e. \textcircled{4} \text{PoPred}(e) * \text{GraphFactsW}(e, x, v, e_{\text{po}}) \right\}_{tid, \Phi} \\
 \\
 \text{SC-HT-MICRO-REGWRITE} \\
 \{ r \mapsto _ \} \text{RegWrite } r \ v \ \{ (). r \mapsto v \}_{tid, \Phi}
 \end{array}$$

Fig. 13. Selected proof rules of AxSL^{SC}

5.2.1 Proof Rules for Microinstructions. Fig. 13 depicts three slightly specialised proof rules for microinstructions MemRead, MemWrite, and RegWrite. We first explain the rule **SC-HT-MICRO-MEMREAD** for MemRead in detail, which is proved sound against the opax rule **S-MEM-READ**.

Our microinstruction Hoare triple for MemRead has the form $\{P\} \text{MemRead } x \ \{v. Q\}_{tid, \Phi}$, which states that, for a MemRead on thread tid following protocol Φ , if one provides the resources specified in the precondition P , then the MemRead results in the updated resources Q of the postcondition, which can refer to the resulting read value v passed to the continuation to continue reasoning about the thread. The event ID of the associated memory read event is existentially quantified in Q as e .

The rule has two main aspects, as do the other low-level rules for MemRead and MemWrite: low-level graph reasoning, and high-level resource transfer.

First, for directly conducting graph reasoning with respect to the axioms of the memory model, we use a bookkeeping assertion $\textcircled{1} \text{PoPred}(e_{\text{po}})$ to capture the e_{po} parts of the opax thread state T . Intuitively, the assertion keeps track of the event that will become the po immediate predecessor of the memory event e that associates with the next microinstruction, and thus allow us to conclude facts about new incoming po edge e_{po} po e , that is included in $\text{GraphFactsR}(e, x, v, e_w, e_{\text{po}})$. It gets updated accordingly in the postcondition, as $\textcircled{5}$. The predicate $\textcircled{2} \text{GraphFactsR}$ includes all graph assertions we can conclude for the read event:

$$\text{GraphFactsR}(e, x, v, e_w, e_{\text{po}}) \triangleq e:\text{R } x \ v * e_{\text{po}} \text{ po } e * e_w \text{ rf } e * e_w:\text{W } x \ v$$

Besides the incoming po edge, it also includes an event assertion $e:\text{R } x \ v$ indicating that the event is assigned with ID e , and another edge assertion $e_w \text{ rf } e$ to relate the read with the write e_w that it is reading from. It is worth noting that all graph assertions are persistent knowledge that can be freely duplicated, which reflects the fact that the graph is fixed once guessed in opax. Furthermore, we can stop the reasoning (by contradiction) if some graph facts gathered together imply a violation of the validity predicate of the axiomatic memory model,

since this further implies that we are reasoning about a graph that does not represent a valid execution result of the program. We demonstrate this idiom in the message passing example below.

Second, to support high-level reasoning via resource transfer, the precondition has a user-supplied $\textcircled{4} P(e, x, v, e_w, e_{po})$ on the right side of an Iris view shift \Rightarrow (a separation implication that permits resource update), which also appears in the post condition, meaning that P is obtained after the read. This view shift constrains what we can conclude in $\textcircled{4}$ given the facts about the new read and the protocol resource $\textcircled{3}$ received from the write e_w being read. Note that we have to give back the same protocol resource Φ after concluding P , to ensure that the resource remains available for other potential reads of the same write e_w . In a sense, this view shift allows us to *temporarily* obtain the ownership of the protocol resource of e_w at this memory read, akin to CSL invariants that can only be opened for a single program step. As we will see later in $\text{AxSL}^{\text{SCExt}}$, this view shift is a specialised form of the notion of flow implication.

The rule $\text{SC-HT-MICRO-MEMWRITE}$ for MemWrite is similar, except that Φ is on the right side of the view shift, indicating that it is sent away; and that the predicate $\text{GraphFactsW}(e, x, v, e_{po})$ includes the graph information for the write:

$$\text{GraphFactsW}(e, x, v, e_{po}) \triangleq e : \text{W } x \ v \ * \ e_{po} \ \text{po } e$$

The rule $\text{SC-HT-MICRO-REGWRITE}$ updates a points-to assertion for register r to mirror the change to $T.\text{regs}$ in the opax rule A-REG-WRITE .

5.2.2 Proof Rules for Instructions. Moving from microinstructions to the instructions composed out of them, we can write a derived high-level proof rule for each instruction. These high-level rules can be further specialised to specific programming idioms and their assumptions to make reasoning practical.

Before looking at the proof rules, we make our treatment of instructions precise. As usual, reasoning about a machine with instructions in (and fetched from) specific addresses in memory, rather than a language with an abstract syntax of statements, causes a slight impedance mismatch with Hoare logic: the thread state does not include instructions, but merely the address of the “current” instruction. However, a normal-looking Hoare triple can be recovered by using some indirection [Myreen et al. 2007; Myreen and Gordon 2007; Myreen et al. 2008][Myreen 2009, §3.4][Erbsen et al. 2021, §4.3][Liu et al. 2023]. We use Hoare triples for presentation, but use weakest preconditions in our formalisation. Our Hoare triples for a single instruction i are of the form $\{P\} a : i \{a' : Q\}_{tid, \Phi}$, where a is the address of the instruction, and a' is the address of the next instruction. This instruction triple is implemented using register points-to of pc: the precondition P is combined with $\text{pc} \mapsto a$, and Q is combined with $\text{pc} \mapsto a'$ for the appropriate a' — which is $a + 4$ (as per the elaboration of IncPC) except for branch instructions.

For presentation purposes, for programs without branching (and thus no looping), we can (as illustrated in Figure 14) conflate instruction instances with instructions, and thus conflate instruction identifiers a, b, c , etc. with numerical addresses for instructions in memory $a, a + 4, a + 8$, etc. For languages where an instruction instance leads to a single memory event (as we have so far), we can conflate instruction instance identifiers with memory event identifiers. In other cases, we use the counters of the opax semantics, although they can often be quantified over in reasoning rather than considered in detail, merely keeping the information that they are smaller than the current counter (and thus po-before the current event).

For instance, the rule SC-HT-INST-LDR in Fig. 15 is for a load instruction at address a with an immediate address x , which is elaborated into MemRead x followed by RegWrite followed by IncPC.

This rule is slightly specialised to only taking an immediate address, and is derived by a $\text{SC-HT-MICRO-MEMREAD}$ followed by and $\text{SC-HT-MICRO-REGWRITE}$. Similarly, one can prove a specialised instruction rule SC-HT-INST-STR for a store with immediate operators by $\text{SC-HT-MICRO-MEMWRITE}$. We use these two instruction rules in the message passing example below.

concrete addresses	symbolic addresses	instruction instances	memory events
1000:	a :	$\text{str } [x] \ 42$	$a:W \ x \ 42$
$1000 + 4$:	b :	$\text{str}_{\text{rel}} [y] \ 1$	$b:W_{\text{rel}} \ y \ 1$
$1000 + 8$:	c :	$r := \text{ldr } [z]$	$c:R \ z \ 2$

Fig. 14. Conflating (left columns) a numerical instruction address (1000, 1000+4, 1000+8) with a symbolic instances address (a, b, c), and (right columns) an instruction instance ($\text{str } [x] \ 42, \dots$) with its unique memory event ($W \ x \ 42, \dots$).

$$\begin{array}{l}
 \text{SC-HT-INST-LDR} \\
 \left\{ r \xrightarrow{\text{I}} _ * \text{PoPred}(e_{\text{po}}) * \forall e, v, e_w. \left(\begin{array}{l} \text{GraphFactsR}(e, x, v, e_w, e_{\text{po}}) * \Phi(x, v, e_w) \\ \Rightarrow P(e, x, v, e_w, e_{\text{po}}) * \Phi(x, v, e_w) \end{array} \right) \right\} \\
 a: r := \text{ldr } [x] \\
 \left\{ a + 4: r \xrightarrow{\text{I}} v * \exists e_w. \text{PoPred}(a) * \text{GraphFactsR}(a, x, v, e_w, e_{\text{po}}) * P(a, v, e_w) \right\}_{\text{tid}, \Phi} \\
 \\
 \text{SC-HT-INST-STR} \\
 \left\{ \text{PoPred}(e_{\text{po}}) * \forall e. (\text{GraphFactsW}(e, x, v, e_{\text{po}}) \Rightarrow \Phi(x, v, e)) \right\} \\
 a: \text{str } [x] \ v \\
 \left\{ a + 4: \text{PoPred}(a) * \text{GraphFactsW}(a, x, v, e_{\text{po}}) \right\}_{\text{tid}, \Phi}
 \end{array}$$

Fig. 15. Two instruction proof rules with instruction triples. $a + 4$ in the postcondition indicates the address of the next instruction.

5.2.3 *Message Passing.* We now demonstrate the graph reasoning capability that opax-based logics have, and how we achieve local reasoning with protocols, by exercising AxSL^{SC} on a message-passing litmus test.

The example has two threads: one sending, and one receiving. The sending thread writes a value (in this case 42) to a ‘data’ address in order to transfer it between threads, then writes 1 to a ‘flag’ address to indicate that the data write has been completed. The receiving thread reads from the flag address to check whether a message has been passed to it, and then reads from the data location.

$$\begin{array}{l}
 a: \text{str } [data] \ 42 \quad \parallel \quad c: r_1 := \text{ldr } [flag] \\
 b: \text{str } [flag] \ 1 \quad \parallel \quad d: r_2 := \text{ldr } [data]
 \end{array}$$

Fig. 16. MP in SC: $r_1 = 1 \rightarrow r_2 = 42$

Specification. We want to be able to prove that if the load of the flag reads 1, then the load of the data will read 42; formally, $r_1 \xrightarrow{\text{I}} v * r_2 \xrightarrow{\text{I}} v' * (v = 1 \rightarrow v' = 42)$.

Picking the protocol. The proof sketch of Fig. 17 relies on extensive graph reasoning. We first specify the protocol used in the proof. For the data address, we pick $\Phi(\text{data}, v, e) \triangleq \text{Initial}(e) \vee v = 42$, where $\text{Initial}(e)$ denotes that the event is an initial write that necessarily has value 0. It is not possible to require that v be 42 in all cases, because the initial write would then not satisfy the protocol. For the flag address, we pick

$$\Phi(\text{flag}, v, e) \triangleq \text{Initial}(e) \vee (v = 1 * e:W \ \text{flag} \ 1 * \exists e'. e':W \ \text{data} \ 42 * e' \ \text{po} \ e)$$

requiring that a non-initial write to the flag address is only allowed if it is a write of value 1, on the sending thread, which is po after a write of 42 to the data address.

Note that the only information we pass between threads with this protocol are persistent graph facts. It means that we can always duplicate and keep the complete protocol resources in the thread when reading. With the protocol in hand, we can give a proof sketch for each thread.

```

Sending thread:
1  {PoPred(-)}
2  GraphFactsW(a, data, 42, -)
   ⇒
3  Φ(data, 42, a) * a:W data 42
4  a: str [data] 42
5  {PoPred(a) * a:W data 42}
6  GraphFactsW(b, flag, 1, a) * a:W data 42
   ⇒
7  Φ(flag, 1, b)
8  b: str [flag] 1
9  {PoPred(b) * a:W data 42 * b:W flag 1 * a po b}
Receiving thread:
10 {PoPred(-) * r1 ↦1 _ * r2 ↦2 _}
11  ∀v, ew. GraphFactsR(c, flag, v, ew, -) * Φ(flag, v, ew)
   ⇒
12  Φ(flag, v, ew) * Φ(flag, v, ew) * c:R flag v * ew rf c
13 c: r1 := ldr [flag]
14  {PoPred(c) * ∃v, e. r1 ↦ v * Φ(flag, v, e) * c:R flag v * e rf c}
15  ∀v', ew. GraphFactsR(d, data, v', ew, c) * Φ(data, v', ew)
   ⇒
16  Φ(data, v', ew) * (v = 1 → v' = 42)
17 d: r2 := ldr [data]
18  {PoPred(d) * ∃v, v'. r1 ↦ v * r2 ↦ v' * (v = 1 → v' = 42)}

```

Fig. 17. Proof sketch of MP in AxSL^{SC}.

Sending thread. On the sending thread, we are required to show first that the data write a satisfies the protocol on $data$, which we can do straightforwardly because the second disjunct of the protocol only requires that the write has value 42. We apply **SC-HT-INST-STR** twice for the two writes. At the write of the data, we learn $\text{PoPred}(a)$ and $a:W\ data\ 42$ in line 5. We are then required to show that the flag write satisfies the flag protocol. We can do so by instantiating the existential on the right hand side of the protocol with a , which is a write to $data$ of 42, as required, and can be shown to be po-before the current event because it is the current po-predecessor. We illustrate how to unfold the two instructions to resources, in particular graph facts, in the proof in Fig. 18.

Receiving thread. The receiving thread is where interesting graph reasoning happens. We apply **SC-HT-INST-LDR** twice for the two reads. We read from the flag, learning $r_1 \mapsto v$ and $\Phi(flag, v, e)$ for some write e in line 14 (we get Φ by duplicating it in line 12). Finally, we consider the data read. We learn $r_2 \mapsto v'$ and $\Phi(data, v', e_w)$ for some v' and e_w , and are required to prove $v = 1 \rightarrow v' = 42$. The graph reasoning in line 15 & 16 starts by case splitting on $\Phi(data, v', e_w)$. We have $v' = 42$ immediately in the right case. In the left case, we derive a contradiction from $\text{Initial}(e_w)$ and $v = 1$ with graph facts. That is, an fr from d to e' , the write to data in the sending thread, can be

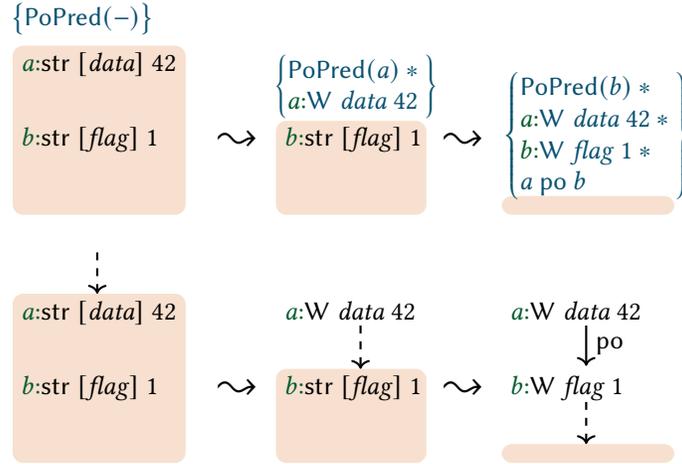


Fig. 18. The proof (here, of the writer side of MP) in progress (proof steps are indicated by \rightsquigarrow) unfolds the program into a graph: either implicitly via assertions $\{ \dots \}$ in the top row, or visualised as an explicit graph in the bottom row.

derived given e_w being the initial write, which closes a hb cycle

$$d \text{ fr } e' \text{ po } e \text{ rf } c \text{ po } d$$

violating the SC axiom.

5.3 Tracking Flow of Resources: The AxSL^{SCExt} Logic

AxSL^{SCExt} extends the syntax of AxSL^{SC} with a notion of tied-to assertion: $a \leftrightarrow P$ means that P is *tied* to event a . This extension allows us to track sources of resources and how resources *flow* between individual events precisely, instead of mixing resources regardless of their sources as in AxSL^{SC}. It is worth noting that this does not add any additional expressive power to the logic - in fact AxSL^{SC} is already a fine logic for SC - but is a robust solution that is also applicable to the relaxed concurrency of Arm-A for which an AxSL^{SC}-like construction is unsound. Recall that in SC, sc is the acyclic synchronisation relation, thus passing resources along it (or its subrelations) is sound. In fact, for thread local reasoning, it suffices to flow resources along po , which is included in sc and is also the reasoning order. In AxSL^{SC} (and other CSLs for SC), we implicitly unify the resource flowing order and the reasoning order, but in AxSL^{SCExt} we separate them by making the former *explicit* with the tied-to assertion. We elaborate this idea with the AxSL^{SCExt} version of AxSL^{SC} rules depicted in Fig. 13.

5.3.1 Rules for Tied-to assertions. To demonstrate how tied-to assertions work, we present part of their user interface in Fig. 19. The presented rules can be used to update the assertions as needed at any point in a program

$$\begin{array}{c}
 \text{TIED-SEP-SPLIT} \\
 \frac{a \leftrightarrow (P * Q)}{\models^i (a \leftrightarrow P * a \leftrightarrow Q)} \\
 \\
 \text{TIED-SEP-COMBINE} \\
 \frac{(a \leftrightarrow P) * (a \leftrightarrow Q)}{\models^i a \leftrightarrow (P * Q)} \\
 \\
 \text{TIED-WAND} \\
 \frac{\Box (P * Q) \quad a \leftrightarrow P}{\models^i a \leftrightarrow Q} \\
 \\
 \text{HT-IUPD} \\
 \frac{\{\models^i P\} e \{Q\}_{tid, \Phi}}{\{P\} e \{Q\}_{tid, \Phi}}
 \end{array}$$

Fig. 19. Selected rules for tied-to assertions and their interactions with the Hoare triple

proof. The update rules are relatively standard except for the required interpretation modality \Rightarrow^i : rules **TIED-SEP-SPLIT** and **TIED-SEP-COMBINE** make tied-to distributive over separation conjunction; **TIED-WAND** updates tied assertions with persistent resources. The logical meaning of tied-to assertions relies on their logical interpretation (using higher-order ghost states) that resides within the Hoare triple (see §6.4.3), so updating tied-to assertions requires access to this interpretation. Following previous work[Gähler et al. 2022; Spies et al. 2021; Timany et al. 2024], we use the interpretation modality to abstract this technicality: the rules require only the modality, which is then stripped using rule **HT-IUPD**. Other essential tied-to update operations, such as allocation, require additional soundness constraints and can only be performed by proof rules.

5.3.2 Proof Rules for Microinstructions. There are substantial similarities between $\text{AxSL}^{\text{SCExt}}$ rules and their AxSL^{SC} counterparts. We have the same bookkeeping assertions for the immediate po predecessor, which is updated in the post condition; and the same clauses universally quantified by the new event e accompanied by the graph fact about it. The two main distinctions are the use of the tied-to assertions, and the new FlowSCX predicates for high-level resource transfer between nodes. Let us take a closer look at them in the rule **SCEXT-HT-MICRO-MEMREAD** in Fig. 20.

$$\begin{array}{c}
\text{SCEXT-HT-MICRO-MEMREAD} \\
\left\{ \text{PoPred}(e_{\text{po}}) * \textcircled{1} *_{(e'_{\text{po}} \mapsto P_{\text{po}}) \in m} (e'_{\text{po}} \rightsquigarrow P_{\text{po}}) * \right. \\
\left. \forall e, v, e_w. \left(\text{GraphFactsR}(e, x, v, e_w, e_{\text{po}}) \textcircled{2} * \text{Po}(\text{dom}(m), e) * \text{FlowSCR}(\Phi, e, x, v, e_w, m, P) \right) \right\} \\
\text{MemRead } x \\
\{v. \exists e, e_w. \text{PoPred}(e) * \text{GraphFacts}(e, x, v, e_w, e_{\text{po}}) * \textcircled{4} e \rightsquigarrow P(x, v, e_w)\}_{\text{tid}, \Phi} \\
\\
\text{SCEXT-HT-MICRO-MEMWRITE} \\
\left\{ \text{PoPred}(e_{\text{po}}) * *_{(e_{\text{po}} \mapsto P_{\text{po}}) \in m} (e_{\text{po}} \rightsquigarrow P_{\text{po}}) * \right. \\
\left. \forall e. \left(\text{GraphFactsW}(e, x, v, e_{\text{po}}) \textcircled{2} * \text{Po}(\text{dom}(m), e) * \text{FlowSCW}(\Phi, e, x, v, m, P) \right) \right\} \\
\text{MemWrite } x \ v \\
\{(). \exists e. \text{PoPred}(e) * \text{GraphFacts}(e, x, v, e_{\text{po}}) * e \rightsquigarrow P(x)\}_{\text{tid}, \Phi} \\
\\
\text{FlowSCR}_{\Phi}(e, x, v, e_w, m, P) \triangleq \left(\left(*_{(_ \mapsto P_{\text{po}}) \in m} P_{\text{po}} \right) * \Phi(x, v, e_w) \right) \Rightarrow \Phi(x, v, e_w) * P(x, v, e_w) \\
\text{FlowSCW}_{\Phi}(e, x, v, m, P) \triangleq \left(*_{(_ \mapsto P_{\text{po}}) \in m} P_{\text{po}} \right) \Rightarrow \Phi(x, v, e_w) * P(x)
\end{array}$$

Fig. 20. $\text{AxSL}^{\text{SCExt}}$ version of the AxSL^{SC} rules shown in Fig. 13. The new syntax and changes are highlighted. The user provides m , a thread-local map from events to the resources consumed.

This rule allows us to explicitly reason about resources flowing to e along po edges. If there is such an incoming edge, say from e' to e , and we have an $e' \rightsquigarrow P$, then the flow implication FlowSCR can use P in its premise. In total, the resources that flow into e , and thus are considered in the flow implication for e , consist of (the separating conjunction of) all such local resources (which need not be persistent) that flow along po edges, as collected in the partial event-to-resource map m (for thread-internal resource flow), combined with the (usually persistent) resources flowing from external events (here, the quantified external write e_w that e is reading from), as specified by the protocol Φ .

To apply the rule, the user has to supply a finite map m to specify how to flow thread-local resources to the current node to show the flow implication (③). Assertion ① $\ast_{(e'_{po} \mapsto P_{po}) \in m} (e'_{po} \rightsquigarrow P_{po})$, collecting the resources in m for the premise of the flow implication. The map m is constrained by assertion ② in the hypothetical reasoning, which requires that an event e'_{po} can only occur in the domain of m when there will be (given the graph facts) a po edge to the new MemRead event. Finally, as a result of the hypothetical reasoning on the last line of the precondition, we get the (user-supplied) result $P(x, v, e_w)$ of the flow implications ③, tied to the new memory event e , as ④. This flow implication FlowSCR replicates the view shift of **SC-HT-MICRO-MEMREAD** except for the now explicit local resource transfer from po predecessors (the iterated separation conjunction). Both this and FlowSCW for MemWrite are instances of the general definition of the flow implication.

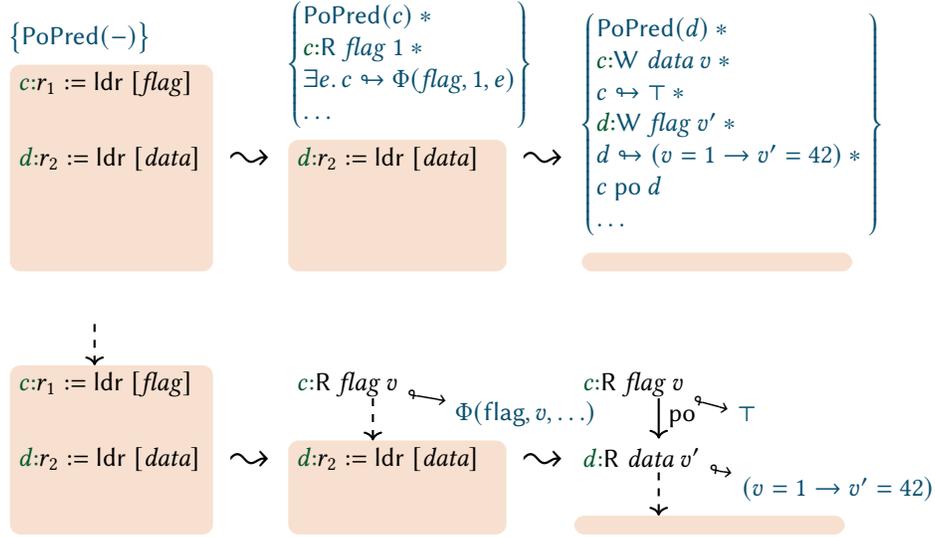


Fig. 21. The proof (here, of the reader side of MP, in AxSL^{SCExt}) in progress. Resources tied-to a node ‘dangle’ off it. Between the second \rightsquigarrow , we flow $\Phi(flag, v, \dots)$ from c to d along po, by consuming it from c and using it to conclude $(v = 1 \rightarrow v' = 42)$ at d .

