

Automated Conversion of Axiomatic to Operational Models: Theory and Practice

Adwait Godbole^{*} , Yatin A. Manerkar[†], and Sanjit A. Seshia[†] 

^{*}[†]University of California Berkeley, Berkeley, USA [†]University of Michigan, Ann Arbor, USA

Abstract—A system may be modelled as an *operational model* (which has explicit notions of state and transitions between states) or an *axiomatic model* (which is specified entirely as a set of invariants). Most formal methods (e.g., IC3, invariant synthesis, etc) are designed for operational models and are largely inaccessible to axiomatic models. Furthermore, no prior method exists to automatically convert axiomatic models to operational ones, so operational equivalents to axiomatic models had to be manually created and proven equivalent.

In this paper, we advance the state-of-the-art in axiomatic to operational model conversion. We show that general axioms in the μspec axiomatic modelling framework cannot be translated to equivalent finite-state operational models. We also derive restrictions on the space of μspec axioms that enable the feasible generation of equivalent finite-state operational models for them. As for practical results, we develop a methodology for automatically translating μspec axioms to equivalent finite-state automata-based operational models. We demonstrate the efficacy of our method by using the models generated by our procedure to prove the correctness of ordering properties on three register-transfer-level (RTL) designs.

I. INTRODUCTION

When modelling hardware or software systems using formal methods, one traditionally uses *operational* models (e.g. Kripke structures [1]), which have explicit notions of state and transitions. However, one may also model a system *axiomatically*, where instead of a state-transition relation, the system is specified entirely by a set of axioms (e.g., invariants) that it maintains. Executions that obey the axioms are allowed, and those that violate one or more axioms are forbidden. The vast majority of formal methods works use the operational modelling style. However, axiomatic models have been used to great effect in certain domains such as memory models, where they have shown order-of-magnitude improvements in verification performance over equivalent operational models [2].

Operational and axiomatic models each have their own advantages and disadvantages [3]. Operational models can be more intuitive as they typically resemble the system that they are modelling. Hence one is not required to reason about invariants to write the model. On the other hand, axiomatic models tend to be more concise and potentially offer faster verification [2].

Many formal methods (e.g., refinement procedures [4], invariant synthesis, IC3/PDR [5], [6]) are set up to use operational models. Axiomatic models are largely or completely incompatible with these techniques, as the axioms constrain full traces rather than a step of the transition relation. One way to take advantage of these techniques when using axiomatic

models is to create and use operational models equivalent to the axiomatic models. The only prior method of doing this was to first manually create the operational model and then manually prove it equivalent to the axiomatic model. There have been several works doing so [2], [7], [8], [9], [10].

Manually creating an operational model and proving equivalence is cumbersome and error-prone. The ability to automatically generate operational models equivalent to a given axiomatic model would be beneficial, eliminating both the time spent creating the operational model as well as the need for tedious manual equivalence proofs. Generated models can then be fed into techniques currently requiring operational models (e.g. IC3/PDR).

To this end, we make advances in this paper towards the automatic conversion of axiomatic models to equivalent operational models, on both theoretical and practical fronts. In our work, we focus specifically on μspec [11], a well-known axiomatic framework for modelling microarchitectural orderings, which has been used in a wide range of contexts [12], [13], [14], [15], [16] including memory consistency, cache coherence and hardware security.

On the theoretical front, we show that it is impossible to convert general μspec axioms to equivalent finite-state operational models. However, we show that it is feasible to generate equivalent operational models for a specific subset of μspec (henceforth referred to as μspecRE). On the practical side, we develop a method to automatically translate universal axioms¹ in μspecRE into equivalent finite-state operational models comprised of building blocks we term as *axiom automata* (finite automata that monitor whether an axiom has been violated). Furthermore, for arbitrary μspec axioms, our method can generate operational models that are equivalent to the axioms up to a program-size bound.

To evaluate our technique, we convert axioms for three RTL designs to their corresponding operational models: an in-order multicore processor (`multi_vsacle`), a memory-controller (`sdram_ctrl`), and an out-of-order single-core processor (`tomasulo`). We showcase how the generated models can be used with procedures like BMC and IC3/PDR which are usually inaccessible for axiomatic models, and we produce both bounded and unbounded proofs of correctness.

Overall, the contributions of this work are as follows:

- We prove that generation of equivalent finite-state operational models for arbitrary μspec axioms is impossible.

¹Axioms that do not contain \exists quantifiers.

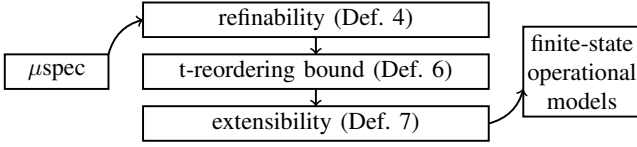


Fig. 1: Roadmap to obtain finite-state operational models.

- We provide a procedure for generating equivalent finite-state operational models for universal axioms in μspecRE .
- We propose the *axiom-automata* formulation to generate equivalent finite-state operational models from universal axioms in μspecRE (or from arbitrary μspec axioms if only guaranteeing equivalence up to a bounded program size).
- We evaluate our method for operational model generation by using our generated models to prove the (bounded/unbounded) correctness of ordering properties on three RTL designs: `multi_vscale`, `tomasulo`, and `sdram_ctrl`.

Generality. While axiomatic models enforce constraints over complete executions, operational models do this local to each transition. Ensuring that behaviours generated by the latter are also allowed by the former requires performing non-local consistency checks which are hard to reason about, especially for unbounded executions. This has been observed in manual operationalization works as well. Taking the example of [7], (which operationalizes C11), we address issues of eliminating consistent executions too early [7, §3] and repeatedly checking consistency [7, §4] by developing concepts such as *t-reordering boundedness* (Def. 6) and *extensibility* (Def. 7). Though we focus on μspec , we believe many of the underlying challenges and concepts carry over to frameworks such as Cat [2].

Outline. §II covers the syntax and semantics of μspec used in this paper. §III covers the formulation of the space of operational models we consider. They have finite control-state and read-only input tapes for the instruction streams (programs) executed by each core. §IV defines our notions of soundness, completeness, and equivalence when comparing operational and axiomatic models. In §V, we show that it is impossible to synthesise equivalent finite-state operational models from arbitrary axiomatic models. We develop an underapproximation, called *t-reordering boundedness*, that addresses this by bounding the depth of reorderings possible. In §VI we restrict μspec further by requiring *extensibility* (preventing current events from influencing orderings between previous events). Restricting μspec by *t-reordering boundedness* and *extensibility* is sufficient to enable the automatic generation of equivalent finite-state operational models (Thm. 2). §VII describes our conversion procedure based on axiom automata. §VIII evaluates our technique by using it to generate operational models, which are then used for checking properties of RTL designs. §IX covers related work, and §X concludes, with §XI suggesting avenues for future work. This paper is accompanied by an extended version which contains supplementary material and proofs [17].

II. μSPEC SYNTAX AND SEMANTICS

A. μspec Syntax

$$\begin{aligned}
 \langle AX \rangle &:= \forall i \ AX \mid \exists i \ AX \mid \phi(i_1, \dots, i_m) \\
 \langle \phi \rangle &:= \phi \wedge \phi \mid \phi \vee \phi \mid \neg \phi \mid \langle \text{atom} \rangle \\
 \langle \text{atom} \rangle &:= i_1 <_r i_2 \mid \text{hb}(i_1.\text{st}, i_2.\text{st}) \mid P(i_1, \dots) \\
 \langle \text{st} \rangle &:= \text{Fet} \mid \text{Dec} \mid \text{Exe} \mid \text{WB} \mid \dots
 \end{aligned}$$

Fig. 2: μspec Syntax.

μspec [11] is a domain-specific language used for specifying microarchitectural orderings. A μspec model consists of *axioms* that enforce first-order constraints over execution graphs; each axiom quantifies over instructions and is required to be a sentence (i.e. not have any free variables). Execution graphs that satisfy the axioms and are acyclic are deemed as valid executions. While ISA-level models [2], [18], [19] treat single instructions as atomic entities, μspec decomposes the execution of an instruction into a set of atomic *events*. Each instruction i and stage st is associated with an event $i.\text{st}$. A program execution is viewed as a directed acyclic graph called a micro-architectural happens-before graph (μhb graph) [12]. Such a graph for a given program has nodes corresponding to events of form $i.\text{st}$ for each instruction i in the program and each stage st prescribed by the model. Edges in the graph correspond to the *happens-before* (**hb**) relation: $\text{hb}(e_1, e_2)$ says that e_1 happened before e_2 . Thus, a cyclic μhb graph corresponds to an impossible scenario where an event happens before itself, and thus represents an execution that cannot occur on the microarchitecture.