5.3.3 Message Passing. We revisit MP to showcase the resource flow reasoning with tied-tos in AxSL^{SCExt}. This time, we use the same protocol, but adapt the specification (changes **highlighted**) to account for explicit resource flowing, as follows:

$$r_1 \vdash v \ast r_2 \vdash v' \ast \exists e. e \rightsquigarrow (v = 1 \rightarrow v' = 42) \ast (\exists e'. e \text{ po } e' \vee e = e' \ast \text{PoPred}(e'))$$

In general, P holds whenever $a \rightsquigarrow P$ appears in a post condition, since one can flow P from a to a hypothetical terminating (po-last) event along po. Our actual Hoare triple allows omitting the tied-to assertion, so that one can write post conditions as usual, see §6.4.5 for how the model of the logic enables this. We however write tied-to assertions explicitly here for better clarity. In this new specification, the implication $v = 1 \rightarrow v' = 42$ is embedded in a tied-to for an event e , which is po-before the hypothetical terminating event, as stated in the last clause.

We look at a proof sketch of the receiving thread depicted in Fig. 22 (with an illustration in Fig. 21, and proof rules in Fig. 20); the proof of the sending thread is nearly identical to that of AxSL^{SC} as no local resource flowing is happening in the thread. This time in the receiving thread, to get the protocol resource $\Phi(flag, \dots)$ from the *flag* load, we have to justify FlowSCR with argument P being instantiated with Φ , as in line 2 & 3. After that,

in line 4, the transferred Φ is tied to the memory read c , which claims a constraint that one can only access Φ if is po after c . To conclude the specification with the graph facts in Φ , we have to flow it to the second load. We therefore let the user-assigned m of the read rule be a singleton map [$c \mapsto \Phi(flag, \dots)$], and show that c is indeed a po predecessor of d by $\text{PoPred}(c)$ ². Then, we show the flow implication FlowSCR for the load of *data* in line 6 & 7, with P being $v = 1 \rightarrow v' = 42$, where we perform the same graph reasoning as in the AxSL^{SC} MP proof. Finally, we have the desired implication tied to event d , which suffices to show the specification.

$$\begin{array}{l}
1 \quad \{\text{PoPred}(-) * r_1 \mapsto _ * r_2 \mapsto _ \} \\
2 \quad \forall v, e_w. \Phi(flag, v, e_w) \\
\quad \Rightarrow \\
3 \quad \Phi(flag, v, e_w) * \Phi(flag, v, e_w) \\
4 \text{ c: } r_1 := \text{ldr } [flag] \\
5 \quad \{\text{PoPred}(c) * \exists v, e. r_1 \mapsto v * c \mapsto \Phi(flag, v, e) * c:\text{R } flag \ v * e \text{ rf } c \} \\
6 \quad \forall v', e_w. \text{GraphFactsR}(d, data, v', c) * \Phi(data, v', e_w) * \Phi(flag, v, e) \\
\quad \Rightarrow \\
7 \quad \Phi(data, v', e_w) * (v = 1 \rightarrow v' = 42) \\
8 \text{ d: } r_2 := \text{ldr } [data] \\
9 \quad \{\text{PoPred}(d) * \exists v, v'. r_1 \mapsto v * r_2 \mapsto v' * d \mapsto (v = 1 \rightarrow v' = 42) \}
\end{array}$$

Fig. 22. Proof sketch of the receiving thread of MP in $\text{AxSL}^{\text{SCExt}}$. The explicit reasoning of resource flow is highlighted.

5.4 Handling Arm-A Concurrency: The AxSL^{Arm} Logic

Using $\text{AxSL}^{\text{SCExt}}$ for TinySc as the base, we build AxSL^{Arm} for TinyArm. The changes we make to $\text{AxSL}^{\text{SCExt}}$ to obtain AxSL^{Arm} are twofold. First, we extend the assertion language. Similar to how we extend the opax semantics of TinySc to obtain the counterpart of TinyArm, we add new assertions dedicated to reason about Arm's dependencies, etc., and augment existing ones. Second, we change the resource flowing order. We shift from reasoning about resource flow along po to Arm's lob to reflect the change of the synchronisation order from SC's sc to Arm's ob. Arm-A's relaxed concurrency is fundamentally different from SC in the sense that po in Arm-A does not impose an intra-thread ordering, which implies that transferring resources along po is unsound. Thus, in AxSL^{Arm} , we instead rely on lob, the local fragment of which enforces synchronisation. We elaborate on these two changes with the rule for MemRead in §5.4.1.

5.4.1 Proof Rules for Microinstructions. We now explain one of the key proof rules $\text{HT-MICRO-MEMREAD-RDEP-EXT}$ in Fig. 23 for a MemRead with ordering strength os , variety vr , address x , and address dependencies d , focusing on the new components introduced to handle Arm-A's memory model. To keep the exposition manageable, the rule is specialised to a read from a distinct thread (an external read), with empty intra-instruction dependencies (i.e. only syntactic register dependencies). We illustrate this rule schematically in Fig. 24.

The first two clauses of the precondition capture the specialisation mentioned above: ① $\text{NoLocalWrites}(x)$ captures the fact that there are no thread-local writes on the same address x up until this point in the program order; with this in hand, one knows that only external writes can be read from by this MemRead, making resource transfer along rfe possible. This matters because rfe enforces inter-thread synchronisation, whereas rfi in Arm-A does not. This fact is unchanged after the MemRead, and is thus restored by the postcondition as ⑧. ② $d = (\text{dom}(regs), \emptyset)$ requires that this MemRead depends on *exactly* the registers in the domain of the (partial) register file $regs$ of ④ (non-involved registers can be framed off to apply this rule via the normal

²In this example, the protocol resource we flow is persistent, but we can also flow non-persistent resources

$$\begin{array}{l}
 \text{HT-MICRO-MEMREAD-RDEP-EXT} \\
 \left\{ \begin{array}{l}
 \textcircled{1} \text{NoLocalWrites}(x) * \textcircled{2} d = (\text{dom}(\text{regs}), \emptyset) * \text{PoPred}(e_{\text{po}}) * \textcircled{3} \text{CtrlPreds}(\text{srcs}_{\text{ctrl}}) * \\
 \textcircled{4} *_{(r \mapsto (v, E)) \in \text{regs}} r \mapsto v @ E * *_{(e_{\text{lob}} \mapsto P_{\text{lob}}) \in m} (e_{\text{lob}} \rightsquigarrow P_{\text{lob}}) * \\
 \forall e, v, e_w. \left(\textcircled{5} \text{GraphFacts}(e, \text{os}, \text{vr}, x, v, e_w, e_{\text{po}}, \text{srcs}_{\text{ctrl}}, d, \text{regs}) \text{ -*} \right) \\
 \left(\textcircled{6} \text{Lob}(\text{dom}(m), e) * \textcircled{7} \text{FlowR}(\Phi, e, x, v, e_w, m, P) \right)
 \end{array} \right\} \\
 \text{MemRead os vr x d} \\
 \left\{ \begin{array}{l}
 \exists e, e_w. D = \{e\} * \textcircled{8} \text{NoLocalWrites}(x) * \text{PoPred}(e) * \text{CtrlPreds}(\text{srcs}_{\text{ctrl}}) * \\
 (v, D). \textcircled{9} *_{(r \mapsto (v, E)) \in \text{regs}} r \mapsto v @ E * \text{GraphFacts}(e, \text{os}, \text{vr}, x, v, e_w, e_{\text{po}}, \text{srcs}_{\text{ctrl}}, d, \text{regs}) * \\
 e \rightsquigarrow P(x, v, e_w)
 \end{array} \right\} \text{tid}, \Phi
 \end{array}$$

Fig. 23. A proof rule of AxSL^{Arm} for the MemRead microinstruction, specialised for a thread that has no writes to the location read. The changes to SCEXT-HT-MICRO-MEMREAD are highlighted.

frame rule familiar from separation logics), and not on any intra-instruction memory read (the \emptyset of event IDs). The collection of register points-to for regs of $\textcircled{4}$ is unchanged and thus restored in the postcondition as $\textcircled{9}$. Note that a register points-to assertion now also maps a register to a set of events E that are the sources of its data, with notation $@E$. This change captures the corresponding extension to the register file in the opax semantics. The GraphFacts predicate in $\textcircled{5}$ now takes more arguments, and gives more graph assertions like the dependency edges. Bookkeeping assertion $\textcircled{3}$ CtrlPreds($\text{srcs}_{\text{ctrl}}$) captures the $\text{srcs}_{\text{ctrl}}$ part of the opax thread state T which works like PoPred. The user-supplied m constrained by Po($\text{dom}(m), e$) in SCEXT-HT-MICRO-MEMREAD is now constrained by $\textcircled{6}$ Lob which instead requires the domain of m to consist of lob predecessors of the read. The flow equation $\textcircled{7}$ and FlowSCR explained in AxSL^{SCExt} share the same definition.

The post condition is parametrised by an additional dependency set D (the set of event IDs that v stems from) for intra-instruction dependencies. In this rule, this set is $\{e\}$, where e is the existentially quantified event ID of the associated memory read event, which is passed to the continuation, to, for instance, a register read to establish the dependency from this read event to the register's data.

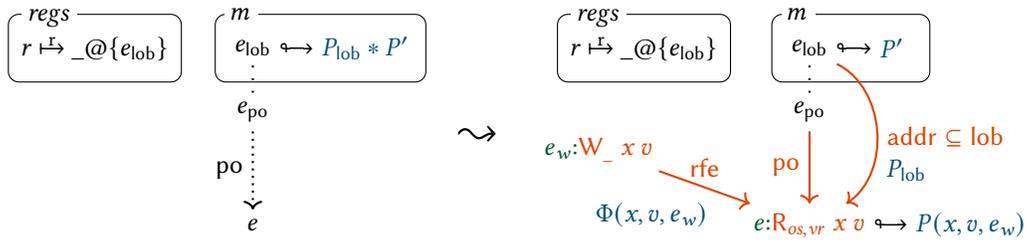


Fig. 24. A visualisation of HT-MICRO-MEMREAD-RDEP-EXT for external read event e , with logical annotations and newly gained graph facts highlighted. \rightsquigarrow indicates the update of resources from the precondition to the postcondition. The protocol specifies $\Phi(x, v, e_w)$ transferred along the rfe edge to e , which is returned to be transferred to other potential reads. Local events e_{lob} , which are shown to become lob predecessors due to the use of some register r , let the corresponding P_{lob} resource, previously tied to e_{lob} , flow along the newly-learned lob edge, to be combined with the protocol, and the result $P(x, v, e_w)$ gets tied to e .

5.4.2 *Specialising Proof Rules for Examples.* The Arm-A memory model is intrinsically complex, so there is no ‘perfect’ rule that is both simple and general. Therefore, we again derive some further specialised instruction rules for reasoning about examples.

We explain how AxSL^{Arm} can be used to reason about Arm-A concurrency compositionally, using two key examples: message-passing (§5.5) and versions of load buffering (§5.6), describing our high-level proof rules along the way. These examples also demonstrate two styles of proof: a low-level proof for (Arm’s version of) MP that is similar to that of $\text{AxSL}^{\text{SCExt}}$, showing how one can tackle subtle reasoning about the Arm-A memory model if needed; and a high-level proof for LB that abstracts physical state with simple, easy-to-use ghost state, demonstrating the convenience of AxSL^{Arm} . For MP, we only sketch the proof of the receiving thread, highlighting the differences to the MP proof done in $\text{AxSL}^{\text{SCExt}}$. For LB, we go through the proof in more detail, and present the specialised rules for the instructions used in the proof along the way, and explain how it fails (as desired) if the necessary synchronisation is removed.

5.5 Graph Reasoning in AxSL^{Arm} : MP+rel+addr

We give an example illustrating more complex reasoning about the memory event graph, on an Arm-A version of message passing. Comparing to the SC counterpart, two synchronisations are inserted to ensure intra-thread order. The flag write of the sending thread is a release write to ensure that the two writes are suitably ordered (this can also be achieved in other ways, for example with a dmb st). The two reads of the receiving thread are ordered by using the result of the flag read to compute the address of the data read, resulting in an address dependency between the two reads. In this example, the address dependency is artificial, but similar shapes arise naturally when a message-passing idiom is used to transfer a pointer to a data structure between threads.

$$\begin{array}{l} a: \text{str } [data] \ 42 \\ b: \text{str}_{\text{rel}} [flag] \ 1 \end{array} \parallel \begin{array}{l} c: r_1 := \text{ldr } [flag] \\ d: \\ r_2 := \text{ldr } [data + r_1 - r_1] \end{array}$$

Fig. 25. MP+rel+addr: $r_1 = 1 \rightarrow r_2 = 42$

Overview. The proof has roughly the same shape as the previous ones, in particular we conclude r_2 is 42 with a contradiction on the validity of the graph in the case of reading the initial value 0 from *data* when the flag is set.

Protocol. The protocol we specify is almost identical to the one used for the SC variant, except for two minor adjustments (highlighted in orange) at *flag* due to the now relaxed memory model.

$$\Phi(flag, v, e) \triangleq \text{Initial}(e) \vee (v = 1 * e:W \ flag \ 1 * (tid(e) = 1) * \exists e'. e':W_{\text{rel}} \ data \ 42 * e' \text{ po } e)$$

First, we additionally require in the right disjunct that the write is from Thread 1 (the sending thread) with $tid(e) = 1$. This allows us to conclude that the load of *flag* on the receiving side has to read from this *external* write, resulting in an rfe edge which contributes to the later graph reasoning because it ensures inter-thread synchronisation. Second, for the same reason, we strengthen the write to *flag* to be a release write, which allows us to conclude a lob edge between the two writes of the sending thread.

Proof Sketch. We look at the proof sketch of the receiving thread, focusing on the reasoning of the data dependency between two loads, and the Arm version of the contradiction proof. We learn $r_1 \mapsto v @ \{c\}$ after reading from the flag, and $c \mapsto \Phi(flag, v, b)$ for some b . The $@\{c\}$ part of the register points-to tells us that the value v comes from c . Therefore, when we consider the data read, we have $c \ data \ d$ because of the artificial data dependency using r_1 . The fact that data is in lob allows us to use the resources tied to c . We learn $r_2 \mapsto v' @ \{d\}$ and $d \mapsto \Phi(data, v', e)$ for some v' and e , and are required to prove $v = 1 \rightarrow v' = 42$. Case splitting on $\Phi(data, v', e)$, we have $v' = 42$ immediately in the right case. In the left case, we derive a contradiction from $\text{Initial}(e)$ and $v = 1$. In detail, the contradiction is as follows: from $\Phi(flag, 1, b)$ we learn $b:W_{\text{rel}} \ flag \ 1$ and $a:W \ data \ 42$ for some a such that $a \text{ po } b$. Since e is an initial write to the same address as a , we know $e \text{ co } a$, and since $e \text{ rf } d$, we know $d \text{ fr } a$

and therefore $d \text{ ob } a$, since d and a are on different threads. Because b is a release write po-after a , we have $a \text{ ob } b$; because c reads from b on a different thread (recall the new bits of the protocol), we have $b \text{ ob } c$; and finally, because there is a data dependency between c and d , we have $c \text{ ob } d$. By transitivity of ob , we obtain $a \text{ ob } d$, and together with $d \text{ ob } a$, we obtain $d \text{ ob } d$, which contradicts the irreflexivity of ob .

5.6 High-Level Reasoning in AxSL^{Arm}: Load Buffering

We detail how to obtain a high-level proof of LB+artificialdata+data (Fig. 26), a version of the LB litmus test with an artificial (but still architecturally respected) data dependency on Thread 1, and a normal data dependency on Thread 2. This version of LB is interesting because this specification cannot be proven using just an invariant about the *values* of the write, as it can for LB+datas [Jeffrey and Riely 2016, §6]: it requires reasoning about the *order* of the writes.

Specification. We would like to show that the example exhibits no load buffering behaviour on Arm-A, i.e., that Thread 1 cannot read 1 (while Thread 2 can read either the initial value 0, or the value 1 that Thread 1 writes), as in the Fig. 26 specification. We give two versions of the postcondition: one still involving tied assertions, corresponding to the final state of the threads, and one where the assertions have been pulled out, as used in our formal definition of weakest preconditions, as we describe in §6.4.5.

Thread 1	Thread 2
$\{r_1 \mapsto _ \}$	$\{r_2 \mapsto _ \}$
$a: r_1 := \text{ldr } [x]$	$c: r_2 := \text{ldr } [y]$
$b: \text{str } [y] \ (1 + r_1 - r_1)$	$d: \text{str } [x] \ r_2$
$\{\exists v_1. r_1 \mapsto v_1 @ \{a\} * a \leftrightarrow (v_1 = 0)\}$	$\{\exists v_2. r_2 \mapsto v_2 @ \{c\} * c \leftrightarrow (v_2 = 0 \vee v_2 = 1)\}$
$\{r_1 \mapsto 0 @ _ \}$	$\{\exists v_2. r_2 \mapsto v_2 @ _ * (v_2 = 0 \vee v_2 = 1)\}$

Fig. 26. LB+artificialdata+data and its (informal) specification

Protocol. The first step of the proof is to come up with an appropriate protocol Φ that abstracts the interference of the threads, and thus enables thread-modular reasoning. For this LB shape, it suffices to transfer the information that the write of 1 by Thread 1 has been executed, in the sense that this write is ob -before the event to which this information is tied. To capture this logically, we use a simple form of ghost state: the ‘oneshot resource algebra’ of Iris [Jung et al. 2018, §2]. The oneshot has two states: pending represents the exclusive permission to make a decision to choose a value, and $\text{shot}(v)$ represents the information that the decision has been made to choose value v . In particular, $\{\text{pending}\} * \{\text{pending}\}$ is a contradiction, and so is $\{\text{pending}\} * \{\text{shot}(v)\}$, but we can view-shift pending into $\text{shot}(v)$, $\{\text{pending}\} \Rightarrow \{\text{shot}(v)\}$, to express the logical decision to commit to v . Using this, we can formalise our protocol: for both locations x and y , either the value is 0, or it is 1, in which case we also have $\{\text{shot}(1)\}$.

Proof sketch. Using this protocol, the proof follows the sketch in Fig. 27 using the specialised instruction proof rules explained later in §5.6.1. On line 1, we give Thread 1 the exclusive permission pending to choose a value, which it will need when it writes 1; moreover, we will use pending in the flow implication of the load from x in line 4 to show that it must read 0.³ Line 2 states the incoming resources of the flow implication: the disjunction obtained from the protocol Φ , and the pending from the context. In the flow implication, we can then do a case

³We also start Thread 1 with the knowledge that it has made no writes to location x , and symmetrically Thread 2 to y , to exclude an internal reads-from. Internal reads-from do not imply ob on Arm-A, and thus has a significantly weaker premise for its flow implication, without $\Phi(x, v_1, _)$.

$$\Phi(_, v, _) \triangleq v = 0 \vee (v = 1 * \boxed{\text{shot}(1)})$$

Thread 1:

- 1 $\{r_1 \mapsto _ * \text{NoLocalWrites}(x) * \boxed{\text{pending}} * \dots\}$
- 2 $\left(v_1 = 0 \vee \left(v_1 = 1 * \boxed{\text{shot}(1)} \right) \right) * \boxed{\text{pending}}$
- \Rightarrow
- 3 $v_1 = 0 * \boxed{\text{pending}}$
- 4 $a: r_1 := \text{ldr } [x]$
- 5 $\left\{ \exists v_1. r_1 \mapsto v_1 @ \{a\} * a \rightsquigarrow \left(v_1 = 0 * \boxed{\text{pending}} \right) * \text{NoLocalWrites}(x) * \dots \right\}$
- 6 $\boxed{\text{pending}}$
- \Rightarrow
- 7 $\boxed{\text{shot}(1)}$
- 8 $b: \text{str } [y] (1 + r_1 - r_1)$
- 9 $\{ \exists v_1. r_1 \mapsto v_1 @ \{a\} * a \rightsquigarrow (v_1 = 0) \}$

Thread 2:

- 10 $\{r_2 \mapsto _ * \text{NoLocalWrites}(y) * \dots\}$
- 11 $c: r_2 := \text{ldr } [y]$
- 12 $\{ \exists v_2. r_2 \mapsto v_2 @ \{c\} * c \rightsquigarrow \Phi(y, v_2, c) * \text{NoLocalWrites}(y) * \dots \}$
- 13 $d: \text{str } [x] r_2$
- 14 $\left\{ \exists v_2. r_2 \mapsto v_2 @ \{c\} * c \rightsquigarrow \left(v_2 = 0 \vee \left(v_2 = 1 * \boxed{\text{shot}(1)} \right) \right) \right\}$

Fig. 27. Proof sketch of LB+artificialdata+data.

analysis on the disjunction, and in the case of $v_1 \neq 0$, we can derive a contradiction by combining pending with $\text{shot}(v)$; therefore, we must be in the case $v_1 = 0$, still with pending in hand, as per line 3. On line 5, because we have used the pending from the context in the flow implication for a , we get it back, but tied to a . This deals with the load. Now, for the store on line 8, we need, as part of the flow implication, to establish the protocol for the written value, 1. For our protocol, we have to take the second disjunct, and so we have to provide $\text{shot}(1)$. Because of the artificial dependency, the flow implication gives access to the resources tied to a , and thus to pending, as per line 6. The pending can be view-shifted, as part of the flow implication, into any $\text{shot}(v)$, and thus in particular into $\text{shot}(1)$, as per line 7. This satisfies the flow implication, and thus concludes the proof sketch for Thread 1.

Thread 2 relies on the same dependencies, but is much simpler: given our protocol Φ , we have $\Phi(y, v_2, c) = \Phi(x, v_2, d)$, so Thread 2 merely forwards this $\Phi(_, v_2, _)$ from the load to the store, which the flow implication for the write allows because of the data dependency.

Abstraction. Using the oneshot resource algebra allows us to reason thread-modularly: the proof of Thread 1 does not involve any graph reasoning about the intricacies of the Arm-A memory model, merely reasoning about abstract state. This is already useful for this small proof sketch (and the corresponding mechanised proof). Thread 2 does not require any inspection of the value or the resource being forwarded, merely plumbing through the flow implication, and the derivation of the contradiction in the impossible case of the load of Thread 1 relies on simple ghost theory. Moreover, the proof is quite flexible: the proof sketch only requires trivial modifications if, e.g., we replace the store of $1 + r_1 - r_1$ by a store of r_1 .

The proof sketch above crucially relies on the artificial data dependency of the store on the load, as it should: without it, if the store were merely $\text{str } [y] 1$, it could be executed out of order with respect to the load, thus

making the relaxed load behaviour observable. More concretely, without the dependency, the flow implication for the store would not include the resource tied to the memory read event a , and would therefore be of the shape $\top \Rightarrow \llbracket \text{shot}(1) \rrbracket$, which is not provable.

5.6.1 Proof Rules for Instructions. We explain the specialised instruction proof rules used in the LB proof, focusing on (highlighted in orange) how the rules for Arm-A (Fig. 28) differ from the (non-specialised) rules for SC (Fig. 15).

Load. The proof rule that we use for the load in both threads, **HT-INS-LDR-PLN-EXT**, is specialised to the instruction: a plain, non-exclusive load with an immediate address x . It is further specialised to the assumption that there are no prior writes to x by this thread, as otherwise the memory model guarantees no synchronisation and thus makes resource transfer unsound.

The first line of the precondition deals with bookkeeping of the po-predecessor, the local writes, and the register r that will be written to by the load.

The next line requires the flow implication for this instruction, which flows the resources of protocol Φ for this address to R that will be tied to this event — with the appropriate address, value, and memory event parameters. The postcondition then updates the bookkeeping accordingly, and keeps the fact that this thread has no writes to x . The key part of the postcondition, highlighted on the next line, is that R is then tied the new memory read event e (①), which is the source of the contents of register r (②).

Store. The proof rule that we use for the store in Thread 1, **HT-INS-STR-PLN-ARTIFICIALDATA**, is similarly specialised to the instruction: a plain, non-exclusive store with an immediate address x and an artificial value dependency on register r with result v . It is further specialised to the assumption that there is an assertion P that is tied to the source of register r . Again, the first lines of the precondition deal with the po bookkeeping and the latest write to the address. The key of the precondition, highlighted on line 2, requires ① ownership of register r together with the knowledge that its source is some memory event e_d , ② P is tied to the source memory event e_d , and ③ the flow implication, which flows the resource P into the protocol for the written value. Since P has been consumed by the instruction, it is absent from the postcondition.