Fig. 2 specifies μspec syntax. It has three types of atoms:

- (i) $\text{hb}(i_1.\text{st}, i_2.\text{st})$: happens-before predicate
- (ii) $i_1 <_r i_2$: the reference order (typically the program order)
- (iii) $P(i_1, \dots)$: instruction predicate atoms

Atoms of type (ii) capture the order in which instructions appear in a given program thread. Atoms of type (iii) are predicates over instructions which capture instruction properties, e.g. opcode, source/destination registers. We note that μspec models in literature [11] also make use of the `NodeExists` predicate which identifies event nodes that occur in the execution. We do not model `NodeExists` in this paper, but our approach can be augmented to incorporate it (see [17]).

We identify two types of axioms of interest: *Universal axioms* are of the form: $\forall i_1 \dots \forall i_k \phi(i_1, \dots, i_k)$, and represent constraints applied symmetrically over all tuples of instructions in a program. *Predicate-free axioms* are axioms that do not have occurrences of predicate (P) atoms. We extend these terms to an axiomatic semantics if all axioms are of that type. In this work, our theoretical treatment focuses on universal semantics. Practically though, some underlying ideas carry over to arbitrary axioms as we discuss in §VII, §VIII.

```

ax0:  $\forall i_1. \text{hb}(i_1.\text{Exe}, i_1.\text{Com})$ 
ax1:  $\forall i_1, i_2. (i_1 <_r i_2 \wedge \text{DepOn}(i_1, i_2)) \implies \text{hb}(i_1.\text{Exe}, i_2.\text{Exe})$ 
ax2:  $\forall i_1, i_2. \text{SameCore}(i_1, i_2) \implies (\text{hb}(i_1.\text{Exe}, i_2.\text{Exe}) \vee \text{hb}(i_2.\text{Exe}, i_1.\text{Exe}))$ 
ax3:  $\forall i_1, i_2. i_1 <_r i_2 \implies \text{hb}(i_1.\text{Com}, i_2.\text{Com})$ 

```

Fig. 3: An example axiomatic model.

B. Illustrative μspec Example

Consider the four axioms in Fig. 3. In the axioms, i_1, i_2 are instruction variables and Exe, Com are stage names (short for execute and commit respectively). The axiom **ax0** requires that for each instruction, the execute stage (Exe) of that instruction must happen before the commit stage (Com). Intuitively, **ax1** says that when i_2 depends on i_1 (captured by the predicate DepOn), i_1 should be executed before i_2 ; **ax2** says that the execute events of instructions on the same core should be *totally ordered* by hb . The third axiom **ax3** says that when i_1 and i_2 are in program order (denoted by $<_r$), i_1 must be committed before i_2 .

Fig. 4 shows valid and invalid execution graphs for the program snippet in Fig. 5. The snippet is of a 2-core program, with two instructions per core. Instruction i_1 is dependent on the result of i_0 (since its source register is the same as the destination of i_0). In the example axiomatic semantics, **ax1** requires that the execute event of instruction i_0 be before that of i_1 . The execution in Fig. 4b is invalid w.r.t. **ax1** since $i_1.\text{Exe}$ is executed before $i_0.\text{Exe}$. The execution in Fig. 4a is valid even though the $i_2.\text{Exe}$ and $i_3.\text{Exe}$ events are reordered since i_3 does not depend on i_2 . Both executions are valid w.r.t. **ax0, ax2** and **ax3**.

C. Programming Model

We consider multi-core systems with each core executing a straight-line program over a finite domain of operations. This is common in memory models [2], [12], [16], [20] and distributed systems [21] literature.