The proof rule for the store in Thread 2 is almost identical, merely requiring knowing the value v' of register r , and the flow implication requiring the protocol be established for that value, $\Phi(x, v', e)$.

$$\begin{array}{c}
 \text{HT-INS-LDR-PLN-EXT} \\
 \left\{ \text{PoPred}(e_{po}) * \text{NoLocalWrites}(x) * r \mapsto _ * \right. \\
 \left. \forall e, v, e_w. (\Phi(x, v, e_w) \Rightarrow R(x, v, e_w) * \Phi(x, v, e_w)) \right\} \\
 a: r := \text{ldr } [x] \\
 \\
 \left\{ a + 4: \exists v, e_w. \text{① } a \mapsto (R(x, v, e_w)) * \text{② } r \mapsto v@ \{a\} \right\}_{tid, \Phi} \\
 \\
 \text{HT-INS-STR-PLN-ARTIFICIALDATA} \\
 \left\{ \text{PoPred}(e_{po}) * (\text{NoLocalWrites}(x) \vee \text{LastLocalWrite}(x, _)) * \right. \\
 \left. \text{① } r \mapsto _@ \{e_d\} * \text{② } e_d \mapsto P * \text{③ } \forall e. (P \Rightarrow \Phi(x, v, e)) \right\} \\
 a: \text{str } [x] (v + r - r) \\
 \\
 \left\{ a + 4: \text{PoPred}(a) * \text{LastLocalWrite}(x, a) * r \mapsto _@ \{e_d\} \right\}_{tid, \Phi}
 \end{array}$$

Fig. 28. Proof rules for the instructions in the left thread of LB+artificialdata+data.

5.7 Supporting Exclusives

We describe how the logic we have seen so far can be extended, with only minor changes, with new proof rules for Arm-A exclusives. We depict two highly specialised ones in Fig. 29.

$$\begin{array}{c}
\text{HT-INS-LDR-EXCL-EXT} \\
\left\{ \text{PoPred}(e_{\text{po}}) * \text{RmwPred}(x, -) * \text{NoLocalWrites}(x) * r \mapsto _ \right\} \\
a: r := \text{ldr}_x [x] \\
\left\{ a + 4: \exists v, e_w. \textcircled{1} a \mapsto (\text{RMWInv}(e_w, P)) * r \mapsto v@ \{a\} * e_w \text{ rf } a \right\}_{\text{tid}, \Phi} \\
\\
\text{HT-INS-STR-EXCL} \\
\left\{ \begin{array}{l} \text{PoPred}(e_{\text{po}}) * \text{RmwPred}(x, e_r) * \\ \text{NoLocalWrites}(x) * \\ r \mapsto _ @ _ * e_r \mapsto (\text{RMWInv}(e_w, P)) \end{array} \right\} \\
a: r := \text{str}_x [x] v \\
\left\{ a + 4: \exists v. r \mapsto v@ _ * \left(\begin{array}{l} v = 1 * \textcircled{2} a \mapsto P * e_r \text{ rmw } a * \\ \text{PoPred}(a) * \text{RmwPred}(x, -) * \\ \text{LastLocalWrite}(x, a) \end{array} \right) \vee \left(\begin{array}{l} v = 0 * \text{PoPred}(e_{\text{po}}) * \text{RmwPred}(x, e_r) * \\ \text{NoLocalWrites}(x) \end{array} \right) \right\}_{\text{tid}, \Phi}
\end{array}$$

Fig. 29. Highly specialised proof rules for the load/store-exclusive instructions, where we assume $\Phi(x, _ , e) \triangleq \text{RMWInv}(e, P)$. The RmwPred predicate tracks the latest local exclusive load on a location, which is used to induce rmw edges.

Arm-A features atomic read-modify-write operations in two forms: atomic instructions (compare-and-swap, fetch-and-add, etc.), and the combination of load-exclusive/store-exclusive pairs. The rules we have described so far only support ‘non-exclusive’ loads and stores. We now explain how we can give strong rules for read-modify-writes that support transfer of non-duplicable resources: a load exclusive a of v from x should, if the subsequent store exclusive succeeds, give a non-duplicable resource P which does not need to be given back because of the exclusivity.

To support atomic read-modify-writes, we do not need to change our definition of protocols, but merely to pick a protocol whose definition is of a specific shape, namely containing the RMWInv predicate. The $\text{RMWInv}(e, P)$ higher-order predicate, which is essentially an Iris invariant, takes a write e , and an arbitrary non-duplicable resource P whose ownership is released by the write and to be acquired by a successful read-modify-write. Since Iris invariants are duplicable, the RMWInv predicate can be transferred to the exclusive read from the protocol as in the $\text{HT-INS-LDR-EXCL-EXT}$ rule ($\textcircled{1}$). The invariant of RMWInv implements the escrow pattern [Kaiser et al. 2017] to trade an exclusive token for the non-duplicable resource P , with enough bookkeeping to capture the

$$\begin{array}{l}
\text{trylock}(\ell) \triangleq \\
r_1 := \text{ldr}_x [\ell] \\
\text{if } (r_1 \neq \text{unlocked}) \text{ return false} \\
r_2 := \text{str}_x [\ell] \text{ locked} \\
\text{dmb sy} \\
\text{return } (r_2 = \text{success})
\end{array}
\quad
\begin{array}{l}
\text{if trylock}(\ell) \{ \\
\text{str } [x] \ 1 \\
\text{str } [y] \ 1 \\
\text{unlock}(\ell) \\
\}
\quad \parallel \quad
\begin{array}{l}
\text{if trylock}(\ell) \{ \\
r_1 := \text{ldr } [x] \\
r_2 := \text{ldr } [y] \\
\text{unlock}(\ell) \\
\}
\end{array}
\end{array}$$

$\text{unlock}(\ell) \triangleq \text{str}_{\text{rel}} [\ell] \text{ unlocked}$

Fig. 30. Try-lock pseudocode

Fig. 31. Example of mutual exclusion with postcondition $r_1 = r_2$

uniqueness of a successful read-modify-write on a given write. The detailed implementation is in §6.4.4. The token and the token exchange happening at the exclusive store are hidden in the `HT-INS-STR-EXCL` rule. The users simply obtain the resource P tied to the event (②) when the exclusive store instruction succeeds (indicated by register value $v = 1$, in contrast to the Arm convention where $v = 0$ indicates success).

Using our proof rules, we prove the classical higher-order specification for a simple try-lock (see Fig. 30), which we then instantiate to prove a basic mutual exclusion example (in Fig. 31): if a writer takes the lock before writing to two variables, then a reader that takes the lock and reads from the two variables has to read the values before or after, but not a mixture.

5.8 Further Examples

To validate how AxSL^{Arm} works generally with the memory model of Arm-A, how it is likely to continue working with future changes, and how it could be ported to other memory models, we verify further examples that exercise different parts of the memory model. `LB+dmb+data` (Fig. 32) relies on the `po; [dmb.full]; po` clause of `bob` (see Fig. 4) to obtain the key `lob` edge on the left, but the proof is otherwise the same as for `LB+artificialdata+data` in §5.6. `LB+ctrls` (Fig. 33) relies on the `ctrl; [W]` clause of `dob` in both threads, but is otherwise the same. Similarly, `MP+rel+dmb+sy` (Fig. 34) relies on `bob` in the right thread, and so does `MP+rel+ctrl-isb` (Fig. 35) via the more complex `(ctrl|(addr; po)); [ISB]; po; [R]` edge which appears incrementally, but their proofs are otherwise as in §5.5. To illustrate that AxSL^{Arm} can reason about communication between many threads, we verify an iterated version of `MP`, namely `ISA2+rel+data+acq` (Fig. 36) [Sarkar et al. 2011]; the proof is just an iterated version of the proof of `MP`. Finally, to illustrate that reasoning about coherence is still possible, albeit unpleasant (which we discuss further in §8.3), we verify two coherence tests: `CoWW` (Fig. 37) and `CoRR` (Fig. 38); the proofs work by symbolic execution followed by discarding the executions with cycles in `co`.

$$\begin{array}{l} a: r_1 := \text{ldr } [x] \\ b: \text{dmb sy} \\ c: \text{str } [y] 1 \end{array} \parallel \begin{array}{l} d: r_2 := \text{ldr } [y] \\ e: \text{str } [x] r_2 \end{array}$$

Fig. 32. `LB+dmb+data`: $r_1 = 0 \wedge (r_2 = 0 \vee r_2 = 1)$

$$\begin{array}{l} a: r_1 := \text{ldr } [x] \\ b: \text{if } (r_1 == 0) \\ c: \text{str } [y] 1 \end{array} \parallel \begin{array}{l} d: r_2 := \text{ldr } [y] \\ e: \text{if } (r_2 == 1) \\ f: \text{str } [x] r_2 \end{array}$$

Fig. 33. `LB+ctrls`: $r_1 = 0 \wedge (r_2 = 0 \vee r_2 = 1)$

$$\begin{array}{l} a: \text{str } [data] 42 \\ b: \text{str}_{\text{rel}} [flag] 1 \end{array} \parallel \begin{array}{l} c: r_1 := \text{ldr } [flag] \\ d: \text{dmb sy} \\ e: r_2 := \text{ldr } [data] \end{array}$$

Fig. 34. `MP+rel+dmb+sy`: $r_1 = 1 \rightarrow r_2 = 42$

$$\begin{array}{l} a: \text{str } [data] 42 \\ b: \text{str}_{\text{rel}} [flag] 1 \end{array} \parallel \begin{array}{l} c: r_1 := \text{ldr } [flag] \\ d: \text{if } (r_1 == 1) \\ e: \text{isb} \\ f: r_2 := \text{ldr } [data] \end{array}$$

Fig. 35. `MP+rel+ctrl-isb`: $r_1 = 1 \rightarrow r_2 = 42$

$$\begin{array}{l} a: \text{str } [x] 42 \\ b: \text{str}_{\text{rel}} [y] 1 \end{array} \parallel \begin{array}{l} c: r_1 := \text{ldr } [y] \\ d: \text{str } [z] r_1 \end{array} \parallel \begin{array}{l} e: r_2 := \text{ldr } [z] \\ f: r_3 := \text{ldr } [x] \end{array}$$

Fig. 36. `ISA2+rel+data+acq`: $r_2 = 1 \rightarrow r_3 = 42$

$$\begin{array}{l} a: \text{str } [x] \ 37 \\ b: \text{str } [x] \ 42 \end{array}$$
Fig. 37. CoWW: $a \xrightarrow{\text{co}} b$

$$\begin{array}{l} a: \text{str } [x] \ 42 \quad \parallel \quad b: r_1 := \text{ldr } [x] \\ \quad \quad \quad \parallel \quad c: r_2 := \text{ldr } [x] \end{array}$$
Fig. 38. CoRR: $r_1 = 42 \rightarrow r_2 = 42$

6 Model and Soundness

A model of Hoare triples for a language of microinstructions is said to be sound if the following two conditions hold: (1) the model must allow us to show that proof rules are valid with respect to the language semantics (soundness), i.e. that the transformation of logical resources reflects the transition of physical states; (2) we must be able to show that, if a program is verified against a Hoare triple, then all valid executions of the program indeed satisfy certain properties (adequacy).

With these two principles in mind, we present the semantic models of the three logics presented in the previous section. We also sketch how to prove the soundness of proof rules using the models, and briefly touch on how the model definitions contribute to adequacy (we leave the details to §7).

Following the structure of the previous section, we build up the model of AxSL^{Arm} in Iris gradually with several intermediate steps, tackling one challenge at a time. We discuss variants of weakest preconditions, for our three logics at two abstraction levels. We will use superscripts SC, SCEst, and Arm to denote that assertions we are talking about belong to AxSL^{SC} , $\text{AxSL}^{\text{SCEst}}$, and AxSL^{Arm} respectively.

First, we recall the model of usual Iris logics built atop of a heap-based operational semantics, in a setting with SC concurrency and a fixed number of threads. We emphasise some core components of the model that are crucial for soundness and adequacy.

Next, we sketch the model of AxSL^{SC} based on our opax semantics for TinySc. We explain how it deals with the shape of opax semantics to achieve modular graph-based reasoning, and give an intuitive explanation on why models as such do not scale to the relaxed concurrency of Arm.

Then, we present the model of $\text{AxSL}^{\text{SCEst}}$ which is fundamentally distinct, due to our now explicit treatment of resource flowing. We concentrate on the flow implications, on how it works together with tied-to assertions to ensure soundness, and why this idea scales to relaxed concurrency.

Finally, we show the full semantic model of AxSL^{Arm} , which shares the same structure as that of $\text{AxSL}^{\text{SCEst}}$, but differs due to the shift of the synchronisation order from SC's sc to Arm's ob. We omit the soundness proof of the proof rules of AxSL^{Arm} as it mostly only concerns the structure of the semantic model thus similar to the one for $\text{AxSL}^{\text{SCEst}}$.

This section presupposes some basic knowledge of Iris. See [Jung et al. 2018] for background on Iris.

6.1 Preliminary: A General Recipe for Building Logics Using Iris

Iris can be instantiated by an operational semantics respecting a few conditions. The framework comes with a recipe for building logics and their adequacy proofs for typical heap-based operational semantics. We now recall this recipe, which we refer to as the *general recipe*. Readers who are familiar with this general recipe may skip this subsection.

Base operational semantics. To instantiate Iris, one normally needs to provide a concurrent language whose semantics has:

- a set of expressions Exp , and values $v : \text{Val} \subseteq \text{Exp}$
- a notion of physical state $\sigma : \text{St}$ shared among all threads
- a per-thread step relation $\rightarrow_{tid} \subseteq (\text{Exp} \times \text{St}) \times (\text{Exp} \times \text{St})$ that specifies how the program of thread tid may transform the shared state for a primitive step

For simplicity, we only consider a language with a fixed number of threads, i.e. without the ability to fork new threads.

Weakest precondition. Given such a semantics, one can follow the general recipe to first define a notion of weakest precondition. Here, we only show the key parts of the definition, and we refer the reader to [Jung et al. 2018] for a full definition that supports concurrency. Intuitively, $wp\ e\ \{Q\}_{tid}$ requires post condition Q to hold if e terminates with a value in thread tid . A Hoare triple $\{P\} e\ \{Q\}_{tid}$ is then defined as $\Box(P \multimap wp\ e\ \{Q\}_{tid})$. The definition of weakest precondition has two cases, depending on whether the program e is a value:

- If e is a value, the weakest precondition requires that the postcondition holds after a resource update: $\mathbb{H} Q(e)$ where \mathbb{H} is the update modality.
- Otherwise, for all physical updates from global state s , the logical resources are required to be updated to mirror them, using a *state interpretation* (SI) which is a predicate that gives the physical state a logical interpretation in Iris:

$$SI(s) \Rightarrow (\forall e', s'. (e, s) \rightarrow_{tid} (e', s') \Rightarrow SI(s') * wp\ e' \{Q\}_{tid})$$

Formally, it says that for any physical transition $(e, s) \rightarrow_{tid} (e', s')$ of thread tid , there must be a corresponding logical transition formulated as a view shift $SI(s) \Rightarrow SI(s')$. The recursive occurrence of the weakest precondition further requires this correspondence to hold for all subsequent steps of the thread. (Here we omit substantial technical details dealing with invariants and guarded recursion.)

State interpretation. Normally, to achieve local reasoning, one would define SI as an authoritative view of the physical state s , and only reason about fragmental views distributed to threads, separating s so that one does not need to reason about the parts of s that are not touched by a thread (known as framing). For instance, for an SC language with a global abstract heap (a map from locations to values) s , one can define the full view over s as $SI(s) \triangleq [\![\bullet s]\!]_1$ and a fragmental view for an individual location l as $l \mapsto v \triangleq [\![\circ\{l \mapsto v\}]\!]_1$ using the *Auth* ghost state constructor of Iris, such that (1) the two views are consistent: $SI(s) * l \mapsto v \vdash s(l) = v$, and (2) we can update them together: $SI(s) * l \mapsto v \Rightarrow SI(s[l \mapsto v']) * l \mapsto v'$.

Adequacy. Next, we show the statement of a typical adequacy theorem, and highlight the role of the state interpretation SI in the proof of the theorem. The theorem has the following notable assumptions (again, we omit substantial details):

- a thread pool step relation $\rightarrow_{tp} \subseteq (\text{list}(\text{Exp}) \times \text{St}) \times (\text{list}(\text{Exp}) \times \text{St})$ that specifies the scheduling of threads – all possible interleavings of primitive steps,
- a terminating thread pool trace $([e_0, \dots, e_n], s_i) \rightarrow_{tp}^* ([e'_0, \dots, e'_n], s_t)$ where s_i is the initial state, s_t is the terminating state, and all e' 's are values,
- a series of weakest preconditions with post conditions ϕ_0, \dots, ϕ_n , one for each thread

The conclusion is that the post conditions $\phi_i(e'_i)$ for all i hold. This theorem allows one to extract meta-logical results from the program logic, showing that verification done in the program logic is valid with respect to the meta logic.

Proof of adequacy. The proof of adequacy proceeds by allocating the initial logical interpretation $SI(s_i)$ to establish the physical-logical correspondence for s_i . We then continue by induction on the trace to show that the correspondence is preserved throughout the execution. In the induction case, when we have a head step $(e_i, s) \rightarrow_{tid} (e'_i, s')$ for thread tid and $SI(s)$, we unfold the definition of weakest precondition of thread i , which gives us $SI(s) \Rightarrow SI(s')$. Since we own $SI(s)$, we apply the view shift to obtain $SI(s')$. The case is concluded by applying the induction hypothesis which requires $SI(s')$ and the remaining reduction steps. In the end, we reach the final state s_t , at which point we get the post condition by the value case of the weakest precondition.

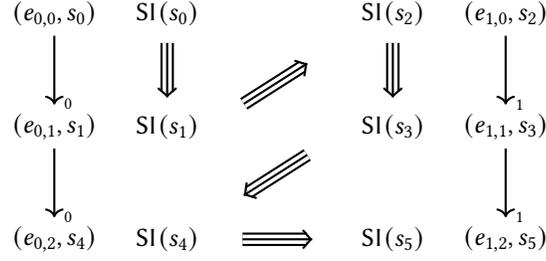


Fig. 39. An example program execution with two threads, above an SC memory model. \rightarrow_0 and \rightarrow_1 are the head reductions (the program order and the reasoning order) of the two threads. The indices of s indicate the global order of updating s (induced from \rightarrow_{tp}), and \Rightarrow indicates the order we update the logical interpretations of s in the adequacy proof. These two orders coincide, and they subsume the reasoning order.

Two orders. A crucial observation about this proof that will become relevant when we adapt the general recipe to work with an opax semantics is that the proof implicitly works with two orders: the order that the physical state evolves in, and the order in which the transformations of the corresponding logical interpretations are collected by weakest preconditions. The first order is induced by the thread pool reduction \rightarrow_{tp} of the semantics, which is a serialisation of all accesses to the physical state (from the perspective of the axiomatic model, a linear extension of sc). The second order is program order. From the model of the weakest precondition, we can see that a view shift is required for each local reduction \rightarrow_{tid} of the semantics. Fig. 39 illustrates the relationship between these two orders. More generally, we observe that, in most logics, the second order is program order, since it is intuitive to reason about programs along program order, whereas the first order may vary depending on the concurrency model.

Now, a crucial observation is the following: the reason we can complete the proof in one pass (with one induction), with a single induction on the former order, is that in the case of SC, the second order is included in the first: $po \subseteq sc$. As discussed in §3, in the case of Arm, these two orders are incompatible, in the sense that their union may be cyclic, and so they together do not form an inductive structure that we can rely on in the proof. This poses a challenge to the adequacy proof, and partially explains why the general recipe does not work for the relaxed concurrency of Arm.

6.2 Model and Soundness of $AxSL^{SC}$

$AxSL^{SC}$ has a simple semantic model and a one-pass adequacy proof similar to most Iris program logics. $AxSL^{SC}$ shares the exact core idea of updating logical resources respecting the transformation of physical states as in the general recipe. Following the recipe, we let the thread state of the opax semantics s be the expression; the complete state $\text{Done } \langle _ \rangle$ be the value; the pair of execution graph and instruction memory $\langle X, I \rangle$ be the physical state σ ; and \xrightarrow{tid}_h be the per-thread step. The notation $s \xrightarrow{tid, X, I}_h s'$ used in Fig. 11 is syntactic sugar for $\langle s, \langle X, I \rangle \rangle \xrightarrow{tid}_h \langle s', \langle X, I \rangle \rangle$. With this configuration, we define a *base weakest precondition* $\text{wpb}_{tid, \Phi}^{SC} s \{Q\}$ that takes an opax state s (the “expression”) and a post condition Q which is a predicate over the terminating state, in the spirit of the aforementioned general recipe, in Fig. 40.

A *clearer interface*. On top of wpb^{SC} , which we will explain soon, we define the weakest precondition $\text{wp}_{tid,\Phi}^{\text{SC}} p \{Q\}$ that merely takes a microinstruction program p , as a cleaner interface hiding s away. We implement this using usual Iris machinery:

$$\text{wp}_{tid,\Phi}^{\text{SC}} p \{Q\} \triangleq \forall T. \textcircled{1} \text{LSI}(T)_{tid} \Rightarrow \text{wpb}_{tid,\Phi}^{\text{SC}} \text{Ctd} \langle p, T \rangle \{T'. Q * \textcircled{2} \text{LSI}(T')_{tid}\}_{tid,\Phi}$$

We have a *local state interpretation* $\textcircled{1} \text{LSI}(T)_{tid}$ interpreting T for the current thread tid . LSI relates (local) logical assertions to the corresponding part of T . It allows us to universally quantify T and use the assertions to only track the parts that are necessary for further reduction from the opax state. The $\textcircled{2} \text{LSI}(T')_{tid}$ in the post condition of the wp^{SC} requires a new interpretation for the updated local state T' . Note that it is a thread-local predicate which only involves assertions of a thread (indexed by a thread ID) (for Iris experts, thread-local assertions use distinct ghost names).

$$\text{wpb}_{tid,\Phi}^{\text{SC}} s \{Q\} \triangleq \left(s = \text{Done } T \wedge \textcircled{1} Q(T) \right) \vee \left(s = \text{Ctd } C \wedge \left(\begin{array}{l} \forall \sigma. \textcircled{3} \text{Valid}(\sigma.X) * \textcircled{4} (\Box \text{SI}(\sigma)) * \forall s'. \textcircled{5} \langle s, \sigma \rangle \xrightarrow{tid}_h \langle s', \sigma \rangle * \\ \forall e = \langle tid, C.T.IT.cntr \rangle. (\text{IsValidEid}(e, X) * \forall \textcircled{6} s_{pg} \supseteq \text{Sc}(\sigma.X, e) \wedge e \notin s_{pg}. \\ \textcircled{7} \text{SIP}(\Phi, \sigma.X, s_{pg}) \Rightarrow \text{SIP}(\Phi, \sigma.X, s_{pg} \cup \{e\}) * \textcircled{8} \text{wpb}_{tid,\Phi}^{\text{SC}} s' \{Q\}) \\ \vee \textcircled{9} \neg \text{IsValidEid}(e, X) * \text{wpb}_{tid,\Phi}^{\text{SC}} s' \{Q\}) \end{array} \right) \right)$$

Fig. 40. The model of wpb^{SC} . Technical details regarding guarded recursion are omitted.

Overall structure. Following the general recipe, the definition of wpb^{SC} has two cases, depending on whether s is a “value”. In the value case, we just get the local state T and assert post condition $Q(T)$ after a ghost update. In the other case, there are two aspects of the definition which are somewhat non-standard compared to the general recipe: (1) it maintains consistency between the execution of the thread in the opax semantics and the global execution graph, to ensure sound graph reasoning; (2) it enforces rely-guarantee protocols on the graph to implement resource transfer.

Persistent graph as memory. Having an ongoing execution $\text{Ctd } C$, we obtain assertions over the “physical state” σ . However, unlike in the recipe when we just assume $\text{SI}(\sigma)$ for heap σ , now we have $\textcircled{4} \Box \text{SI}(\sigma)$. It is a *persistent* interpretation of σ , a pair of execution graph (the “shared memory”) and instruction memory, reflecting the fact that in opax both of them are constant. Additionally, $\textcircled{3} \text{Valid}(\sigma.X)$ assumes the validity condition of the execution graph, which helps us rule out ill-formed or inconsistent graphs – we discard a graph when we obtain information from the semantics that conflicts with this assumption (as demonstrated at the logic level in examples in §5). Next, as in the general recipe, we assume that we can make progress by taking a per-thread step ($\textcircled{5}$), which updates the state to s' . The weakest precondition should hold recursively for s' (as $\textcircled{8}$) – the weakest precondition predicate is defined as the (guarded) fixed point satisfying the recursive equation.⁴ We also need a case distinction using IsValidEid to check whether the event ID e corresponds to an event. In the case when we run out of microinstructions and have to reload (the microinstruction program of) the next instruction (namely e is invalid, as $\textcircled{9}$), we simply proceed with s' . This case distinction is just our ad-hoc way of handling the reloading step of opax, and there might be other more systematic solutions.