1) *Cores*: The system consists of n processor cores: $\text{Cores} = [n]$. Each core executes operations from a finite set \mathbb{O} . The axiomatic model \mathcal{A} assigns predicates from \mathbb{P} an interpretation over the universe \mathbb{O} . We denote this interpretation as $\mathbb{P}^{\mathcal{A}} \subseteq \mathbb{O}^k$ for an arity- k predicate.

2) *Instruction streams*: An *instruction stream* \mathcal{I} is a word over \mathbb{O} : $\mathcal{I} \in \mathbb{O}^*$. A program \mathcal{P} is a set of *per-core* instruction streams: $\{\mathcal{I}_c\}_{c \in \text{Cores}}$. For a core c and label $0 \leq j < |\mathcal{I}_c|$, we call the triple $(c, j, \mathcal{I}_c[j])$ an *instruction*². We denote components of instruction $i = (c, j, \mathcal{I}_c[j])$, as: $c(i) = c$, label $\lambda(i) = j$ and operation $\text{op}(i) = \mathcal{I}_c[j]$. The set of instructions occurring in \mathcal{P} is: $\text{instrsOf}(\mathcal{P}) = \{(c, j, \mathcal{I}_c[j]) \mid c \in$

²Note the terminology: operations are commands that the core can execute. Since we interpret predicates over \mathbb{O} we require $|\mathbb{O}|$ to be a finite set for computability reasons. Instructions are operations combined with the label and core identifier (and hence form an infinite set).

$\text{Cores}, 0 \leq j < |\mathcal{I}_c|$) and the set of all possible instructions as $\mathbb{I} = \text{Cores} \times \mathbb{Z}^{\geq 0} \times \mathbb{O}$.

3) *Instruction stages*: Instruction execution in μspec is decomposed into stages. The set of stages, Stages , is a parameter of the semantics. Instruction i performing in stage st , (i.e. $i.st$) is an atomic event in an execution. The execution of \mathcal{P} is composed of the set of events: $\text{eventsOf}(\mathcal{P}) = \{i.st \mid i \in \text{instrsOf}(\mathcal{P}), st \in \text{Stages}\}$. The set of all possible events is $\mathbb{E} = \{i.st \mid i \in \mathbb{I}, st \in \text{Stages}\}$.

Definition 1 (Event). *An event e is of the form $i.st$. It represents the instruction $i \in \mathbb{I}$, (atomically) performing in stage $st \in \text{Stages}$.*

Example 1. *Following the example in Fig. 5 we consider an architecture with two opcodes: **add, lw** for add and load respectively. For each of these, we may have several actual operations (with different operands), thus giving us the set \mathbb{O} . The program \mathcal{P} in Fig. 5 has two cores: $\text{Cores} = \{c_0, c_1\}$ and four instructions: $\text{instrsOf}(\mathcal{P}) = \{i_0, i_1, i_2, i_3\}$. We have, for example, $c(i_1) = c_0, \lambda(i_1) = 1, \text{op}(i_1) = \text{add } r3, r2, r1$ while $c(i_2) = c_1, \lambda(i_2) = 0$. The instruction stream for core c_0 is $\mathcal{I}_0 = i_0 \cdot i_1$ and that of core c_1 is $\mathcal{I}_1 = i_2 \cdot i_3$.*

Let us suppose that this program is executed on a 4-stage microarchitecture with $\text{Stages} = \{\text{Fet}, \text{Dec}, \text{Exe}, \text{WB}, \text{Com}\}$. The events corresponding to the program are given by $\text{eventsOf}(\mathcal{P}) = \{i_0.\text{Fet}, i_0.\text{Dec}, \dots, i_3.\text{Com}\}$ with $|\text{eventsOf}(\mathcal{P})| = 4 \times 5 = 20$.

D. Formal μspec Semantics

We now define the formal semantics of μspec axioms.

Definition 2 (μhb graph). *For a program \mathcal{P} , a μhb graph is a directed acyclic graph, $G(V, E)$, with nodes $V = \text{eventsOf}(\mathcal{P})$ representing events and edges representing the happens-before relationships, i.e. $(e_1, e_2) \in E \equiv \text{hb}(e_1, e_2)$.*

Validity of μhb graph w.r.t. an axiomatic semantics: Consider an axiomatic semantics \mathcal{A} (i.e. a set of axioms). A μhb graph $G = (V, E)$ is said to represent a valid execution of program \mathcal{P} under \mathcal{A} if it satisfies all the axioms in \mathcal{A} . We denote the validity of a μhb graph G by $G \models_{\mathcal{P}} \mathcal{A}$.