⁴Our definition does not require that s can take a step in that case; thus, our definition does not enforce progress.

Enforcing protocols. Finally, we have to show that for the current microinstruction (the first one in $C.p$), the associated event with ID e preserves the protocol Φ for the rely-guarantee style reasoning. Concretely, we use the *progress set* s_{pg} , a set of event IDs, to track how far we are from enforcing the protocol on all nodes of the guessed graph $\sigma.X$, and from checking the consistency between the guessed graph $\sigma.X$ and the program. It contains the set of the events that have been confirmed to conform the protocol and have corresponding microinstructions. We need to show that for an s_{pg} that contains (at least) all sc predecessors but not e , as assumed by ⑥, the protocol holds on e (so we make progress by adding e) given it holds on all events in s_{pg} (as ⑦), as illustrated in Fig. 41. The predicate $SIP(\Phi, X, s_{pg})$ is an interpretation of the progress set given a protocol Φ , which enforces the protocol on all write events in s_{pg} , as captured by:

$$SIP(\Phi, X, s_{pg}) \triangleq \bigstar_{e \in s_{pg}} \forall x, v. (X.\text{lab}(e) = W \ x \ v) \Rightarrow \Phi(x, v, e)$$

The view shift ⑦ means that we can *rely* on all sc-before events (i.e. those that are visible to the current event) conforming the protocol to *guarantee* that e also conforms the protocol. This view shift is crucial for proving the soundness of resource transfer happening in the proof rules. In particular, in the case of e being a read event, $s_{pg} \supseteq \text{Sc}(X, e)$ of ⑥ ensures that the protocol holds at all possible writes that it may read from (since they are all sc-before), which allows us to transfer the protocol resource along the rf from the actual write to this read. In the case of e being a write event, $SIP(\dots, s_{pg} \cup \{e\})$ requires the protocol resource of e .

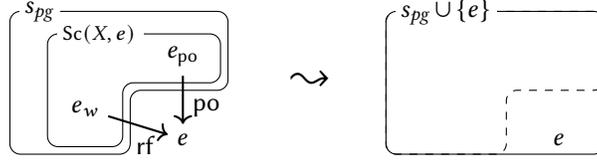


Fig. 41. An illustration on the evolving of the progress set s_{pg} at a read event e .

6.2.1 Soundness of Proof Rules. In the elaboration of the model above, we have intentionally left the precise definition of the predicates LSI and SI undefined, since their implementations are irrelevant to the model (and to the adequacy proof), and only pertain to the soundness of the proof rules. Indeed, to show soundness of the proof rules for AxSL^{SC} , it suffices that the predicates satisfy the selection of properties in Fig. 42. We refer readers to §C for a concrete definition of the predicates that satisfies all these properties.

Single-step weakest preconditions. Weakest preconditions specify the behaviours of a whole microinstruction program execution, while we need a mechanism to specify the behaviours of a single microinstruction to formulate the proof rules for it. Inspired by previous verification work for low-level languages [Erbesen et al. 2021; Jensen et al. 2013; Liu et al. 2023; Myreen and Gordon 2007], we use *single-step* base weakest precondition $\text{sswpb}_{tid, \Phi}^{\text{SC}} s \{s'. Q\}$ which means that Q holds after taking one reduction step (namely executing one microinstruction) from state s . We define sswpb^{SC} simply by replacing the recursive occurrence of wpb^{SC} with Q in the definition of wpb^{SC} . The single-step base weakest precondition formally corresponds (as expressed by the bi-directional entailment $\dashv\vdash$) to a single unfolding of the base weakest precondition, as captured by:

$$\begin{array}{l} \text{SSWPB-WPB} \\ \text{sswpb}_{tid, \Phi}^{\text{SC}} s \{s'. \text{wpb}_{tid, \Phi}^{\text{SC}} s' \{Q\}\} \dashv\vdash \text{wpb}_{tid, \Phi}^{\text{SC}} s \{Q\} \end{array}$$

Crucially, the right-to-left direction allows us to decompose a microinstruction program to only focus on one microinstruction at a time. This essentially plays the role of the bind rule in logics for high-level languages.

$$\begin{array}{c}
 \text{LSI-REG-AGREE} \\
 \text{LSI}(T)_{tid} * r \mapsto v \vdash T.\text{regs}(r) = v \\
 \\
 \text{LSI-PO-AGREE} \\
 \text{LSI}(T)_{tid} * \text{PoPred}(e) \vdash \langle tid, T.IT.cntr \rangle > e \\
 \\
 \text{SI-EDGE-AGREE} \\
 \frac{\text{Rel is a relation}}{\text{SI}(\sigma) * a \text{ Rel } b \vdash (a, b) \in \sigma.X.\text{Rel}}
 \end{array}
 \qquad
 \begin{array}{c}
 \text{LSI-REG-UPDATE} \\
 \frac{T'.\text{regs} = T.\text{regs}[r \mapsto v']}{\text{LSI}(T)_{tid} * r \mapsto v \Rightarrow \text{LSI}(T')_{tid} * r \mapsto v'} \\
 \\
 \text{LSI-PO-UPDATE} \\
 \frac{\langle tid, T'.IT.cntr \rangle > e'}{\text{LSI}(T)_{tid} * \text{PoPred}(e) \Rightarrow \text{LSI}(T')_{tid} * \text{PoPred}(e')} \\
 \\
 \text{SI-EDGE-ALLOC} \\
 \frac{\text{Rel is a relation} \quad (a, b) \in \sigma.X.\text{Rel}}{\text{SI}(\sigma) \vdash a \text{ Rel } b}
 \end{array}$$

Fig. 42. Selected properties of state interpretation and local state interpretation. Generally speaking, for mutable fields of T , we require agreement and update rules, while for the immutable graph X , we do not require an update rule.

Further, we define the single-step weakest precondition $\text{sswp}^{\text{SC}} p \{Q\}_{tid, \Phi}$, and finally define the microinstruction Hoare triple used by the proof rule for microinstruction i as

$$\{P\} i \{Q\}_{tid, \Phi} \triangleq \square \left(\forall p. (P * \exists K. (\text{Next } i \ K) = p) \multimap \text{sswp}_{tid, \Phi}^{\text{SC}} p \{Q\} \right)$$

where we require that the first microinstruction of p is i .

We use single-step microinstruction Hoare triples to specify proof rules, and then show the soundness result of AxSL^{SC} :

THEOREM 6.1. *The AxSL^{SC} proof rules for microinstructions are sound.*

PROOF SKETCH. We describe the general approach for a proof of soundness of a proof rule here. There are four major steps in the proof after unfolding the model of assertions. First, we use the agreement rules (Fig. 42) between the assertions in the precondition and the interpretation to partially recover the state s . Next, we take the opax step for the microinstruction, obtaining new facts about the execution graph and an updated state s' . At the same time, we perform resource transfer according to the protocol Φ . Finally, we allocate and update assertions to mirror the update of the local state (again using the LSI and SI rules in Fig. 42) and the resource transfer. \square

6.3 Model and Soundness of $\text{AxSL}^{\text{SCExt}}$

The $\text{AxSL}^{\text{SCExt}}$ logic extends the assertion language of AxSL^{SC} with tied-to assertions, which make it possible to reason about resource flowing between nodes. As a logic also built atop of an opax shape semantics, its model shares substantial similarities with that of AxSL^{SC} . The key difference is how it keeps track of all tied-to assertions to enable sound transfer of tied resources between events. As noted, explicitly tracking resource transfer between events does not add more expressive power to $\text{AxSL}^{\text{SCExt}}$, but makes it more general in the sense that both its model and adequacy result are robust for more (relaxed) concurrency models. This is reminiscent of how invariants are tracked in the weakest precondition for the standard Iris program logic [Jung et al. 2018]. We present the model in Fig. 43 and highlight how the model manages tied-to assertions below.

Interpretation for tied resources. SI_{\top} of $\textcircled{1}$ interprets the full authoritative view of all tied assertions, τ , which we keep as a logical map from event IDs to Iris propositions. Since τ mentions Iris propositions and itself is interpreted as an Iris proposition, we leverage Iris' support for higher-order ghost states to implement SI_{\top} . The

$$\begin{aligned}
 \text{wpb}_{tid, \Phi}^{\text{SCExt}} s \{Q\} &\triangleq (s = \text{Done } T \wedge \dot{\equiv} Q(T)) \vee \\
 &\left(s = \text{Ctd } C \wedge \right. \\
 &\quad \left(\forall \sigma. \text{Valid}(\sigma.X) * (\Box \text{SI}(\sigma)) * \forall s'. \langle s, \sigma \rangle \xrightarrow{tid}_h \langle s', \sigma \rangle * \right. \\
 &\quad \quad \left. \forall e = \langle tid, C.T.IT.cnt \rangle, \tau. (\text{IsValidEid}(e, X) * \textcircled{1} \text{SI}_{\tau}(\tau) \Rightarrow \right. \\
 &\quad \quad \quad \left. \exists \tau'. \textcircled{2} \text{SI}_{\tau'}(\tau') * \textcircled{3} \text{FlowImp}(\sigma.X, \Phi, e, \tau, \tau') * \text{wpb}_{tid, \Phi}^{\text{SCExt}} s' \{Q\} \right) \\
 &\quad \quad \left. \vee (\neg \text{IsValidEid}(e, X) * \text{wpb}_{tid, \Phi}^{\text{SCExt}} s' \{Q\}) \right) \\
 \text{FlowImp}(e, X, \Phi, \tau, \tau') &\triangleq \exists \tau_{in}, \tau_{res}, R. \textcircled{4} \text{Detach}(X, e, \tau, \tau_{in}, \tau_{res}) * \textcircled{5} \tau' = \tau_{res}[e \mapsto R] * \\
 &\quad \forall s_{pg} \supseteq \text{PredOf}(X.sc, e) \wedge e \notin s_{pg}. \\
 &\quad \textcircled{6} (*_{(e \mapsto R_{in}) \in \tau_{in}} R_{in}) * \text{SIP}(\Phi, X, s_{pg}) \Rightarrow \text{SIP}(\Phi, X, s_{pg} \cup \{e\}) * \textcircled{7} R
 \end{aligned}$$

Fig. 43. The model of base weakest precondition of $\text{AxSL}^{\text{SCExt}}$. The key changes to that of AxSL^{SC} are highlighted in orange. Again, the details handling guarded recursion are omitted.

predicate SI_{τ} is defined in such a way that, together with the fragmental tied-to-assertions, it enjoys agreement and update rules similar to those of the register map shown before. (For Iris experts we remark that the rules are slightly different since τ is higher-order: specifically, we can only obtain the agreement *later*.) The next line essentially allows us to update a fragment of τ (to τ' , as ②) and the associated tied assertions. Crucially, the update has to follow the flow implication ③ FlowImp , which we now explain.

Flow implication. The $\text{FlowImp}(X, \Phi, e, \tau, \tau')$ predicate regulates the flow and update of resources (from τ to τ') that may happen at memory event e . Intuitively, the rule expresses that the sum of resources pushed to e along its incoming sc edges implies (with a view shift) the resources given out along outgoing sc edges, plus the leftovers tied to e .

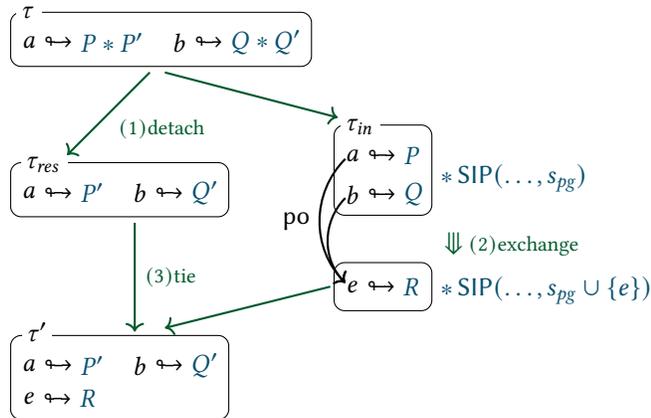


Fig. 44. A visualisation of the three actions of updating τ to τ' , where we flow P and Q to e along po from a and b respectively. R is eventually tied to e after exchange.

We divide the update from resource map τ to τ' into a sequence of three actions: detach, exchange, and tie, as illustrated in Fig. 44. The ④ Detach predicate captures the detachment:

$$\text{Detach}(X, e, \tau, \tau_{in}, \tau_{res}) \triangleq \text{dom}(\tau_{in}) \subseteq \text{PredOf}(X.\text{po}, e) * \forall(e \mapsto R_{in}) \in \tau_{in}. \\ \exists R_{res} = \tau_{res}(e), R = \tau(e). R \text{ -* } (R_{in} * R_{res})$$

τ is split into τ_{in} and τ_{res} , where τ_{in} is a fragment of τ which represents the resources *detached* from the po predecessors of e , as determined by the user-provided tied assertions, and τ_{res} is the remaining global map, after detaching τ_{in} .⁵ The last two lines do the resource *exchange*, which is an augmented version of the view shift presented in the model of AxSL^{SC}. It now has the local resources of τ_{in} explicitly given on the left as ⑥, and the resource ⑦ R remaining on event e on the right. The map update ⑤ does the tying: the remaining resource R is *tied* to e , by extending the map τ_{res} .

6.3.1 Stratification. The key difference between AxSL^{SC} and AxSL^{SCExt} is the tied-to assertions: $a \rightsquigarrow P$. The intent of $a \rightsquigarrow P$ is to express that P holds at event a , which is po-before the current event (if we ignore the case where P is transferred to another thread). In AxSL^{SCExt}, this then allows us to transfer P to the current event, or to any future event of the thread, and so $a \rightsquigarrow P$ amounts exactly to P . However, when we consider an actual relaxed memory model in AxSL^{Arm}, this transfer step will be conditional on ordering (in lob) from a to the current event, and the indirection induced by $a \rightsquigarrow P$ allows us to express that P is available under this condition.

The effect of tying P to a in AxSL^{SCExt} is to constrain the use of P : it isolates P from the reasoning along po that we do using $\text{wpb}^{\text{SCExt}}$. This stratification is reflected in the model $\text{wpb}^{\text{SCExt}}$: the update of tied resources as performed by the view shift in FlowImp cannot affect the recursive $\text{wpb}^{\text{SCExt}}$ for s' : we cannot use the updated resource to directly reason about s' . In Fig. 43, this is reflected by $\text{wpb}^{\text{SCExt}}$ being pulled out of the view shift. This is to contrast with AxSL^{SC} (see Fig. 40), where wpb^{SC} is on the right side of the view shift, and can therefore observe the update. This stratification is crucial for how we structure adequacy of AxSL^{Arm}, as we will show in the next section.

6.3.2 Soundness of Proof Rules. Like in AxSL^{SC}, we need to implement logical interpretation predicates and define a notion of single-step weakest precondition and Hoare triple. The soundness proofs for proof rules are then unsurprising, thanks to the substantial similarity between the models of the two wpbs, except that now more effort is needed to show that the update of τ to τ' satisfies the FlowImp predicate.

THEOREM 6.2. *The AxSL^{SCExt} proof rules for microinstructions are sound.*

6.4 Model and Soundness of AxSL^{Arm}

The model for AxSL^{SCExt} is almost parametric in the memory model, in the sense that it almost works directly for AxSL^{Arm}, a logic for TinyArm with relaxed Arm-A concurrency. We elaborate the two key changes we make to adapt the AxSL^{SCExt} model to Arm-A below. This demonstrates that the structure of the model can easily be adapted to another axiomatic memory model in which the synchronisation order (ob) and its local fragment (lob) are specified.

6.4.1 Hoare Triples. As in the other two logics, we define our Hoare triple of AxSL^{Arm} for a microinstruction program using its weakest precondition wp^{Arm} ; and define the weakest precondition using a base weakest precondition wpb^{Arm} for an ongoing local execution state.

The model of the base weakest precondition wpb^{Arm} has the same overall structure as in AxSL^{SCExt}, with a notable difference: because it targets the relaxed memory model of Arm-A, resources can only flow along the ob ordering. To enforce this, we make two changes in the definition of wpb^{Arm} , as depicted in Fig. 45:

⁵Since τ is higher-order, the official definition actually includes a later modality on the right of the separation implication in the definition of Detach, but we have omitted that from the presentation for simplicity.

- In FlowImp, we now require the quantified progress set s_{pg} to include only ob predecessors, so that one can only flow resources from nodes ordered with respect to the current event.
- In Detach, we require the user-provided τ_{in} to mention only lob predecessors of e , to further constrain the flow of local resources.

$$\begin{aligned}
\text{wpb}_{tid,\Phi}^{\text{Arm}} s \{Q\} &\triangleq (s = \text{Done } T \wedge \dot{\equiv} Q(T)) \vee \\
&\left(s = \text{Ctd } C \wedge \right. \\
&\quad \left. \left(\begin{array}{l} \forall \sigma. \text{Valid}(\sigma.X) * (\Box \text{SI}(\sigma)) * \forall s'. \langle s, \sigma \rangle \xrightarrow{tid}_h \langle s', \sigma \rangle -* \\ \forall e = \langle tid, C.T.IT.cnt \rangle, \tau. (\text{IsValidEid}(e, X) * \text{SI}_T(\tau) \Rightarrow \\ \exists \tau'. \text{SI}_T(\tau') * \text{FlowImp}(\sigma.X, \Phi, e, \tau, \tau') * \text{wpb}_{tid,\Phi}^{\text{Arm}} s' \{Q\}) \\ \vee (\neg \text{IsValidEid}(e, X) * \text{wpb}_{tid,\Phi}^{\text{Arm}} s' \{Q\}) \end{array} \right) \right) \\
\text{FlowImp}(e, X, \Phi, \tau, \tau') &\triangleq \exists \tau_{in}, \tau_{res}. R. \text{Detach}(X, e, \tau, \tau_{in}, \tau_{res}) * \tau' = \tau_{res}[e \mapsto R] * \\
&\quad \forall s_{pg} \supseteq \text{PredOf}(X.\text{ob}, e) \wedge e \notin s_{pg}. \\
&\quad (*_{(e \mapsto R_{in}) \in \tau_{in}} R_{in}) * \text{SIP}(\Phi, X, s_{pg}) \Rightarrow \text{SIP}(\Phi, X, s_{pg} \cup \{e\}) * R \\
\text{Detach}(X, e, \tau, \tau_{in}, \tau_{res}) &\triangleq \text{dom}(\tau_{in}) \subseteq \text{PredOf}(X.\text{lob}, e) * \forall (e \mapsto R_{in}) \in \tau_{in}. \\
&\quad \exists R_{res} = \tau_{res}(e), R = \tau(e). R -* (R_{in} * R_{res})
\end{aligned}$$

Fig. 45. The model of base weakest precondition of AxSL^{Arm} . The diff from $\text{AxSL}^{\text{SCExt}}$ (highlighted) reflects the shift of the synchronisation order from sc to ob.

6.4.2 *Soundness of Proof Rules.* With these minor modifications from $\text{AxSL}^{\text{SCExt}}$, this model works for AxSL^{Arm} , and we can prove soundness of the proof rules for microinstructions:

THEOREM 6.3. *The AxSL^{Arm} proof rules for microinstructions are sound.*

PROOF SKETCH. The proof is the same as the one for $\text{AxSL}^{\text{SCExt}}$, except that now we flow resources along lob/ob instead of po/sc (which is possible thanks to the stratification implemented by the model, as we remarked in §6.3), and we need to deal with the dependency edges of Arm-A. \square

6.4.3 *Supporting Framing and Invariants.* In the same way that one is used to splitting resources in separation logic, one would expect to be able to split $a \wp (P * Q)$ into $(a \wp P) * (a \wp Q)$. Recall that it is useful for proving examples (see Fig. 8). Modeling such splitting is, however, quite challenging due to its higher-order nature. (Interested readers may find our complex implementation using a combination of various Iris CMRAs in §C.2) We address this challenge by first implementing splitting with the help of the interpretation SI_T : $e \wp (P * Q) -* (\forall \tau. \text{SI}_T(\tau) \Rightarrow (\text{SI}_T(\tau) * e \wp P * e \wp Q))$. Here, the view shift allows us to update tied assertions and the interpretation without changing the value of the global map τ . Then, the view shift structure is hidden by adapting the definition of the weakest precondition to close it under this pattern. We define this view shift pattern as the interpretation modality $\dot{\equiv}^i$:

$$\dot{\equiv}^i P \triangleq (\forall \tau. \text{SI}_T(\tau) \Rightarrow (\text{SI}_T(\tau) * P))$$

Supporting invariants, on the other hand, only requires minor updates to the weakest precondition definition: we just replace the plain view shift in FlowImp with a more expressive view shift (a combination of the ‘later’ modality and the ‘fancy update’ modality) that allows opening and closing of invariants, similar to [Jung et al. 2018]. Importantly, these invariants do not enable resource transfer (which would be unsound), as they did in CSL

for non-relaxed concurrency. Instead, we mainly use them to construct persistent wrappers for non-persistent resources that we want to transfer via AxSL^{Arm} protocols, like escrows in GPS [Kaiser et al. 2017; Turon et al. 2014].

6.4.4 Supporting Exclusives. We use the escrow pattern to support transferring non-duplicable resources at a successful read-modify-write of a given write, as outlined in §5.7.

Given a location x on which to use read-modify-writes to transfer an exclusive resource P , we put P into the following Iris invariant RMWInv, as part of the protocol of x :

$$\text{RMWInv}(e_w, P) \triangleq \boxed{P \vee (\exists e_r, e'_w. e_w ((\text{rf}; [e_r]; \text{rmw}) \& \text{co}) e'_w * \text{ExTok}(e'_w))}$$

where the notation \boxed{I} is an invariant containing I . The invariant asserts a disjunction: either P is released by the store e_w to this location (the left disjunct), or P has already been exchanged for an exclusive token $\text{ExTok}(e'_w)$ by some successful store exclusive e'_w that is paired with the load exclusive e_r reading from e_w (right disjunct). The exclusive token satisfies rule **EXTOK-EXCL** to capture that e'_w is unique and the exchange occurs once.

$$\begin{array}{c} \text{SIT-EXTOK-ALLOC} \\ \frac{e \notin \text{dom}(\tau) \quad \tau' = \tau[e \mapsto _]}{\text{Sl}_\tau(\tau) \Rightarrow \text{Sl}_\tau(\tau') * \text{ExTok}(e)} \end{array} \qquad \begin{array}{c} \text{EXTOK-EXCL} \\ \text{ExTok}(e) * \text{ExTok}(e) \vdash \perp \end{array}$$

While multiple load exclusives may read from e_w and obtain the invariant, none can exploit it on its own. Only a paired successful store exclusive can make use of it: a successful store exclusive e''_w that pairs with a po-earlier load exclusive e_r is guaranteed to be the unique such store exclusive across all loads reading from e_w . This unicity implies $e''_w = e'_w$ by graph reasoning, which allows us to open the invariant (which it can obtain from e along $\text{rmw} \subseteq \text{ob}$) and refute the right disjunct: The exclusive token $\text{ExTok}(e''_w)$ allocated for the successful store exclusive contradicts the same token from the protocol indicating that this exchange has already occurred by **EXTOK-EXCL**. Therefore, a successful store exclusive can extract P from the left disjunct when opening the invariant and close it with the right disjunct by depositing its exclusive token.

We now describe the needed minor adaption to the language semantics and the logic. (1) We add an extra bookkeeping field srcs_{rmw} (of type $\text{option}(\text{Eid})$) to the thread state T of the opax. This new field is to remember the candidate load exclusive e that the next store exclusive may pair with, which we track with a new bookkeeping assertion $\text{RmwPred}(e)$, similar to how we track $\text{srcs}_{\text{ctrl}}$ with CtrlPreds . (2) We extend Sl_τ with a new interpretation for the domain of the global tied-to map τ , such that it allows us to allocate exclusive tokens for successful store exclusives with the **SIT-EXTOK-ALLOC** rule.

6.4.5 Pulling Out Tied Resources. In some situations, for example when considering the postcondition of a whole program, knowing exactly which event which resource comes from is not helpful. Instead, one can use a ‘normal’ postcondition by pulling out the tied resources from tied-to assertions (which means that we cannot reuse the specification in the proof of a larger program anymore). Assertions can only be pulled out after the completion of the verification, for a closed program. Stripping the tied-to from P otherwise, namely in the middle of a program, is unsound, since such an operation would allow flowing resources from ob back to later po (and thus may lead to circular reasoning). Implementing pulling out only requires a minor change to the semantics of wpb^{Arm} : we replace the $\dot{\Rightarrow} Q$ in the case handling termination with $\text{PullOutTied}(\text{tid}, Q)$:

$$\text{PullOutTied}(\text{tid}, Q) \approx \forall \tau. \text{Sl}_\tau(\tau) * (\text{Sl}_\tau(\tau) * (*_{\{R \mid (e \mapsto R) \in \tau \wedge \text{tid} = e.\text{tid}\}} R \Rightarrow Q))$$

PullOutTied requires us to establish the postcondition Q , assuming that the predicates pulled out from local tied assertions hold. Technically, the definition of $\text{PullOutTied}(\text{tid}, Q)$ makes use of the agreement rule for a local event e , which roughly says that $\text{Sl}_\tau(\tau) * e \mapsto R$ implies that $e \mapsto R$ is in τ and thus that R holds (this is an approximate description, the formal details are a bit more subtle because of the higher-order nature of the τ map

mentioned above) — we have already seen an example of how this is used, namely in the final reasoning step (in each thread) in Fig. 26.

6.4.6 Stuckness and Infinite Executions. As described in §4.3.5, we assume non-stuckness in the model of our weakest preconditions. We do not need to show that the program does not get stuck, since we only consider terminated opax traces, as we will see in the adequacy statement in the next section.

Besides, because of the open problem with infinite executions in the memory model (discussed in §4.3.6), our definitions of weakest preconditions does not take measures to handle infinite executions either.

7 Adequacy

The adequacy of Iris logics is usually expressed as a meta-level theorem. Generally speaking, this theorem about a program logic (in our case, Iris) extracts, from a program specifications proven in the logic, a result in the meta logic in which the program logic is implemented (in our case, this meta logic is Rocq). This theorem shows that the program logic is sound, in the sense that the properties of programs proved in the logic hold in the meta logic.

The crux of the adequacy proof is to compose thread-local reasoning results, specified as weakest preconditions. For classic concurrent separation logics (including most Iris-based program logics using the standard weakest precondition construction, including iGPS), the proof of adequacy works by induction on the program execution trace, as described in the general recipe in §6.1. The adequacy proofs of AxSL^{Arm} differs from — and, we argue, generalises — the general recipe (basically, the general recipe is our approach, instantiated with the same order twice, see §7.2.4 for an example) in two respects: our novel opax semantics, and the semantics model.

In this section, we first explain how to work with an opax semantics by showing the adequacy of AxSL^{SC} . In this adequacy proof, we need to handle the opax shape, and the rely-guarantee style resource transfer. Next, we show the adequacy proof of AxSL^{Arm} ; we skip $\text{AxSL}^{\text{SCExt}}$, since it has an almost identical semantic model, and thus proof, to AxSL^{Arm} . This proof significantly differs from the previous general recipe, due to the new semantic model that enforces stratification (as described in §6.3.1). Because the stratification isolates the reasoning order (po) and the resource flowing order (ob), we have to do two separate inductions on these two separate orders.

7.1 Adequacy of AxSL^{SC}

The statement of adequacy of AxSL^{SC} is as follows:

THEOREM 7.1 (ADEQUACY OF AxSL^{SC}). *For any initial thread states \vec{C} (each is a pair $\langle p, T \rangle$), meta-level propositions \vec{P} (one for each thread), and consistent, well-formed execution graph X , we have*

$$\begin{aligned} & \left(\textcircled{1} \bigwedge_{tid=1}^n \text{Ctd } \vec{C}(tid) \xrightarrow{tid, X, I}^*_{\text{h}} \text{Done } \langle _ \rangle \right) \Rightarrow \\ & \textcircled{2} \exists \Phi. \vdash \left(\textcircled{3} \text{InitRes}(\Phi) * \left(\textcircled{4} \square \text{SI}(\langle X, I \rangle) \right) * \right. \\ & \quad \left. *_{tid=1}^n \left(\textcircled{5} \text{LSI}(\vec{C}(tid).T)_{tid} * \textcircled{6} \text{wp}_{tid, \Phi}^S \vec{C}(tid).p \left\{ \left[\vec{P}(tid) \right] \right\} \right) \right) \Rightarrow \\ & \left(\textcircled{7} \bigwedge_{tid=1}^n \vec{P}(tid) \right) \end{aligned}$$

Here, \vec{T} is a sequence of initial thread states (one for each thread). The first hypothesis (Ⓚ) ensures that the memory graph X reflects the behaviours of a complete program by assuming terminating executions of all threads starting from $\text{Ctd } \vec{C}$. The second line assumes that we have proofs in AxSL^{SC} of weakest preconditions (as Ⓞ), one for each thread, using the same protocol Φ . This is where we require that the *same* protocol Φ is agreed upon between the thread-local proofs of the weakest preconditions for each thread (as Ⓜ), as well as agreement about the execution graph X and instruction memory I via the state interpretation (Ⓨ). In addition, we also require the allocation of initial protocol resources ⓐ (for all initial writes). The weakest preconditions Ⓞ

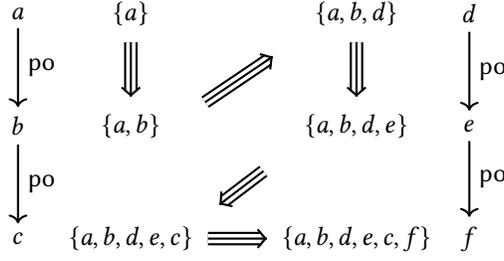


Fig. 46. An example program execution graph with two threads, above an SC memory model. Assuming sc edges from b to d , from e to c , and from c to f , \Rightarrow indicates the order we update $\text{SIP}(\Phi, X, s_{pg})$, the interpretation of the progress set in the adequacy proof, where the sets of nodes indicates the advance of the progress omitting initial nodes.

are for microinstructions, which takes the microinstruction program $\vec{C}(tid).p$. The post conditions are assumed to be the meta-level propositions \vec{P} , lifted to AxSL^{SC} by $[_]$. (This lifting is similar to what happens in other Iris-based program logics: it simply embeds a proposition P from the meta-level as the corresponding proposition in AxSL^{SC} .) Next to the weakest precondition of the thread, we require the local state interpretation ⑤ for the initial thread state $\vec{C}(tid).T$. From these, adequacy tells us that ⑦ \vec{P} hold in the meta-logic as well.

The value of the adequacy theorem is that it means that we do not need to trust or even understand the intricacies of AxSL^{SC} : once we have proved weakest preconditions for each thread using the AxSL^{SC} proof rules, then the adequacy theorem guarantees that the postconditions \vec{P} do indeed hold at the meta-level (assuming each thread's execution terminates).

PROOF SKETCH. The overall proof strategy is that we first show $\models \left[\bigwedge_{tid=1}^n \vec{P}(tid) \right]$, that is the goal lifted to AxSL^{SC} , and then show that all lifted meta level propositions (under certain Iris modalities) hold in the meta logic. We only look at the first part, as the latter is exactly the soundness result of the Iris base logic.

The proof starts by allocating $\text{SIP}(\Phi, X, s_{pg})$, where s_{pg} is the set of all initial (write) events, which means that we have to establish the protocol resources on all initial nodes. This is derivable from $\text{InitRes}(\Phi)$ which essentially states the same thing. Next, we can unfold the model of every wp. Given the local interpretation LSI, we obtain one wpb from one wp. We then do induction on the sc order. The proof then proceeds in a way that closely follows the adequacy proof of the general recipe, as illustrated in Fig. 46. That is, we follow the sc order, which is the order that the program executes, to update logical resources using the view shift which we obtain by unfolding the model of wpb. One major difference is that, unlike SI in the general recipe, the resource we update now is SIP. Respecting the sc order, we collect events into s_{pg} via the update, and at the same time ensure that the protocol Φ is maintained by all nodes. \square

Crucially, the proof would not go through if we did induction on the program execution trace, which is what the general recipe does. This is because traces of the opax semantics do not contain interleaving (they are merely thread-local traces), but the resource transformation depends on interleaving. We therefore have to do induction on a structure that contains the interleaving information, which is the sc relation of the execution graph. Technically, the $s_{pg} \supseteq \text{Pred}(X.sc, e)$ condition in the model of wpb enforces that one can only perform the resource update for e (making a progress) after all sc predecessors of e have been handled.

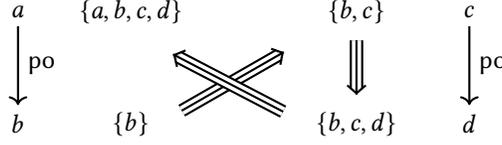


Fig. 47. An example program execution graph with two threads, above Arm-A memory model. Assuming ob edges from b to c , from c to d , and from d to a , \Rightarrow indicates the order we update $\text{SIP}(\Phi, X, s_{pg})$ in the second phase of the adequacy proof, where the sets of nodes indicate the advance of s_{pg} .

7.2 Adequacy of $\text{AxSL}^{\text{SCExt}}$ and AxSL^{Arm}

We cover the adequacy of $\text{AxSL}^{\text{SCExt}}$ and AxSL^{Arm} together, as the similarities between their semantic models mean that their adequacy proofs also proceed similarly. We focus on the adequacy proof of AxSL^{Arm} , and only describe its difference to the one for $\text{AxSL}^{\text{SCExt}}$ at the end of this subsection.

The *statement* of adequacy for AxSL^{Arm} is identical to Theorem 7.1, except for that we now additionally require an SI_\top (highlighted) for an empty tied-to map in the second hypothesis, meaning that, at the beginning, no resources are associated with any events. However, the *proof* of adequacy significantly diverges from that of Theorem 7.1, because of the stratification enforced by the semantic model of AxSL^{Arm} .

THEOREM 7.2 (ADEQUACY OF AxSL^{ARM}). *For any initial thread states \vec{C} , meta-level propositions \vec{P} (one for each thread), and consistent, well-formed execution graph X , we have*

$$\begin{aligned} & \left(\bigwedge_{tid=1}^n \text{Ctd } \vec{C}(tid) \xrightarrow{tid, X, I}_h^* \text{Done } \langle _ \rangle \right) \Rightarrow \\ & \exists \Phi. \vdash \left(\text{InitRes}(\Phi) * \text{SI}_\top(\emptyset) * (\Box \text{SI}(\langle X, I \rangle)) * \right. \\ & \quad \left. *_{tid=1}^n \left(\text{LSI}(\vec{C}(tid).T) * \text{wp}_{tid, \Phi}^\wedge \text{Ctd } \vec{C}(tid).p \left\{ \left[\vec{P}(tid) \right] \right\} \right) \right) \Rightarrow \\ & \quad \left(\bigwedge_{tid=1}^n \vec{P}(tid) \right) \end{aligned}$$

7.2.1 Overview of the Proof. Following the stratification of the semantics model, our novel adequacy proof is also stratified into two phases. We now outline the two phases informally.

Phase one, informally. First, like in some previous logics based on axiomatic models, including RSL [Vafeiadis and Narayan 2013] and GPS [Turon et al. 2014], we construct an annotated execution graph. This follows directly from the fact that the opax semantics is essentially an operational wrapper on top of an axiomatic model. Then we construct an annotation of the execution graph using flow implications from our weakest precondition. At each reduction step of the semantics, the weakest precondition remembers a flow implication for each node in the graph, and ensures that all these flow implications can be chained.

Thus, we can collect the needed flow implications at all nodes by unfolding the weakest preconditions of all threads along program order, and then connect them together to obtain a full annotation. During the annotation construction, our protocol plays a crucial role, since it specifies how resources flow across threads and is agreed upon between them, guaranteeing that the flow implications of different threads are compatible.

Phase two, informally. Second, to get an Iris-style adequacy statement in which we show all pure postconditions hold in the meta logic, we need to actually perform all the resource transformations (namely the flow implications) of the annotated graph, starting from the initial resources. This step of our proof is novel, since usually the resource transformations are simply performed *on the fly* in stronger settings, thanks to the acyclicity of $\text{po} \cup \text{rfe}$ in these memory models. In those settings, it is possible to update resources between program points according to

the flow implications (i.e. apply them to the resources) in $po \cup rfe$ during the unfolding, since this order includes the po order in which weakest preconditions collect the flow implications.

On the other hand, when working with a relaxed model like that of Arm-A, the order that we rely on is ob , which does not include po , and our tied-to assertions are employed to restrict po resource flow in $AxSL^{Arm}$. However, we still want syntax-oriented weakest preconditions which collect flow implications along po , as illustrated in Fig. 47. This tension poses the main challenge to the proof of adequacy, since one cannot do induction on the potentially cyclic $po \cup ob$. We resolve this tension by *stratifying* the usual proof procedure described in the general recipe into two phases. We argue that this stratification, and the semantic model of $AxSL^{Arm}$ that enables it, are a generalisation of the usual Iris-based CSLs and their one-pass adequacy proofs.

In the rest of the section, we explain our novel two-phase adequacy proof in more detail, and leave the discussion on the relation to other adequacy proofs to related work (§9).

7.2.2 Phase One. The goal of phase one is to show the following lemma, from which we show the final adequacy statement in phase two.

LEMMA 7.3. *For a valid execution graph X ,*

$$\begin{aligned} & \exists \epsilon, \tau. \textcircled{1} \text{Sl}_T(\tau) * \textcircled{2} \text{dom}(\tau) = \text{dom}(\epsilon) = \text{AllNodes}(X) * \\ & \quad \textcircled{3} \forall e \mapsto m \in \epsilon. \text{dom}(m) \subseteq \text{PredOf}(X.\text{lob}, e) * \\ & \quad \forall e \mapsto R \in \tau. \\ & \quad \forall s_{pg} \supseteq \text{PredOf}(X.\text{ob}, e) \wedge e \notin s_{pg}. (\text{SIP}(\Phi, X, s_{pg}) * \textcircled{4} *_{\rightarrow R_i \in \epsilon[e]} R_i) \Rightarrow \\ & \quad (\text{SIP}(\Phi, X, s_{pg} \cup \{e\}) * \textcircled{5} *_{e_o \in \text{SuccOf}(X.\text{lob}, e)} \epsilon[e_o][e] * R) \end{aligned}$$

Generally speaking, the lemma requires a well-formed annotation on every lob edge of the execution graph. The edge annotation ϵ (of type $Eid \rightarrow Eid \rightarrow iProp$) records the history of how resources evolve and are transferred soundly along lob in a complete program execution (e.g. $\epsilon[e][e']$ is the annotation of e' lob e), as inspired by RSL/FSL. The first line of the lemma also requires a tied-to map τ , and its logical interpretation $\textcircled{1} \text{Sl}_T(\tau)$. This map is the final one after checking all events, which records the resources that remain on the events after flowing all the resources. To enforce that this map is indeed final, $\textcircled{2}$ requires that the domain of τ and ϵ is all the nodes of the execution graph X . In the next line, we impose a well-formedness condition on ϵ : $\textcircled{3}$ says that, for all mappings of e to m in ϵ , m is a node-to-resource map specifying the resources that flow to e from its lob predecessors. The last three lines relate τ and ϵ : for all dangling resources R on node e in τ , a variant of the view shift part of FlowImp holds for the protocol Φ , where $\textcircled{4}$ is the local resources flowing into e , and $\textcircled{5}$ is the resources flowing out from e , both along lob .

Thanks to our definition of weakest precondition, the local edge annotations for an event, which is essentially a resource transformer, can be easily derived from the corresponding tied-to map update specified by the flow implication: we take τ_{in} as $\epsilon[e]$ for every node e . Furthermore, the fact that every such update follows the flow implication guarantees that these local annotations can be composed both vertically (in po) and horizontally (between threads) to get a global edge annotation, which is captured as the lemma above.

PROOF SKETCH. In each thread, vertical composition is done by induction on the thread's trace to unfold the recursively defined weakest precondition. This allows us to collect local annotations along po to obtain an annotation for the thread. Next, the derived annotations are glued horizontally (between threads, by taking the union of the local ones) to obtain a global annotation ϵ , which is possible since all resource exchange across threads (as annotated on obs) are specified by the same rely-guarantee protocol Φ . \square

7.2.3 Phase Two. Given the edge annotation ϵ from phase one, phase two stitches the flow implications together by induction on ob , as illustrated in Fig. 47.

PROOF SKETCH. For the base case of the induction, we require the user to show that the resources specified by Φ for the initial events of all memory locations can be established. The allocated resources are then used as the starting point for the application of the flow implications along (an order extension of) ob . Finally, we combine the remaining tied-to assertions at the end of the process to show the postconditions. \square

7.2.4 *Adequacy of $AxSL^{SCExt}$* . The adequacy proof of $AxSL^{SCExt}$ differs from that of $AxSL^{Arm}$ only in phase two: we just need to make the induction of phase two follow sc instead of ob . This proof is effectively the proof of adequacy of $AxSL^{SC}$, but divided into two phases due to the stratification imposed by the model. Recall that in the $AxSL^{SC}$ proof, we do induction on the sc relation to collect flow implications and apply them on-the-fly. Here in this proof, we collect the flow implications in phase one along po and apply them in the sc order in phase two.

8 Technical Remarks and Limitations

8.1 Technical Improvements to the Original $AxSL^{Arm}$

Although the syntax of $AxSL^{Arm}$ in §5 is almost identical to the original $AxSL^{Arm}$ [Hammond et al. 2024], the model presented in §6 is new, and factored cleanly around well-defined abstractions. This has two benefits: first, it makes the proof of adequacy significantly simpler, and second, it allows us to apply the same recipe to different memory models, which we illustrate using SC in $AxSL^{SCExt}$.

Concretely, in the old model, we enforce the protocol on exactly the obs predecessors of the current node, which causes two major inconveniences. First, in the proof of adequacy, we have to unfold the (old) $FlowImp$ predicate, and reason about the protocol resources tied to the obs predecessors of a node explicitly. For instance, we have to reason about whether an obs predecessor of a node e is also an obs predecessor of another node e' (e.g., when two loads read from the same write) in the phase two of the adequacy proof. This adds extra complexity to the proof.

Second, users of the logic can only send away persistent resources with the (old) proof rule for stores. This is because we have to require the protocol resources to be persistent, ensuring that they remained unchanged after applying the flow implication (so they can be transferred to other nodes) in the adequacy proof.

With the improved model of this paper, we use the progress set s_{pg} to track the set of events on which we have ensured that the protocol holds, and abstract the enforcement of the protocol on s_{pg} with SIP . This abstraction avoids unfolding the definition of SIP and reduces explicit resources reasoning.

8.2 Proof Effort for $AxSL^{Arm}$

Much of the effort was in the overall design of the logic to overcome the challenge to soundness posed by load buffering. Shaping the idea to fit Iris, and detailing the model of assertions and the definition of weakest preconditions, took over a person-year, but the result has been robust to small changes, for example to add exclusives. The adequacy theorem has the most significant proof: it took two or three person-months to mechanise the original proof of the POPL 2024 paper after initial design work. Afterwards, the proof was simplified thanks to the improvement to the model described in §8.1, which took around a week. Writing and proving an instruction proof rule takes about half a person-day now that we have enough of them to flesh out a pattern. Finding an overall proof structure for a new shape of litmus test takes a few person-days; in fact, it is very similar to what is needed for example in RSL. Adapting a proof from one variant of a litmus test to another is just a few hours' work, and less than a hundred lines of Rocq.

Overall, the mechanisation for $AxSL^{Arm}$ is divided as follows:

Item	LoC
Prelude (incl. outcome interface and infrastructure)	~4800
Language definition and lemmas	~1100
Axiomatic model and lemmas	~3000
Iris CMRAs	~900
WPs and assertions	~3700
Proof rules and their soundness proofs	~2700
Adequacy	~1100
Examples	~3300
Total	~20000

The purpose of this table is to give readers an impression on the scale of the mechanisation. The numbers in this tabular are not intended for direct numerical comparison with the numbers reported in the POPL 2024 paper [Hammond et al. 2024], since the mechanisation has experienced refactoring and is therefore more flexible than before.

8.3 Mixed ordering reasoning

The memory model of Arm-A involves two main axioms: *external*, which requires *ordered-before* to be acyclic, and *internal*, which requires *po-loc | ca | rf* to be acyclic, which effectively enforces per-location sequential consistency (the atomicity axiom is much more ‘local’ and easier to use, as per §5.7). This latter order is sometimes also called *coherence*, or *extended coherence* to avoid confusion with the coherence *relation*, *co*, which is merely part of it.

The way AxSL^{Arm} is defined in Iris above the opax semantics using the memory model of Fig. 4 means that it captures both axioms. However, the design of the logic focuses on the external axiom, and leverages the acyclicity of *ob* to allow sound transfer of resources along *ob*. This means that AxSL^{Arm} cannot soundly allow transfer resources along the potentially cyclic combination of *ob* and extended coherence, as the phase two induction proof of the adequacy needs an acyclic order. It is always possible to reason about extended coherence by brute force in AxSL^{Arm} , by explicit graph reasoning (as in §5.5), but this is unsatisfactory. This is a definite limitation of AxSL^{Arm} , as many common programming and reasoning idioms rely on reasoning about the combination of ordered-before with extended coherence, in particular full GPS protocols and the notion of non-atomics.

9 Related Work

There is extensive work on verification of relaxed memory models. Here we mainly discuss the line of work on separation logics starting from RSL, and only mention some other work that is closely related.

Overview. RSL [Vafeiadis and Narayan 2013], GPS [Turon et al. 2014]/GPS+ [He et al. 2016], and FSL [Doko and Vafeiadis 2016]/FSL++ [Doko 2021; Doko and Vafeiadis 2017] are defined with respect to an axiomatic memory model, namely that of C11/RC11, and their proofs of soundness are built from the ground up using non-standard models of separation logic, which (as described by Kaiser et al. [Kaiser et al. 2017, §1.2]) requires significant effort. Later logics like iGPS [Kaiser et al. 2017], Cosmos [Mével et al. 2020], and iRC11 [Dang et al. 2020], rely on (re)formulating the target relaxed memory model as an operational model to obtain an Iris-based logic with advanced features like higher-order ghost states ‘for free’.

Very relaxed hardware memory models. The proofs of adequacy of RSL and FSL are somewhat similar to ours, also being based on chaining flow implications along a global acyclic synchronisation relation, C11’s *happens-before* (*hb*). However, the memory model of C11 is substantially different from that of Arm, and in particular, despite a similar role, *hb* is substantially different from *ob*: it is defined as $\text{hb} \triangleq (\text{sb} \cup \text{sw})^+$, where *sb* is C11’s

counterpart to po , and sw is C11’s loose counterpart to obs . This allows them to freely persist resources along program order, and so their flow implications refer to immediate program order successors and predecessors of instructions, which means that they have their postcondition immediately in hand, and do not need to collect tied resources. It also means that, unlike ours, their proof of adequacy can be along program order.

FSL [Doko and Vafeiadis 2016] extends RSL to make it possible to transfer resources using C11 ‘relaxed’ access that are suitably fenced by guarding resources with modalities: a relaxed load, which imposes little order by itself, merely obtains ∇P , which is not usable by itself, but which an acquire fence turns into P . Symmetrically, P itself cannot be sent away by a mere relaxed store, but a release fence turns P into ΔP , which a relaxed store can send away. Our $a \rightsquigarrow P$ assertion can be seen as an indexed version of ∇P , keeping track of the source of the assertion in a way that is compatible with ghost state, even with cycles in $po \cup rf$.

FSL++ [Doko 2021; Doko and Vafeiadis 2017] extends FSL with a form of ghost state (albeit one not as expressive as that of Iris), but this makes it unsound for memory models that exhibit load buffering, and so FSL++ targets RC11 [Lahav et al. 2017], a significant strengthening of C11 that requires that $po \cup rf$ is acyclic.

In a sense, FSL and FSL++ both allow reasoning along po , but put some guards to limit transfer of physical resources. For FSL, which has no other resources, this limits its expressivity but poses no soundness problem. For FSL++, this imposes strong requirements on the underlying memory model.

Our flow implications are inspired by those of RSL and FSL. However, thanks to the phrasing of the memory model of Arm-A, we give a single, generic definition, instead of one based on case analysis of instructions. In addition, the pervasive effect of undefined behaviour (stemming from data races on non-atomics and uninitialised reads) in C11 substantially complicates the definition of flow implications of RSL and FSL.

A substantial fragment of RSL, FSL, and FSL++ is encoded into the Viper deductive verification tool [Summers and Müller 2018] for automated verification.

Very relaxed programming language models. SLR [Svendsen et al. 2018] targets the Promising Semantics [Kang et al. 2017] designed to fix the out-of-thin-air problem of C11. SLR takes advantage of the extra strength to enable coherence (sc-per-location) reasoning on relaxed accesses, but does not allow resource transfer using relaxed accesses. SLR features an assertion to keep track of writes that is somewhat similar to our NoLocalWrites and LastLocalWrite assertions: $W^\pi(x, X)$ imposes a lower bound X on the set of writes done so far to location x ; however, they use it for coherence reasoning, rather than to bound internal reads.

Dalvandi et al. [2020, 2022] developed a mechanised Owicki-Gries-style logic for a subset of C11 [Doherty et al. 2019], which was later adapted by Wright et al. [2021, 2023] to the modular semantic dependency model of Paviotti et al. [2020] — another proposed solution to the out-of-thin-air problem. Following this line of work, Piccolo [Lahav et al. 2023] is a rely/guarantee logic that claims to provide a uniformed account for these logics. The logic of Wright et al. [2021, 2023] supports sound reasoning in the presence of $po \cup rf$ cycles using an approach that differs from AxSL in several key aspects. First, their logic is not syntax-directed, meaning that it does not support reasoning about programs by following their syntactic structure along program order, which is one of the key benefits of program logics, in particular for interactive proving. Second, inheriting from Owicki-Gries, their logic lacks thread-modularity — a crucial property that enables scalable compositional verification. Third, the general form of auxiliary variables in Owicki-Gries logics is unsound in the setting of relaxed memory [Lahav and Vafeiadis 2015], and so they only feature a restricted form, whereas AxSL allows arbitrary yet sound use of higher-order ghost state, as proven by our adequacy theorem. On the other hand, their simple notion of auxiliary state lends itself well to automation, as they demonstrate. Fourth, since C11 treats non-atomic accesses specially (‘catching fire’ in case of data races on non-atomics), their underlying semantics helps their logic, whereas such a notion would have to be derived for Arm-A.

Less relaxed memory models. GPS [Turon et al. 2014] targets the subset of C11 featuring release stores, acquire loads, and non-atomic accesses (on which data races are undefined behaviour). GPS features ghost state (which is

sound because $po \cup rf$ is acyclic), per-location protocols (carrying state transition system tokens), and escrows. Kaiser et al. [Kaiser et al. 2017] describe how the protocols and escrows of GPS can be defined in terms of simpler components (like invariants) in Iris, and this is part of our motivation for using Iris.

To enrich GPS with the expressive power of Iris ‘for free’, iGPS [Kaiser et al. 2017] is based on an operational reformulation of release-acquire. One of our contributions is to show how to reason about an axiomatic memory model directly in Iris, avoiding this operational reformulation. iRC11 [Dang et al. 2020] combines iGPS with FSL++. It targets ORC11, an operational reformulation of a fragment of RC11.

Compass [Dang et al. 2022] is a specification framework for the ORC11 memory model that gives programs specifications in terms of event graphs the inter-thread edges of which are induced by data structure operations, for example from an enqueue to the dequeue of the same value, which generalises rf .

Cosmo [Mével et al. 2020] is a logic for the multicore OCaml memory model. The OCaml memory model is stronger and simpler than that of C11 in many respects which Cosmo leverages extensively to derive simple but powerful reasoning rules. Following iGPS, Cosmo is an instantiation of the standard Iris framework with an operational semantics, following the general recipe, and has a layered design featuring a base logic exposing memory model details and a high-level logic with almost standard CSL proof rules.

Very strong models. For other stronger models like TSO that have a simple operational model, working that close to the axiomatic model is probably more a burden, although possible. Significantly different abstractions would need to be built on top to capture this strength so that the logic is usable [Jacobs 2014; Ridge 2010; Sieczkowski et al. 2015; Wehrman 2012; Wehrman and Berdine 2011; Zhang and Feng 2014].

Weak and persistent memory. There are efforts to build logics for weak and persistent memory models [Bila et al. 2022; Raad et al. 2020; Stefanescu et al. 2024; Vindum and Birkedal 2023]. The out-of-order behaviours of these models are relatively mild (e.g., TSO and the release-acquire fragment of C11), while handling persistency and the intricate combination of weak memory and persistent aspects in a logic is challenging.

Per-location protocols. As the name suggests, a per-location protocol in GPS and iGPS [Kaiser et al. 2017; Turon et al. 2014] is a logical assertion dedicated to resource transfer between memory operations of a location, in contrast to invariants, which are implicitly shared among all locations. Per-location protocols make it possible to bind the physical value of a location to an abstract state of a state transition system (STS), and ensures that the evolution of the value is consistent with the transitions of the STS. This abstraction usually enables more high-level proofs (compared to their counterparts in AxSL), and can be implemented using basic Iris building blocks (invariants and ghost states), as in iGPS. The idea is based on the memory model assumption that the accesses to individual locations have SC behaviours (SC-per-location is also known coherence) which is often phrased in axiomatic memory models as an acyclicity requirement on the (extended) coherence order, eco (e.g. the Internal visibility requirement of Arm-A in Fig. 4). Intuitively, given a pre-agreed STS, the axiom prevents reading from an old state and forces to make a valid and consistent transition when writing. In the fragments of C11 that GPS and iGPS are based on, the synchronisation order $po \cup rf$ is included in eco , and so threads can exchange resources describing the abstract states of the same location along the eco edges by rely-guarantee reasoning.

As noted in §5.1, our AxSL protocol Φ is heavily inspired by GPS, but tailored to only support relatively simple resource transfer, due to the limited support for the coherence reasoning described in §8.3. Concretely, our logic is parametric by a concrete stateless protocol whose implementation does not rely on the coherence axiom. We leave implementing full stateful protocols in AxSL as future work, since enabling meaningful proof rules with stateful protocols requires a non-trivial extension to AxSL that solves more challenging circularity issues. Fundamentally, such an extended logic would need to support flowing resources not only along the coherence

order and the synchronisation order respectively, but also between the two orders — even though the union of these two orders is potentially cyclic, as in Arm-A.

Dealing with backtracking. The trick we use to express an axiomatic memory model as an operational semantics of the shape that Iris expects is inspired by how Islaris treats its event-enriched SMT language [Sammler et al. 2022]. A $(\text{declare-const } x); s$ program can take a step to $s[x \mapsto v]$ for any value v , and an $(\text{assert } e); s$ program can take a step in one of two ways: if e evaluates to true, the program takes a step to s ; and if e evaluates to false, the assert steps to the ‘execution discarded’ state. The definition of postcondition in Islaris then ignores discarded states.

Tracking ordering and flowing. The Lace logic [Bornat et al. 2015] shares the same core idea of explicitly tracking ordering between memory events (which they call ‘laces’), and of flowing assertions along edges (which they call ‘embroidery’) of an axiomatic memory model. Their setup looks quite different on the surface, as their approach to ordering is top-down (in the style of Crary and Sullivan [Crary and Sullivan 2015]), stating which instructions they require ordering from, rather than our bottom-up approach, in which we infer order from the instructions of the program. The main difference is that they try to talk about variables (memory locations), and so, to soundly flow assertions along edges, they need to check for interference on said variables, which is a whole-program check. In addition, they leave supporting separation as future work, and thus feature no notion of transfer of resources. However, Lace features modalities to ease reasoning about coherence, whereas it has to be done by graph reasoning in our logic. Lace was implemented using a custom proof checker, with no proof of its soundness.

Tracking memory events. The Ogre and Pythia invariance proof method [Alglave and Cousot 2017] refines Owicki-Gries without auxiliary variables (which are unsound for relaxed memory [Lahav and Vafeiadis 2015]) by working with memory events (via program counters), and their “pythia” variables keep track of the values of reads. Their method is parameterised by an axiomatic memory model expressed by relational algebra in the .cat format [Alglave et al. 2014]. They show that their proof method is sound and relatively complete, but their invariants are whole-program, and they leave development of abstractions that make proofs tractable as future work.

Tied-tos and flow implications. We conjecture that the unpublished extension of ribbon proofs to relaxed memory mentioned in [Wickerson et al. 2013] would have had some similarities to our setup, with unclosed ribbons standing for tied assertions, and flow implications imposing conditions on when ribbons can be joined.

10 Conclusion

The very relaxed concurrency memory models of hardware architectures like Arm-A, in which synchronisation (ob for Arm-A) does not follow program order, make syntax-directed and thread-modular reasoning challenging. The need to program directly to the hardware architecture for performance and for access to systems features makes this challenge unavoidable. Our family of logics, AxSL, addresses this challenge through assertions that track how synchronisation is induced by the program text, and makes reasoning tractable by flowing higher-order ghost state along synchronisation. This allows us to capture and thus validate key synchronisation idioms. Moreover, as demonstrated by our instantiation of AxSL to different memory models, our approach relies essentially just on the acyclicity of the synchronisation order, and should therefore apply to many hardware architectures.

This opens up a potential approach for reasoning about a wide range of axiomatic models, and there are many important extensions to explore, e.g. to integrate with the full Arm-A or RISC-V ISAs, to cover mixed-size accesses, and to cover systems features including instruction-fetch and virtual memory. An important challenge is to tackle

reasoning involving not only synchronisation but also coherence, in particular as leveraged by non-atomics, even though the union of synchronisation and coherence can have cycles.

A Further Proof Rules

A.1 MemWrite Rule

$$\begin{array}{l}
\text{HT-MICRO-MEMWRITE} \\
\left\{ \begin{array}{l}
\textcircled{1} (\text{LastLocalWrite}(x, v_{\text{last}}) \vee \text{NoLocalWrites}(x)) * d = (\text{dom}(\text{regs}), \emptyset) * \\
\text{PoPred}(e_{\text{po}}) * \text{CtrlPreds}(\text{srcs}_{\text{ctrl}}) * \\
*_{(r \mapsto (v, E)) \in \text{regs}} r \mapsto v @ E * *_{(e_{\text{lob}} \mapsto P_{\text{lob}}) \in m} (e_{\text{lob}} \rightsquigarrow P_{\text{lob}}) * \\
\forall e. \left(\begin{array}{l}
\textcircled{2} \text{GraphFacts}'(os, vr, x, v, e, e_{\text{po}}, \text{srcs}_{\text{ctrl}}, d, \text{regs}) -* \\
(\text{Lob}(\text{dom}(m), e) * \textcircled{3} \text{Flow}'_{\Phi}(e, x, v, m, P))
\end{array} \right)
\end{array} \right\} \\
\text{MemWrite } os \text{ } vr \text{ } x \text{ } v \text{ } d \\
\left\{ \begin{array}{l}
\exists e. D = \{e\} * \textcircled{4} \text{LastLocalWrite}(x, v) * \text{PoPred}(e) * \text{CtrlPreds}(\text{srcs}_{\text{ctrl}}) * \\
(v, D). \textcircled{5} *_{(r \mapsto (v, E)) \in \text{regs}} r \mapsto v @ E * \textcircled{6} \text{GraphFacts}'(os, vr, x, v, e, e_{\text{po}}, \text{srcs}_{\text{ctrl}}, d, \text{regs}) * \\
\textcircled{7} e \rightsquigarrow P(x, v)
\end{array} \right\}_{tid, \Phi}
\end{array}$$

Fig. 48. A proof rule of AxSL^{Arm} for the MemWrite microinstruction. As for the MemRead rule in Fig. 23 the user provides m , a thread-local map from events to the resources consumed.

Perhaps surprisingly, the basic rule for MemWrite is extremely similar to that for MemRead given in Fig. 23, so we only highlight the key differences here. Unlike the read rule, the write rule is not affected by the presence of previous local writes, and so can take either a NoLocalWrites or a LastLocalWrite resource in $\textcircled{1}$. Because a new local write is produced, whichever resource is passed in $\textcircled{1}$, a LastLocalWrite carrying the new event is returned $\textcircled{4}$. The definition of GraphFacts' at $\textcircled{2}$ and $\textcircled{6}$ differs from the version used in Fig. 23 because write events induce a different collection of incoming edges, most importantly lacking an incoming rf edge. Similarly the definition of Flow' $\textcircled{3}$ is changed from showing the incoming tied resources and protocol imply the new tied resource to showing the incoming tied resources imply the new tied resource and the outgoing protocol. Finally we note that the produced tied resource $\textcircled{7}$ is parameterised only on the value written and the identifier of the new node, not on the identifier of some write event being read from.

A.2 MemRead Rule with Local Writes

$$\begin{array}{l}
\text{HT-MICRO-MEMREAD-RDEP-EXT-LOCAL} \\
\left\{ \begin{array}{l}
\textcircled{1} \text{LastLocalWrite}(x, v') * d = (\text{dom}(\text{regs}), \emptyset) * \\
\text{PoPred}(e_{\text{po}}) * \text{CtrlPreds}(\text{srcs}_{\text{ctrl}}) * \\
*_{(r \mapsto (v, E)) \in \text{regs}} r \mapsto v @ E * *_{(e_{\text{lob}} \mapsto P_{\text{lob}}) \in m} (e \rightsquigarrow P_{\text{lob}}) * \\
\forall e, v, e_w. \left(\begin{array}{l}
\text{GraphFacts}(e, os, vr, x, v, e_w, e_{\text{po}}, \text{srcs}_{\text{ctrl}}, d, \text{regs}) -* \\
(\text{Lob}(\text{dom}(m), e) * \text{Flow}_{\Phi}(e, x, v, e_w, m, P))
\end{array} \right)
\end{array} \right\} \\
\text{MemRead } os \text{ } vr \text{ } x \text{ } d \\
\left\{ \begin{array}{l}
\exists e, e_w. D = \{e\} * \textcircled{2} \text{LastLocalWrite}(x, v') * \text{PoPred}(e) * \text{CtrlPreds}(\text{srcs}_{\text{ctrl}}) * \\
(v, D). \text{GraphFacts}(e, os, vr, x, v, e_w, e_{\text{po}}, \text{srcs}_{\text{ctrl}}, d, \text{regs}) * *_{(r \mapsto (v, E)) \in \text{regs}} r \mapsto v @ E * \\
\textcircled{3} ((e \rightsquigarrow P(x, v, e_w)) \vee v = v')
\end{array} \right\}_{tid, \Phi}
\end{array}$$

Fig. 49. A proof rule of AxSL^{Arm} for the MemRead microinstruction, similar to that shown in Fig. 23, but specialised to handle the case where the current thread contains a previous local write to the address being read from.

Fig. 49 shows a variant of the MemRead rule in Fig. 23 which allows for a previous local write to the address being read from. The prerequisites of the rule differ only in that it requires ① LastLocalWrite instead of NoLocalWrites, which is still returned unchanged ②. The conclusion becomes a disjunction ③, if an external write is read from then the rule produces the expected tied resource. However if the local write is read from no resource transfer takes place, because thread local reads do not provide ob ordering, so we learn only that the value read is equal to the value written in the most recent local write. While this does not provide as powerful a mechanism for resource transfer as might be desired, we do expect it to allow some reasoning since in synchronisation primitives it will generally be necessary for an external thread to write distinct values to those written by the local thread.

B Recovering Points-Tos in AxSL^{SC}

In this section, we discuss the connection between the graph reasoning of opax-based logics and the heap reasoning of standard CSLs. Concretely, we show formally that we can build a notion of points-tos with abstractions in AxSL^{SC}.

Standard CSL points-tos. We first recap the points-to assertion and two standard proof rules using it in CSL. A usual CSL points-to assertion $x \hookrightarrow v$ means that the latest value of location x is v on the shared heap. Owning this assertion grants one the exclusive right to access x with the two standard rules in §B. Our goal is to obtain a definition of points-to that satisfies the same rules in AxSL^{SC}.

$$\begin{array}{ll}
 \text{STD-HT-MICRO-MEMREAD} & \text{STD-HT-MICRO-MEMWRITE} \\
 \{x \hookrightarrow v\} \text{MemRead } x \{w. w = v * x \hookrightarrow v\}_{tid, \Phi} & \{x \hookrightarrow v\} \text{MemWrite } x v' \{_. x \hookrightarrow v'\}_{tid, \Phi}
 \end{array}$$

Fig. 50. Standard CSL read and write rules with points-tos (if we omit the protocol Φ). We reformulate them with the microinstruction language and assume a fixed number of hardware threads.

B.0.1 Raw Definition. In AxSL^{SC}, we model $x \hookrightarrow v$ by leveraging the fact that all observed writes on location x are ordered by co, forming a sequence of write events, and that one can only read from the latest write (head) of the sequence with value v : reading from outdated writes would violate the acyclicity requirement. We define a raw points-to to capture these observations:

$$x \hookrightarrow_{raw} v @ e_{lst} \triangleq \exists \gamma_x. [\circ [x \mapsto \gamma_x]]^{! \gamma} * \exists ch. [\bullet (e_{lst} :: ch)]^{! \gamma_x} * \text{IsLastWrite}(e_{lst}, ch) * e_{lst} \cdot W \ x \ v$$

This raw definition asserts that the globally co-latest write of x has event ID e_{lst} and value v . We use two ghost states in the definition to capture this. $[\circ [x \mapsto \gamma_x]]^{! \gamma}$ is a persistent fact binding location x to a unique ghost name γ_x . γ_x is the ghost name of a list of event IDs $[\bullet (e_{lst} :: ch)]^{! \gamma_x}$ which models the sequence of writes of x (the head is the latest). This ghost list ensures that updates of the sequence of the writes are monotone (list only grows), and one can at anytime take (persistent) snapshots of the list: $[\circ \overline{ch'}]^{! \gamma_x}$ such that ch' is a sub-list. The IsLastWrite predicate establishes that e_{lst} is indeed the co-latest write in the list with graph facts, and finally, $e_{lst} \cdot W \ x \ v$ remembers the value v .

Leveraging protocol. To ensure that the sequence of writes is well-synchronised with the program execution, that is, all writes one has reasoned about were added to the sequence, and, dually, reads can only read from the writes in the sequence, we use our rely-guarantee protocol Φ . As a first step, we fix the protocol using a new predicate WriteOf(x, e) as $\Phi(x, v, e) \triangleq \text{WriteOf}(x, e)$ (this can be further generalised, but we keep it simple for

now):

$$\text{WriteOf}(x, e) \triangleq \exists \gamma_x. [\circ [x \mapsto \gamma_x]]^Y * \exists ch. [\circ (e :: ch)]^{Yx} * e : W \ x \ v$$

The new predicate asserts that e is an observed write event of x , and thus is included in the write sequence, depicted as **SC-PT-RAW-AG**. As we will see below, this is crucial for concluding that only the latest write is readable, mimicking the semantics of points-tos.

$$\begin{array}{c} \text{SC-PT-RAW-AG} \\ \frac{x \hookrightarrow_{\text{raw}} v @ e * \text{WriteOf}(x, e')}{e' \text{ co}^* e} \end{array} \quad \begin{array}{c} \text{SC-PT-RAW-SHT} \\ x \hookrightarrow_{\text{raw}} v @ e \Rightarrow \text{WriteOf}(x, e) \end{array} \quad \begin{array}{c} \text{SC-PT-RAW-UPD} \\ \frac{e \text{ co } e' * e' : W \ x \ v'}{x \hookrightarrow_{\text{raw}} v @ e \Rightarrow x \hookrightarrow_{\text{raw}} v' @ e'} \end{array}$$

$$\begin{array}{c} \text{SC-PT-WO-PERS} \\ \frac{\text{WriteOf}(x, e)}{\text{WriteOf}(x, e) * \text{WriteOf}(x, e)} \end{array}$$

Fig. 51. Selected operations of the raw points-to and WriteOf

B.0.2 Full Definition. The full points-to assertion is modeled as a monotone predicate over the po-latest event e of a thread, using the raw definition:

$$\llbracket x \hookrightarrow v \rrbracket \triangleq \lambda e. \exists e_{\text{lst}}. x \hookrightarrow_{\text{raw}} v @ e_{\text{lst}} * e_{\text{lst}} \leq_{\text{sc}} e$$

where the notation $e \leq_{\mathcal{R}} e'$ means that either the two events are identical, or e is \mathcal{R} -ordered before e' (in contrast, \geq means identical or after). We instantiate this relation to sc , requiring that the latest write e_{lst} happens before e . We also need to monotonise all other assertions of the base AxSL^{SC} logic. Most of them are standard, for instance below is how we monotonise the weakest precondition.

$$\begin{aligned} & \llbracket \{P\} i \{w. Q(w)\}_{\text{tid}, \Phi} \rrbracket \triangleq \\ & \lambda e. \forall e' \geq_{\text{po}} e. \{ \llbracket P \rrbracket (e') * \text{PoPred}(e') \} i \{ v. \exists e'' \geq_{\text{po}} e'. \llbracket Q(v) \rrbracket (e'') * \text{PoPred}(e'') \}_{\text{tid}, \Phi} \end{aligned}$$

B.0.3 Proving Standard CSL Rules. We now sketch the soundness proof of the two classic CSL rules in AxSL^{SC} . The proof starts with unfolding all the monotone predicates then proceeds using AxSL^{SC} proof rules.

Store. For **STD-HT-MICRO-MEMWRITE**, after unfolding, we need to show:

$$\begin{aligned} & \{ \text{PoPred}(e) * \exists e_{\text{lst}}. x \hookrightarrow_{\text{raw}} v @ e_{\text{lst}} * e_{\text{lst}} \leq_{\text{sc}} e \} \\ & \text{MemWrite } x \ v' \\ & \{ () . \exists e' \geq_{\text{po}} e. \text{PoPred}(e') * \exists e_{\text{lst}}. x \hookrightarrow_{\text{raw}} v' @ e_{\text{lst}} * e_{\text{lst}} \leq_{\text{sc}} e' \}_{\text{tid}, \Phi} \end{aligned}$$

We proceed with the base rule **SC-HT-MICRO-MEMWRITE**, picking e as e_{po} . We have the following view shift as a sub-goal, which guarantees that the protocol is enforced on the new write (namely the new write is added to the sequence as the latest):

$$\begin{aligned} & \forall e_w. (\text{GraphFactsW}(e_w, x, v', e) * \exists e_{\text{lst}}. x \hookrightarrow_{\text{raw}} v @ e_{\text{lst}} * e_{\text{lst}} \leq_{\text{sc}} e) \\ & \Rightarrow (\text{WriteOf}(x, e_w)) * \exists e_{\text{lst}}. x \hookrightarrow_{\text{raw}} v' @ e_{\text{lst}} * e_{\text{lst}} \leq_{\text{sc}} e_w \end{aligned}$$

We update the raw points-to assertion to $x \hookrightarrow_{\text{raw}} v' @ e_w$ by **SC-PT-RAW-UPD** and then take a snapshot with **SC-PT-RAW-SHT**. The remaining goal $e_w \leq_{\text{sc}} e_w$ is trivial. Finally, we conclude the proof with the rule of consequence.

Load. In the case of **STD-HT-MICRO-MEMREAD**, we need to show:

$$\begin{aligned} & \{ \text{PoPred}(e) * \exists e_{lst}. x \hookrightarrow_{raw} v @ e_{lst} * e_{lst} \leq_{sc} e \} \\ & \text{MemRead } x \\ & \{ w. \exists e' \geq_{po} e. w = v * \text{PoPred}(e') * \exists e_{lst}. x \hookrightarrow_{raw} v @ e_{lst} * e_{lst} \leq_{sc} e' \}_{tid, \Phi} \end{aligned}$$

We proceed by applying **SC-HT-MICRO-MEMREAD**. We need to prove the following view shift which captures that we are reading from a write that is in the sequence:

$$\begin{aligned} & \forall e_r, w, e_w. (\text{GraphFactsR}(e_r, x, v, e_w, e) * \exists e_{lst}. x \hookrightarrow_{raw} v @ e_{lst} * e_{lst} \leq_{sc} e * \text{WriteOf}(x, e_w)) \\ & \Rightarrow w = v * \exists e_{lst}. x \hookrightarrow_{raw} v @ e_{lst} * e_{lst} \leq_{sc} e_r * \text{WriteOf}(x, e_w) \end{aligned}$$

By **SC-PT-RAW-AG**, we know $e_w \leq_{co} e_{lst}$, namely the write e_w that we are reading from is one of the observed writes. We show $w = v$ by showing $e_w = e_{lst}$, that is, we can only read from the latest write. This is done by showing a violation of the acyclicity of sc in the other case when reading from an old write ($e_w co e_{lst}$). The problematic cycle is $e_{lst} sc e po e_r fr e_{lst}$ where $e_r fr e_{lst}$ is induced from $e_w rf e_r$ and $e_w co e_{lst}$. Again, we conclude the proof with the rule of consequence.

C An Implementation of Logical Interpretations and Assertions for AxSL^{Arm}

This section requires knowledge on the built-in CMRA constructors of Iris ⁶.

C.1 Reserved Ghost Names

We have a collection of reserved ghost names that assert ownership of various ghost resources, which are used to implement state interpretations and assertions of AxSL^{Arm}. We list their notations and usages in Fig. 52.

Notation	Global	Notation	Local
γ_G	Graph facts	γ_R	Register points-tos
γ_I	Instruction points-tos	γ_L	Local writes
γ_T	Tied-tos	γ_P	po predecessors
γ_E	Exclusive tokens	γ_C	ctrl predecessors
		γ_M	rmw predecessors

Fig. 52. Each thread has its own set of local ghost names.

C.2 Global Assertions, SI, and SI_T

Persistent facts. The *Ag* constructor of Iris is used to implement graph and instruction memory. The ghost resources and state interpretation SI using it are therefore persistent.

⁶There is not a comprehensive documentation for all the constructors, we therefore refer readers to the Iris Rocq repository for details: <https://gitlab.mpi-sws.org/iris/iris/-/tree/master/iris/algebra>.

$$\begin{aligned}
\text{Sl}(\sigma) &\triangleq [\text{ag}(\sigma.X) : \text{Ag}(\text{Graph})]^{YG} * [\text{ag}(\sigma.I) : \text{Ag}(\text{InstMem})]^{YI} \\
a \text{ R } b &\triangleq \exists X. [\text{ag}(X) : \text{Ag}(\text{Graph})]^{YG} * (a, b) \in X.R \quad (\text{R is a relation}) \\
a:E &\triangleq \exists X. [\text{ag}(X) : \text{Ag}(\text{Graph})]^{YG} * X(a) = E \quad (\text{E is an event}) \\
a \mapsto i &\triangleq \exists I. [\text{ag}(I) : \text{Ag}(\text{InstMem})]^{YI} * I(a) = i
\end{aligned}$$

Tied-to. We interpret the tied-to map τ (of type $\text{Eid} \rightarrow \text{iProp}$) with two levels of indirection. First, instead of constructing a complex CMRA that directly takes the map, we use a simpler CMRA GhostMapAg for a map em that only goes to GName from Eid , where GName is used to manage more ghost states using simpler CMRAs. GhostMapAg is a standard authoritative ghost map CMRA constructor except for being insert-only (the fragmental views are persistent). Second, for every $e : \text{Eid}$, the corresponding ghost name $\gamma = em(e)$ is used to track another map γm from GName to iProp . This map allows us to track fractions of the proposition that tied to e , which is crucial for splitting tied-to assertions. We use SetDisjAuth and SavedProp of Iris for this map. SetDisjAuth is an authoritative ghost set (with disjoint union) CMRA constructor that we use to track a collection of ghost names, each of which owns a fraction of the tied proposition of a node using the SavedProp CMRA.

$$\begin{aligned}
\text{Sl}_\top(\tau) &\triangleq \exists em : \text{Eid} \rightarrow \text{GName}. [\bullet em : \text{GhostMapAg}(\text{Eid}, \text{GName})]^{YI} * \text{dom}(em) = \tau * \\
&\quad *_{(e \mapsto \gamma) \in em} \left(\exists \gamma m : \text{GName} \rightarrow \text{iProp}. [\bullet \text{dom}(\gamma m) : \text{SetDisjAuth}(\text{Eid})]^{Y'} * \right. \\
&\quad \quad \left. *_{(\gamma' \mapsto R) \in \gamma m} [\text{sp}(\frac{1}{2}, R) : \text{SavedProp}]^{Y'} * \triangleright \square(\tau(e) * *_{(_ \mapsto R) \in \gamma m} R) \right) \\
a \rightsquigarrow P &\triangleq \exists \gamma, \gamma'. [\circ e \mapsto \gamma : \text{GhostMapAg}(\text{Eid}, \text{GName})]^{YI} * [\circ \{\gamma'\} : \text{SetDisjAuth}(\text{Eid})]^{Y'} * \\
&\quad [\text{sp}(\frac{1}{2}, P) : \text{SavedProp}]^{Y'}
\end{aligned}$$

Token for exclusives. We extend Sl_\top with $[\bullet(\text{dom}(\tau)) : \text{SetDisjAuth}(\text{Eid})]^{YE}$, and define $\text{ExTok}(e)$ as $[\circ\{e\} : \text{SetDisjAuth}(\text{Eid})]^{YE}$.

C.3 Local Assertions and LSI

Register points-to. We interpret the register file and define a register points-to using the standard authoritative ghost map CMRA constructor of Iris, GhostMap .

$$r \mapsto v @ E \triangleq [\circ r \mapsto (v, E) : \text{GhostMap}(\text{RegName}, \text{Val} \times \text{Set}(\text{Eid}))]^{YR}$$

Local writes. We implement the local write assertions using the same ghost map CMRA, but on a map lm from Addr to $\text{option}(\text{Eid})$, which we associate with the execution graph in LW in LSI. LW contains the authoritative view of the map, requiring that for any location a if $a \mapsto \text{Some}(e) \in lm$ then e is indeed the latest write of a that the thread has been observed at the current program point; and if $a \mapsto \text{None} \in lm$, no writes of a have been observed.

$$\begin{aligned}
\text{LW}(cntr)_{tid} &\triangleq \exists X. [\text{ag}(X) : \text{Ag}(\text{Graph})]^{YG} * \exists lm : \text{Addr} \rightarrow \text{option}(\text{Eid}). \\
&\quad [\bullet lm : \text{GhostMap}(\text{Addr}, \text{option}(\text{Eid}))]^{YL} * \\
&\quad (\forall a \mapsto \text{Some}(e) \in lm. \text{IsLastWriteOn}(a, e, \text{OldWrites}(X, cntr)_{tid})) * \\
&\quad (\forall a \mapsto \text{None} \in lm. \text{NoWritesOn}(a, \text{OldWrites}(X, cntr)_{tid}))
\end{aligned}$$

$$\text{OldWrites}(X, cntr)_{tid} \triangleq \text{filter}(\lambda(e, E). e.tid = tid \wedge e.cntr < cntr \wedge \text{IsWrite}(E)) X.E$$

We define the assertions as just fragmental views of the map.

$$\begin{aligned} \text{NoLocalWrites}(a) &\triangleq \boxed{\circ a \mapsto \text{None} : \text{GhostMap}(\text{Addr}, \text{option}(\text{Eid}))}^{YL} \\ \text{LastLocalWrite}(a, e) &\triangleq \boxed{\circ a \mapsto \text{Some}(e) : \text{GhostMap}(\text{Addr}, \text{option}(\text{Eid}))}^{YL} \end{aligned}$$

po predecessors. We define a new CMRA *MonoNatPair* for implementing PoPred which tracks the latest po predecessor. This CMRA is inspired by the *MonoNat* CMRA of Iris, which ensures that a natural number is increment-only. We take the idea, and modify it to take a pair of natural numbers, and change the order to the lexicographical order. We use the CMRA to track a lower bound of $IT.cnt$, the counter used to generate fresh *Eid*, in an indirect way in the PoP predicate in LSI. We have a tweaked oneshot CMRA \boxed{o}^{YP} where o can be a *splittable* pending or shot(γ). In the latter case, γ owns a half of the authoritative view of cnt : $\boxed{\bullet^{1/2} cnt : \text{MonoNatPair}}^Y$. We use the other half to implement the PoPred assertion. The assertion has a persistent variant PoPred' which we implement using the fragmental view, which can also be used to form po edges.

$$\begin{aligned} \text{PoPred}(a) &\triangleq \exists \gamma. \boxed{\text{shot}(\gamma)}^{YP} * \boxed{\bullet^{1/2} a.cnt : \text{MonoNatPair}}^Y \\ \text{PoPred}(-) &\triangleq \boxed{1/2 \text{ pending}}^{YP} \\ \text{PoPred}'(a) &\triangleq \exists \gamma. \boxed{\text{shot}(\gamma)}^{YP} * \boxed{\circ a.cnt : \text{MonoNatPair}}^Y \end{aligned}$$

$$\begin{aligned} \text{PoP}(cnt)_{tid} &\triangleq \exists X. \boxed{\text{ag}(X) : \text{Ag}(\text{Graph})}^{YG} * \boxed{1/2 \text{ pending}}^{YP} \vee (\exists \gamma, cnt' < cnt. \\ &\quad \boxed{\text{shot}(\gamma)}^{YP} * \boxed{\bullet^{1/2} cnt' : \text{MonoNatPair}}^Y * \text{IsValidEid}(\langle tid, cnt' \rangle, X)) \end{aligned}$$

ctrl predecessors. We use $DfracAgree(\text{Set}(\text{Eid}))$ to implement CtrlPreds. CtrlPreds contains half of the CMRA, and CoP contains the other half to track $T.srsc_{ctrl}$.

$$\begin{aligned} \text{CtrlPreds}(s) &\triangleq \exists s' \supseteq s. \boxed{1/2 s' : DfracAgree(\text{Set}(\text{Eid}))}^{YC} \\ \text{CoP}(srsc_{ctrl}) &\triangleq \boxed{1/2 srsc_{ctrl} : DfracAgree(\text{Set}(\text{Eid}))}^{YC} \end{aligned}$$

rmw predecessors. We use $DfracAgree(\text{option}(\text{Eid}))$ to implement RmwPred. Similar to how CtrlPreds is implemented, we divide the CMRA into two halves, put one half in LSI, and use the other as the assertion.

$$\begin{aligned} \text{RmwPred}(e) &\triangleq \boxed{1/2 \text{ Some}(e) : DfracAgree(\text{option}(\text{Eid}))}^{YM} \\ \text{RmwPred}(-) &\triangleq \boxed{1/2 \text{ None} : DfracAgree(\text{option}(\text{Eid}))}^{YM} \\ \text{RoP}(srsc_{rmw}) &\triangleq \boxed{1/2 srsc_{rmw} : DfracAgree(\text{option}(\text{Eid}))}^{YM} \end{aligned}$$

mrd. We omit *mrd* in the presentation of the proof rules. We can have a ghost state using $DfracAgree(\text{Set}(\text{Eid}))$ to track the set of instruction internal memory reads, or we can just expose it to the weakest preconditions for microinstructions, and later hide it in the proof rules for instructions (since it is always empty at the beginning of an instruction).

Local state interpretation. Finally, we define LSI:

$$\begin{aligned} \text{LSI}(T)_{tid} &\triangleq \boxed{\bullet T.regs : \text{GhostMap}(\text{RegName}, \text{Val} \times \text{Set}(\text{Eid}))}^{YR} * \text{PoP}(T.IT.cnt)_{tid} * \\ &\quad \text{CoP}(T.srsc_{ctrl}) * \text{RoP}(T.srsc_{rmw}) * \text{LW}(T.IT.cnt)_{tid} \end{aligned}$$

Acknowledgments

We thank Amin Timany for helpful discussions concerning framing.

This work was supported in part by Google, through ASPIRE faculty awards and other funding to Birkedal, Pichon-Pharabod, and Sewell; in part by Arm (Sewell); by the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement No. 789108, AdG ELVER, Sewell); by a Villum Investigator grant (no. 25804), Center for Basic Research in Program Verification (CPV), from the VILLUM Foundation (Birkedal); by an AUFF starter grant (Pichon-Pharabod); and by the Innovate UK project Digital Security by Design (DSbD) Technology Platform Prototype, 105694.

References

- Jade Alglave and Patrick Cousot. 2017. Ogre and Pythia: an invariance proof method for weak consistency models. In *Proceedings of the 44th ACM SIGPLAN Symposium on Principles of Programming Languages, POPL 2017, Paris, France, January 18–20, 2017*, Giuseppe Castagna and Andrew D. Gordon (Eds.). ACM, 3–18. doi:10.1145/3009837.3009883
- Jade Alglave, Will Deacon, Richard Grisenthwaite, Antoine Hacquard, and Luc Maranget. 2021. Armed Cats: Formal Concurrency Modelling at Arm. *ACM Trans. Program. Lang. Syst.* 43, 2 (2021), 8:1–8:54. doi:10.1145/3458926
- Jade Alglave, Richard Grisenthwaite, Artem Khyzha, Luc Maranget, and Nikos Nikoleris. 2024. Puss in Boots: Formalizing Arm’s Virtual Memory System Architecture. *IEEE Micro* 44, 6 (2024), 83–91. doi:10.1109/MM.2024.3422668
- Jade Alglave, Luc Maranget, Paul E. McKenney, Andrea Parri, and Alan S. Stern. 2018. Frightening Small Children and Disconcerting Grown-ups: Concurrency in the Linux Kernel. In *Proceedings of the Twenty-Third International Conference on Architectural Support for Programming Languages and Operating Systems, ASPLOS 2018, Williamsburg, VA, USA, March 24–28, 2018*, Xipeng Shen, James Tuck, Ricardo Bianchini, and Vivek Sarkar (Eds.). ACM, 405–418. doi:10.1145/3173162.3177156
- Jade Alglave, Luc Maranget, Susmit Sarkar, and Peter Sewell. 2010. Fences in Weak Memory Models. In *Computer Aided Verification, 22nd International Conference, CAV 2010, Edinburgh, UK, July 15–19, 2010. Proceedings (Lecture Notes in Computer Science, Vol. 6174)*, Tayssir Touili, Byron Cook, and Paul B. Jackson (Eds.). Springer, 258–272. doi:10.1007/978-3-642-14295-6_25
- Jade Alglave, Luc Maranget, and Michael Tautschnig. 2014. Herding Cats: Modelling, Simulation, Testing, and Data Mining for Weak Memory. *ACM Trans. Program. Lang. Syst.* 36, 2 (2014), 7:1–7:74. doi:10.1145/2627752
- Arm Ltd. 2023. *ARM Architecture Reference Manual (for A-profile architecture)*. Arm Ltd. ARM DDI 0487J.a (ID042523), <https://developer.arm.com/documentation/ddi0487/latest/>, Accessed 2023-07-04.
- Alasdair Armstrong, Brian Campbell, Ben Simner, Christopher Pulte, and Peter Sewell. 2021. Isla: Integrating Full-Scale ISA Semantics and Axiomatic Concurrency Models. In *Computer Aided Verification - 33rd International Conference, CAV 2021, Virtual Event, July 20–23, 2021, Proceedings, Part I (Lecture Notes in Computer Science, Vol. 12759)*, Alexandra Silva and K. Rustan M. Leino (Eds.). Springer, 303–316. doi:10.1007/978-3-030-81685-8_14
- Mark Batty, Kayvan Memarian, Kyndylan Nienhuis, Jean Pichon-Pharabod, and Peter Sewell. 2015. The Problem of Programming Language Concurrency Semantics. In *ESOP*. 283–307. doi:10.1007/978-3-662-46669-8_12
- Mark Batty, Kayvan Memarian, Scott Owens, Susmit Sarkar, and Peter Sewell. 2012. Clarifying and compiling C/C++ concurrency: from C++11 to POWER. In *Proceedings of the 39th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2012, Philadelphia, Pennsylvania, USA, January 22–28, 2012*, John Field and Michael Hicks (Eds.). ACM, 509–520. doi:10.1145/2103656.2103717
- Mark Batty, Scott Owens, Susmit Sarkar, Peter Sewell, and Tjark Weber. 2011. Mathematizing C++ concurrency. In *Proceedings of the 38th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2011, Austin, TX, USA, January 26–28, 2011*, Thomas Ball and Mooly Sagiv (Eds.). ACM, 55–66. doi:10.1145/1926385.1926394
- P. Becker (Ed.). 2011. *Programming Languages — C++*. ISO/IEC 14882:2011. <http://www.open-std.org/jtc1/sc22/wg21/docs/papers/2011/n3242.pdf>.
- Eleni Vafeiadi Bila, Brijesh Dongol, Ori Lahav, Azalea Raad, and John Wickerson. 2022. View-Based Owicki-Gries Reasoning for Persistent x86-TSO. In *Programming Languages and Systems - 31st European Symposium on Programming, ESOP 2022, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2022, Munich, Germany, April 2–7, 2022, Proceedings (Lecture Notes in Computer Science, Vol. 13240)*, Ilya Sergey (Ed.). Springer, 234–261. doi:10.1007/978-3-030-99336-8_9
- Hans-Juergen Boehm and Sarita V. Adve. 2008. Foundations of the C++ concurrency memory model. In *Proceedings of the ACM SIGPLAN 2008 Conference on Programming Language Design and Implementation, Tucson, AZ, USA, June 7–13, 2008*, Rajiv Gupta and Saman P. Amarasinghe (Eds.). ACM, 68–78. doi:10.1145/1375581.1375591
- Richard Bornat, Jade Alglave, and Matthew J. Parkinson. 2015. New Lace and Arsenic: adventures in weak memory with a program logic. *CoRR* abs/1512.01416 (2015). arXiv:1512.01416 <http://arxiv.org/abs/1512.01416>
- Stephen Brookes and Peter W. O’Hearn. 2016. Concurrent separation logic. *ACM SIGLOG News* 3, 3 (2016), 47–65. doi:10.1145/2984450.2984457

- Soham Chakraborty and Viktor Vafeiadis. 2019. Grounding thin-air reads with event structures. *Proc. ACM Program. Lang.* 3, POPL (2019), 70:1–70:28. doi:10.1145/3290383
- Minki Cho, Sung-Hwan Lee, Dongjae Lee, Chung-Kil Hur, and Ori Lahav. 2022. Sequential reasoning for optimizing compilers under weak memory concurrency. In *PLDI '22: 43rd ACM SIGPLAN International Conference on Programming Language Design and Implementation, San Diego, CA, USA, June 13 - 17, 2022*, Ranjit Jhala and Isil Dillig (Eds.). ACM, 213–228. doi:10.1145/3519939.3523718
- Karl Crary and Michael J. Sullivan. 2015. A Calculus for Relaxed Memory. In *Proceedings of the 42nd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2015, Mumbai, India, January 15-17, 2015*, Sriram K. Rajamani and David Walker (Eds.). ACM, 623–636. doi:10.1145/2676726.2676984
- Sadegh Dalvandi, Simon Doherty, Brijesh Dongol, and Heike Wehrheim. 2020. Owicki-Gries Reasoning for C11 RAR. In *34th European Conference on Object-Oriented Programming, ECOOP 2020, November 15-17, 2020, Berlin, Germany (Virtual Conference) (LIPICs, Vol. 166)*, Robert Hirschfeld and Tobias Pape (Eds.). Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 11:1–11:26. doi:10.4230/LIPICs.ECOOP.2020.11
- Sadegh Dalvandi, Brijesh Dongol, Simon Doherty, and Heike Wehrheim. 2022. Integrating Owicki-Gries for C11-Style Memory Models into Isabelle/HOL. *J. Autom. Reason.* 66, 1 (2022), 141–171. doi:10.1007/S10817-021-09610-2
- Hoang-Hai Dang, Jacques-Henri Jourdan, Jan-Oliver Kaiser, and Derek Dreyer. 2020. RustBelt meets relaxed memory. *Proc. ACM Program. Lang.* 4, POPL (2020), 34:1–34:29. doi:10.1145/3371102
- Hoang-Hai Dang, Jaehwang Jung, Jaemin Choi, Duc-Thuan Nguyen, William Mansky, Jeehoon Kang, and Derek Dreyer. 2022. Compass: strong and compositional library specifications in relaxed memory separation logic. In *PLDI '22: 43rd ACM SIGPLAN International Conference on Programming Language Design and Implementation, San Diego, CA, USA, June 13 - 17, 2022*, Ranjit Jhala and Isil Dillig (Eds.). ACM, 792–808. doi:10.1145/3519939.3523451
- Will Deacon. 2016. The ARMv8 Application Level Memory Model. <https://github.com/herd/herdtools7/blob/master/herd/libdir/aarch64.cat>.
- Thomas Dinsdale-Young, Lars Birkedal, Philippa Gardner, Matthew J. Parkinson, and Hongseok Yang. 2013. Views: compositional reasoning for concurrent programs. In *The 40th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL '13, Rome, Italy - January 23 - 25, 2013*, Roberto Giacobazzi and Radhia Cousot (Eds.). ACM, 287–300. doi:10.1145/2429069.2429104
- Thomas Dinsdale-Young, Mike Dodds, Philippa Gardner, Matthew J. Parkinson, and Viktor Vafeiadis. 2010. Concurrent Abstract Predicates. In *ECOOP 2010 - Object-Oriented Programming, 24th European Conference, Maribor, Slovenia, June 21-25, 2010. Proceedings (Lecture Notes in Computer Science, Vol. 6183)*, Theo D'Hondt (Ed.). Springer, 504–528. doi:10.1007/978-3-642-14107-2_24
- Simon Doherty, Brijesh Dongol, Heike Wehrheim, and John Derrick. 2019. Verifying C11 programs operationally. In *Proceedings of the 24th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming, PPOPP 2019, Washington, DC, USA, February 16-20, 2019*, Jeffrey K. Hollingsworth and Idit Keidar (Eds.). ACM, 355–365. doi:10.1145/3293883.3295702
- Marko Doko. 2021. *Program Logic for Weak Memory Concurrency*. Ph. D. Dissertation. Kaiserslautern University of Technology, Germany. <https://kluedo.ub.rptu.de/frontdoor/index/index/docId/6679>
- Marko Doko and Viktor Vafeiadis. 2016. A Program Logic for C11 Memory Fences. In *Verification, Model Checking, and Abstract Interpretation - 17th International Conference, VMCAI 2016, St. Petersburg, FL, USA, January 17-19, 2016. Proceedings (Lecture Notes in Computer Science, Vol. 9583)*, Barbara Jobstmann and K. Rustan M. Leino (Eds.). Springer, 413–430. doi:10.1007/978-3-662-49122-5_20
- Marko Doko and Viktor Vafeiadis. 2017. Tackling Real-Life Relaxed Concurrency with FSL++. In *Programming Languages and Systems - 26th European Symposium on Programming, ESOP 2017, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2017, Uppsala, Sweden, April 22-29, 2017. Proceedings (Lecture Notes in Computer Science, Vol. 10201)*, Hongseok Yang (Ed.). Springer, 448–475. doi:10.1007/978-3-662-54434-1_17
- Andres Erbsen, Samuel Gruetter, Joonwon Choi, Clark Wood, and Adam Chlipala. 2021. Integration verification across software and hardware for a simple embedded system. In *PLDI '21: 42nd ACM SIGPLAN International Conference on Programming Language Design and Implementation, Virtual Event, Canada, June 20-25, 2021*, Stephen N. Freund and Eran Yahav (Eds.). ACM, 604–619. doi:10.1145/3453483.3454065
- Robert W. Floyd. 1967. Assigning Meanings to Programs. In *Proceedings of the Symposium in Applied Mathematics*, Vol. 19. American Mathematical Society, 19–32.
- 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 *Proceedings of the 44th ACM SIGPLAN Symposium on Principles of Programming Languages, POPL 2017, Paris, France, January 18-20, 2017*, Giuseppe Castagna and Andrew D. Gordon (Eds.). ACM, 429–442. doi:10.1145/3009837.3009839
- Lennard Gäher, Michael Sammler, Simon Spies, Ralf Jung, Hoang-Hai Dang, Robbert Krebbers, Jeehoon Kang, and Derek Dreyer. 2022. Simuliris: a separation logic framework for verifying concurrent program optimizations. *Proc. ACM Program. Lang.* 6, POPL (2022), 1–31. doi:10.1145/3498689
- Kourosh Gharachorloo. 1995. *Memory Consistency Models for Shared-Memory Multiprocessors*. Ph. D. Dissertation. Stanford University.

- Kathryn E. Gray, Gabriel Kerneis, Dominic P. Mulligan, Christopher Pulte, Susmit Sarkar, and Peter Sewell. 2015. An integrated concurrency and core-ISA architectural envelope definition, and test oracle, for IBM POWER multiprocessors. In *Proceedings of the 48th International Symposium on Microarchitecture, MICRO 2015, Waikiki, HI, USA, December 5-9, 2015*, Milos Prvulovic (Ed.). ACM, 635–646. doi:10.1145/2830772.2830775
- Angus Hammond, Zongyuan Liu, Thibaut Pérami, Peter Sewell, Lars Birkedal, and Jean Pichon-Pharabod. 2024. An Axiomatic Basis for Computer Programming on the Relaxed Arm-A Architecture: The AxSL Logic. *Proc. ACM Program. Lang.* 8, POPL, Article 21 (jan 2024), 34 pages. doi:10.1145/3632863
- Mengda He, Viktor Vafeiadis, Shengchao Qin, and João F. Ferreira. 2016. Reasoning about Fences and Relaxed Atomics. In *24th Euromicro International Conference on Parallel, Distributed, and Network-Based Processing, PDP 2016, Heraklion, Crete, Greece, February 17-19, 2016*. IEEE Computer Society, 520–527. doi:10.1109/PDP.2016.103
- Lisa Higham, LillAnne Jackson, and Jalal Kawash. 2007. Specifying memory consistency of write buffer multiprocessors. *ACM Trans. Comput. Syst.* 25, 1 (2007), 1. doi:10.1145/1189736.1189737
- C. A. R. Hoare. 1969. An Axiomatic Basis for Computer Programming. *Commun. ACM* 12, 10 (1969), 576–580. doi:10.1145/363235.363259
- Bart Jacobs. 2014. *Verifying TSO Programs (Report CW660)*. Technical Report.
- Alan Jeffrey and James Riely. 2016. On Thin Air Reads Towards an Event Structures Model of Relaxed Memory. In *Proceedings of the 31st Annual ACM/IEEE Symposium on Logic in Computer Science, LICS '16, New York, NY, USA, July 5-8, 2016*, Martin Grohe, Eric Koskinen, and Natarajan Shankar (Eds.). ACM, 759–767. doi:10.1145/2933575.2934536
- Jonas B. Jensen, Nick Benton, and Andrew Kennedy. 2013. High-Level Separation Logic for Low-Level Code. In *Proceedings of the 40th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (Rome, Italy) (POPL '13)*. Association for Computing Machinery, New York, NY, USA, 301–314. doi:10.1145/2429069.2429105
- Ralf Jung, Robbert Krebbers, Jacques-Henri Jourdan, Ales Bizjak, Lars Birkedal, and Derek Dreyer. 2018. Iris from the ground up: A modular foundation for higher-order concurrent separation logic. *J. Funct. Program.* 28 (2018), e20. doi:10.1017/S0956796818000151
- Ralf Jung, David Swasey, Filip Sieczkowski, Kasper Svendsen, Aaron Turon, Lars Birkedal, and Derek Dreyer. 2015. Iris: Monoids and Invariants as an Orthogonal Basis for Concurrent Reasoning. In *Proceedings of the 42nd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2015, Mumbai, India, January 15-17, 2015*, Sriram K. Rajamani and David Walker (Eds.). ACM, 637–650. doi:10.1145/2676726.2676980
- Jan-Oliver Kaiser, Hoang-Hai Dang, Derek Dreyer, Ori Lahav, and Viktor Vafeiadis. 2017. Strong Logic for Weak Memory: Reasoning About Release-Acquire Consistency in Iris. In *31st European Conference on Object-Oriented Programming, ECOOP 2017, June 19-23, 2017, Barcelona, Spain (LIPIcs, Vol. 74)*, Peter Müller (Ed.). Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 17:1–17:29. doi:10.4230/LIPIcs.ECOOP.2017.17
- Jeehoon Kang, Chung-Kil Hur, Ori Lahav, Viktor Vafeiadis, and Derek Dreyer. 2017. A promising semantics for relaxed-memory concurrency. In *Proceedings of the 44th ACM SIGPLAN Symposium on Principles of Programming Languages, POPL 2017, Paris, France, January 18-20, 2017*, Giuseppe Castagna and Andrew D. Gordon (Eds.). ACM, 175–189. doi:10.1145/3009837.3009850
- Prince Kohli, Gil Neiger, and Mustaque Ahamad. 1993. A Characterization of Scalable Shared Memories. In *Proceedings of the 1993 International Conference on Parallel Processing, Syracuse University, NY, USA, August 16-20, 1993. Volume I: Architecture*, C. Y. Roger Chen and P. Bruce Berra (Eds.). CRC Press, 332–335. doi:10.1109/ICPP.1993.15
- Ori Lahav, Brijesh Dongol, and Heike Wehrheim. 2023. Rely-Guarantee Reasoning for Causally Consistent Shared Memory. In *Computer Aided Verification - 35th International Conference, CAV 2023, Paris, France, July 17-22, 2023, Proceedings, Part I (Lecture Notes in Computer Science, Vol. 13964)*, Constantin Enea and Akash Lal (Eds.). Springer, 206–229. doi:10.1007/978-3-031-37706-8_11
- Ori Lahav and Viktor Vafeiadis. 2015. Owicki-Gries Reasoning for Weak Memory Models. In *Automata, Languages, and Programming - 42nd International Colloquium, ICALP 2015, Kyoto, Japan, July 6-10, 2015, Proceedings, Part II (Lecture Notes in Computer Science, Vol. 9135)*, Magnús M. Halldórsson, Kazuo Iwama, Naoki Kobayashi, and Bettina Speckmann (Eds.). Springer, 311–323. doi:10.1007/978-3-662-47666-6_25
- Ori Lahav, Viktor Vafeiadis, Jeehoon Kang, Chung-Kil Hur, and Derek Dreyer. 2017. Repairing sequential consistency in C/C++11. In *Proceedings of the 38th ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI 2017, Barcelona, Spain, June 18-23, 2017*, Albert Cohen and Martin T. Vechev (Eds.). ACM, 618–632. doi:10.1145/3062341.3062352
- Leslie Lamport. 1977. Proving the Correctness of Multiprocess Programs. *IEEE Trans. Software Eng.* 3, 2 (1977), 125–143. doi:10.1109/TSE.1977.229904
- Sung-Hwan Lee, Minki Cho, Anton Podkopaev, Soham Chakraborty, Chung-Kil Hur, Ori Lahav, and Viktor Vafeiadis. 2020. Promising 2.0: global optimizations in relaxed memory concurrency. In *Proceedings of the 41st ACM SIGPLAN International Conference on Programming Language Design and Implementation, PLDI 2020, London, UK, June 15-20, 2020*, Alastair F. Donaldson and Emina Torlak (Eds.). ACM, 362–376. doi:10.1145/3385412.3386010
- Zongyuan Liu, Sergei Stepanenko, Jean Pichon-Pharabod, Amin Timany, Aslan Askarov, and Lars Birkedal. 2023. VMSL: A Separation Logic for Mechanised Robust Safety of Virtual Machines Communicating above FF-A. In *Proceedings of the 44th ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI 2023, Orlando, Florida, June 17-21, 2023*. ACM, 1438–1462. doi:10.1145/3591279

- Yatin A. Manerkar, Caroline Trippel, Daniel Lustig, Michael Pellauer, and Margaret Martonosi. 2016. Counterexamples and Proof Loophole for the C/C++ to POWER and ARMv7 Trailing-Sync Compiler Mappings. *CoRR* abs/1611.01507 (2016). arXiv:1611.01507 <http://arxiv.org/abs/1611.01507>
- Paul E. McKenney, Ulrich Weigand, Andrea Parri, Boqun Feng, and Alan Stern. 2020. Linux-Kernel Memory Model. ISO/IEC JTC1 SC22 WG21 P0124R7 <https://www.open-std.org/jtc1/sc22/wg21/docs/papers/2020/p0124r7.html>.
- Glen Mével, Jacques-Henri Jourdan, and François Pottier. 2020. Cosmo: a concurrent separation logic for multicore OCaml. *Proc. ACM Program. Lang.* 4, ICFP, Article 96 (aug 2020), 29 pages. doi:10.1145/3408978
- F.L. Morris and C.B. Jones. 1984. An Early Program Proof by Alan Turing. *Annals of the History of Computing* 6, 2 (1984), 139–143. doi:10.1109/MAHC.1984.10017
- Magnus O. Myreen. 2009. *Formal verification of machine-code programs*. Ph. D. Dissertation. University of Cambridge.
- Magnus O. Myreen, Anthony C. J. Fox, and Michael J. C. Gordon. 2007. Hoare Logic for ARM Machine Code. In *Fundamentals of Software Engineering (FSEN)*, Farhad Arbab and Marjan Sirjani (Eds.). Springer, 272–286.
- Magnus O. Myreen and Michael J. C. Gordon. 2007. Hoare Logic for Realistically Modelled Machine Code. In *Tools and Algorithms for the Construction and Analysis of Systems (TACAS)*, Orna Grumberg and Michael Huth (Eds.). Springer, 568–582.
- Magnus O. Myreen, Michael J. C. Gordon, and Konrad Slind. 2008. Machine-Code Verification for Multiple Architectures - An Application of Decompile into Logic. In *Formal Methods in Computer-Aided Design (FMCAD)*, Alessandro Cimatti and Robert B. Jones (Eds.). IEEE, 1–8.
- Peter Naur. 1966. Proofs of algorithms by general snapshots. *BIT* 6, 310–316.
- Peter W. O’Hearn, John C. Reynolds, and Hongseok Yang. 2001. Local Reasoning about Programs that Alter Data Structures. In *Computer Science Logic, 15th International Workshop, CSL 2001. 10th Annual Conference of the EACSL, Paris, France, September 10-13, 2001, Proceedings (Lecture Notes in Computer Science, Vol. 2142)*, Laurent Fribourg (Ed.). Springer, 1–19. doi:10.1007/3-540-44802-0_1
- Scott Owens, Susmit Sarkar, and Peter Sewell. 2009. A Better x86 Memory Model: x86-TSO. In *Theorem Proving in Higher Order Logics, 22nd International Conference, TPHOLS 2009, Munich, Germany, August 17-20, 2009. Proceedings (Lecture Notes in Computer Science, Vol. 5674)*, Stefan Berghofer, Tobias Nipkow, Christian Urban, and Makarius Wenzel (Eds.). Springer, 391–407. doi:10.1007/978-3-642-03359-9_27
- Susan S. Owicki and David Gries. 1976. An Axiomatic Proof Technique for Parallel Programs I. *Acta Informatica* 6 (1976), 319–340. doi:10.1007/BF00268134
- Marco Paviotti, Simon Cooksey, Anouk Paradis, Daniel Wright, Scott Owens, and Mark Batty. 2020. Modular Relaxed Dependencies in Weak Memory Concurrency. In *Programming Languages and Systems - 29th European Symposium on Programming, ESOP 2020, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2020, Dublin, Ireland, April 25-30, 2020, Proceedings (Lecture Notes in Computer Science, Vol. 12075)*, Peter Müller (Ed.). Springer, 599–625. doi:10.1007/978-3-030-44914-8_22
- Jean Pichon-Pharabod and Peter Sewell. 2016. A concurrency semantics for relaxed atomics that permits optimisation and avoids thin-air executions. In *Proceedings of the 43rd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2016, St. Petersburg, FL, USA, January 20 - 22, 2016*, Rastislav Bodík and Rupak Majumdar (Eds.). ACM, 622–633. doi:10.1145/2837614.2837616
- Christopher Pulte. 2018. *The Semantics of Multicopy Atomic ARMv8 and RISC-V*. Ph. D. Dissertation. University of Cambridge, UK. <https://doi.org/10.17863/CAM.39379>
- Christopher Pulte, Shaked Flur, Will Deacon, Jon French, Susmit Sarkar, and Peter Sewell. 2018. Simplifying ARM concurrency: multicopy-atomic axiomatic and operational models for ARMv8. *Proc. ACM Program. Lang.* 2, POPL (2018), 19:1–19:29. doi:10.1145/3158107
- Christopher Pulte, Jean Pichon-Pharabod, Jeehoon Kang, Sung Hwan Lee, and Chung-Kil Hur. 2019. Promising-ARM/RISC-V: a simpler and faster operational concurrency model. In *Proceedings of the 40th ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI 2019, Phoenix, AZ, USA, June 22-26, 2019*, Kathryn S. McKinley and Kathleen Fisher (Eds.). ACM, 1–15. doi:10.1145/3314221.3314624
- Thibaut Pérami, Thomas Bauereiss, Brian Campbell, Zongyuan Liu, Nils Lauermaun, Alasdair Armstrong, and Peter Sewell. 2026. ArchSem: Reusable Rigorous Semantics of Relaxed Architectures. *Proc. ACM Program. Lang.* POPL (2026). doi:10.1145/3776650
- Azalea Raad, Ori Lahav, and Viktor Vafeiadis. 2020. Persistent Owicki-Gries reasoning: a program logic for reasoning about persistent programs on Intel-x86. *Proc. ACM Program. Lang.* 4, OOPSLA (2020), 151:1–151:28. doi:10.1145/3428219
- John C. Reynolds. 2002. Separation Logic: A Logic for Shared Mutable Data Structures. In *17th IEEE Symposium on Logic in Computer Science (LICS 2002), 22-25 July 2002, Copenhagen, Denmark, Proceedings*. IEEE Computer Society, 55–74. doi:10.1109/LICS.2002.1029817
- Tom Ridge. 2010. A Rely-Guarantee Proof System for x86-TSO. In *Verified Software: Theories, Tools, Experiments, Third International Conference, VSTTE 2010, Edinburgh, UK, August 16-19, 2010. Proceedings (Lecture Notes in Computer Science, Vol. 6217)*, Gary T. Leavens, Peter W. O’Hearn, and Sriram K. Rajamani (Eds.). Springer, 55–70. doi:10.1007/978-3-642-15057-9_4
- Michael Sammler, Angus Hammond, Rodolphe Lepigre, Brian Campbell, Jean Pichon-Pharabod, Derek Dreyer, Deepak Garg, and Peter Sewell. 2022. Islaris: verification of machine code against authoritative ISA semantics. In *PLDI ’22: 43rd ACM SIGPLAN International Conference on Programming Language Design and Implementation, San Diego, CA, USA, June 13 - 17, 2022*, Ranjit Jhala and Isil Dillig (Eds.). ACM, 825–840. doi:10.1145/3519939.3523434

- Susmit Sarkar, Kayvan Memarian, Scott Owens, Mark Batty, Peter Sewell, Luc Maranget, Jade Alglave, and Derek Williams. 2012. Synchronising C/C++ and POWER. In *ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI '12, Beijing, China - June 11 - 16, 2012*, Jan Vitek, Haibo Lin, and Frank Tip (Eds.). ACM, 311–322. doi:10.1145/2254064.2254102
- Susmit Sarkar, Peter Sewell, Jade Alglave, Luc Maranget, and Derek Williams. 2011. Understanding POWER multiprocessors. In *Proceedings of the 32nd ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI 2011, San Jose, CA, USA, June 4-8, 2011*, Mary W. Hall and David A. Padua (Eds.). ACM, 175–186. doi:10.1145/1993498.1993520
- Susmit Sarkar, Peter Sewell, Francesco Zappa Nardelli, Scott Owens, Tom Ridge, Thomas Braibant, Magnus O. Myreen, and Jade Alglave. 2009. The semantics of x86-CC multiprocessor machine code. In *Proceedings of the 36th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2009, Savannah, GA, USA, January 21-23, 2009*, Zhong Shao and Benjamin C. Pierce (Eds.). ACM, 379–391. doi:10.1145/1480881.1480929
- Filip Sieczkowski, Kasper Svendsen, Lars Birkedal, and Jean Pichon-Pharabod. 2015. A Separation Logic for Fictional Sequential Consistency. In *Programming Languages and Systems - 24th European Symposium on Programming, ESOP 2015, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2015, London, UK, April 11-18, 2015. Proceedings (Lecture Notes in Computer Science, Vol. 9032)*, Jan Vitek (Ed.). Springer, 736–761. doi:10.1007/978-3-662-46669-8_30
- Ben Simner, Alasdair Armstrong, Thomas Bauereiss, Brian Campbell, Ohad Kammar, Jean Pichon-Pharabod, and Peter Sewell. 2025. Precise exceptions in relaxed architectures. In *Proceedings of the 52nd Annual International Symposium on Computer Architecture, ISCA 2025, Tokyo, Japan, June 21-25, 2025*. ACM, 211–224. doi:10.1145/3695053.3731102
- Ben Simner, Alasdair Armstrong, Jean Pichon-Pharabod, Christopher Pulte, Richard Grisenthwaite, and Peter Sewell. 2022. Relaxed virtual memory in Armv8-A. In *Programming Languages and Systems - 31st European Symposium on Programming, ESOP 2022, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2022, Munich, Germany, April 2-7, 2022. Proceedings (Lecture Notes in Computer Science, Vol. 13240)*, Ilya Sergey (Ed.). Springer, 143–173. doi:10.1007/978-3-030-99336-8_6
- Ben Simner, Shaked Flur, Christopher Pulte, Alasdair Armstrong, Jean Pichon-Pharabod, Luc Maranget, and Peter Sewell. 2020. ARMv8-A System Semantics: Instruction Fetch in Relaxed Architectures. In *Programming Languages and Systems - 29th European Symposium on Programming, ESOP 2020, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2020, Dublin, Ireland, April 25-30, 2020. Proceedings (Lecture Notes in Computer Science, Vol. 12075)*, Peter Müller (Ed.). Springer, 626–655. doi:10.1007/978-3-030-44914-8_23
- Simon Spies, Lennard Gäher, Daniel Gratzer, Joseph Tassarotti, Robbert Krebbers, Derek Dreyer, and Lars Birkedal. 2021. Transfinite Iris: resolving an existential dilemma of step-indexed separation logic. In *PLDI '21: 42nd ACM SIGPLAN International Conference on Programming Language Design and Implementation, Virtual Event, Canada, June 20-25, 2021*, Stephen N. Freund and Eran Yahav (Eds.). ACM, 80–95. doi:10.1145/3453483.3454031
- Léo Stefanescu, Azalea Raad, and Viktor Vafeiadis. 2024. Specifying and Verifying Persistent Libraries. In *Programming Languages and Systems - 33rd European Symposium on Programming, ESOP 2024, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2024, Luxembourg City, Luxembourg, April 6-11, 2024. Proceedings, Part II (Lecture Notes in Computer Science, Vol. 14577)*, Stephanie Weirich (Ed.). Springer, 185–211. doi:10.1007/978-3-031-57267-8_8
- Alexander J. Summers and Peter Müller. 2018. Automating Deductive Verification for Weak-Memory Programs. In *Tools and Algorithms for the Construction and Analysis of Systems - 24th International Conference, TACAS 2018, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2018, Thessaloniki, Greece, April 14-20, 2018. Proceedings, Part I (Lecture Notes in Computer Science, Vol. 10805)*, Dirk Beyer and Marieke Huisman (Eds.). Springer, 190–209. doi:10.1007/978-3-319-89960-2_11
- Kasper Svendsen and Lars Birkedal. 2014. Impredicative Concurrent Abstract Predicates. In *Programming Languages and Systems - 23rd European Symposium on Programming, ESOP 2014, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2014, Grenoble, France, April 5-13, 2014. Proceedings (Lecture Notes in Computer Science, Vol. 8410)*, Zhong Shao (Ed.). Springer, 149–168. doi:10.1007/978-3-642-54833-8_9
- Kasper Svendsen, Jean Pichon-Pharabod, Marko Doko, Ori Lahav, and Viktor Vafeiadis. 2018. A Separation Logic for a Promising Semantics. In *Programming Languages and Systems - 27th European Symposium on Programming, ESOP 2018, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2018, Thessaloniki, Greece, April 14-20, 2018. Proceedings (Lecture Notes in Computer Science, Vol. 10801)*, Amal Ahmed (Ed.). Springer, 357–384. doi:10.1007/978-3-319-89884-1_13
- Runzhou Tao, Jianan Yao, Xupeng Li, Shih-Wei Li, Jason Nieh, and Ronghui Gu. 2021. Formal Verification of a Multiprocessor Hypervisor on Arm Relaxed Memory Hardware. In *SOSP '21: ACM SIGOPS 28th Symposium on Operating Systems Principles, Virtual Event / Koblenz, Germany, October 26-29, 2021*, Robbert van Renesse and Nikolai Zeldovich (Eds.). ACM, 866–881. doi:10.1145/3477132.3483560
- Amin Timany, Simon Oddershede Gregersen, Léo Stefanescu, Jonas Kastberg Hinrichsen, Léon Gondelman, Abel Nieto, and Lars Birkedal. 2024. Trillium: Higher-Order Concurrent and Distributed Separation Logic for Intensional Refinement. *Proc. ACM Program. Lang.* 8, POPL (2024), 241–272. doi:10.1145/3632851
- Alan M. Turing. 1949. Checking a large routine. In *Report of a Conference on High Speed Automatic Calculating Machines*. Mathematical Laboratory, Cambridge, UK, 67–69. <https://turingarchive.kings.cam.ac.uk/publications-lectures-and-talks-ambt/amt-b-8>

- Aaron Turon, Viktor Vafeiadis, and Derek Dreyer. 2014. GPS: navigating weak memory with ghosts, protocols, and separation. In *Proceedings of the 2014 ACM International Conference on Object Oriented Programming Systems Languages & Applications, OOPSLA 2014, part of SPLASH 2014, Portland, OR, USA, October 20-24, 2014*, Andrew P. Black and Todd D. Millstein (Eds.). ACM, 691–707. doi:10.1145/2660193.2660243
- Viktor Vafeiadis and Chinmay Narayan. 2013. Relaxed separation logic: a program logic for C11 concurrency. In *Proceedings of the 2013 ACM SIGPLAN International Conference on Object Oriented Programming Systems Languages & Applications, OOPSLA 2013, part of SPLASH 2013, Indianapolis, IN, USA, October 26-31, 2013*, Antony L. Hosking, Patrick Th. Eugster, and Cristina V. Lopes (Eds.). ACM, 867–884. doi:10.1145/2509136.2509532
- Simon Friis Vindum and Lars Birkedal. 2023. Spirea: A Mechanized Concurrent Separation Logic for Weak Persistent Memory. *Proc. ACM Program. Lang.* 7, OOPSLA2 (2023), 632–657. doi:10.1145/3622820
- Andrew Waterman and Krste Asanović. 2019. *The RISC-V Instruction Set Manual, Volume I: Unprivileged ISA, Document Version 20191213*. Technical Report. <https://github.com/riscv/riscv-isa-manual/releases/download/Ratified-IMAFDQC/riscv-spec-20191213.pdf>
- Ian Wehrman. 2012. *Weak-Memory Local Reasoning (Dissertation draft)*. Ph.D. Dissertation. University of Texas at Austin.
- Ian Wehrman and Josh Berdine. 2011. A proposal for weak-memory local reasoning. In *Low-level languages and applications (LOLA)*.
- John Wickerson, Mike Dodds, and Matthew J. Parkinson. 2013. Ribbon Proofs for Separation Logic. In *Programming Languages and Systems - 22nd European Symposium on Programming, ESOP 2013, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2013, Rome, Italy, March 16-24, 2013. Proceedings (Lecture Notes in Computer Science, Vol. 7792)*, Matthias Felleisen and Philippa Gardner (Eds.). Springer, 189–208. doi:10.1007/978-3-642-37036-6_12
- Daniel Wright, Mark Batty, and Brijesh Dongol. 2021. Owicki-Gries Reasoning for C11 Programs with Relaxed Dependencies. In *Formal Methods - 24th International Symposium, FM 2021, Virtual Event, November 20-26, 2021, Proceedings (Lecture Notes in Computer Science, Vol. 13047)*, Marieke Huisman, Corina S. Pasareanu, and Naijun Zhan (Eds.). Springer, 237–254. doi:10.1007/978-3-030-90870-6_13
- Daniel Wright, Sadegh Dalvandi, Mark Batty, and Brijesh Dongol. 2023. Mechanised Operational Reasoning for C11 Programs with Relaxed Dependencies. *Formal Aspects Comput.* 35, 2 (2023), 10:1–10:27. doi:10.1145/3580285
- Li-yao Xia, Yannick Zakowski, Paul He, Chung-Kil Hur, Gregory Malecha, Benjamin C. Pierce, and Steve Zdancewic. 2020. Interaction trees: representing recursive and impure programs in Coq. *Proc. ACM Program. Lang.* 4, POPL (2020), 51:1–51:32. doi:10.1145/3371119
- Yang Zhang and Xinyu Feng. 2014. *Program Logic for Local Reasoning in TSO*. Technical Report.