Satisfaction w.r.t. an axiom: We first define satisfaction for the quantifier-free part, starting at the atoms. Let $s : \mathbb{I}(\text{AX}) \rightarrow \mathbb{I}$ be an assignment for the symbolic instruction variables $\mathbb{I}(\text{AX})$ in axiom AX .

$$G \models i_1[s] <_r i_2[s] \iff c(s(i_1)) = c(s(i_2)) \wedge \lambda(s(i_1)) < \lambda(s(i_2)) \quad \dots(i)$$

$$G \models P(i_1, \dots, i_m)[s] \iff (\text{op}(s(i_1)), \dots, \text{op}(s(i_m))) \in \mathbb{P}^{\mathcal{A}} \quad \dots(ii)$$

$$G \models \text{hb}(i_1.st_1, i_2.st_2)[s] \iff (s(i_1).st_1, s(i_2).st_2) \in E^+ \quad \dots(iii)$$

In (i), the reference order $<_r$ relates instructions i_1, i_2 from the same instruction stream if i_1 is before i_2 . In (ii) we extend predicate interpretations, $\mathbb{P}^{\mathcal{A}}$, (defined over \mathbb{O}) to instructions by taking the $\text{op}(\cdot)$ component. Finally, hb atoms




Fig. 4: Valid (a) and an invalid (b) execution graphs for the program in Fig. 5 and axioms in Fig. 3. All edges represent the hb relation. The red (bold) edge violates $ax1$.

```

i0: lw r1, 42(r0)   |   i2: lw r4, 42(r0)
i1: add r3, r2, r1 |   i3: add r3, r2, r1

```

Fig. 5: Example program snippet

are interpreted as E^+ , i.e. transitive closure of E , as stated in (iii). Operators \wedge, \vee, \neg have their usual semantics.

We now define the satisfaction of a (quantified) axiom AX by a graph G , denoted by $G \models_{\mathcal{P}} AX$ above.

$$\begin{aligned}
G \models_{\mathcal{P}} \phi[s] &\equiv G \models \phi[s] \\
&\quad \text{for quantifier-free } \phi \\
G \models_{\mathcal{P}} \forall i \phi[s] &\equiv G \models_{\mathcal{P}} \phi[s[i \leftarrow i]] \\
&\quad \text{for all } i \in \text{instrsOf}(\mathcal{P}) \setminus \text{range}(s) \\
G \models_{\mathcal{P}} \exists i \phi[s] &\equiv G \models_{\mathcal{P}} \phi[s[i \leftarrow i]] \\
&\quad \text{for some } i \in \text{instrsOf}(\mathcal{P}) \setminus \text{range}(s)
\end{aligned}$$

The base case is $G \models_{\mathcal{P}} \phi[s]$ (where ϕ is quantifier-free) and follows the earlier definitions. We extend $G \models_{\mathcal{P}} \phi$ with (almost) usual quantification semantics: $\forall (\exists)$ quantifies over all (some) instructions in $\text{instrsOf}(\mathcal{P})$. Execution G is a valid execution of \mathcal{P} under semantics \mathcal{A} , denoted as $G \models_{\mathcal{P}} \mathcal{A}$, if $G \models_{\mathcal{P}} AX$ for all axioms AX in \mathcal{A} .

III. OPERATIONAL MODEL OF COMPUTATION

To concretize our claims, we introduce a model of computation that characterizes the models of interest. We choose to focus on finite-state operational models that generate totally ordered traces, where transitions represent (i.st) events. While there are less restrictive models (e.g. event structures [22], [23]), such models require specialized, typically under-approximate, verification techniques (e.g. [24], [25], [26]). Our choice is motivated by the ability to (a) have finite-state implementations of generated models (e.g. in RTL) and (b) verify against these models with off-the-shelf tools (e.g. model checkers using BDD and SMT-based backends).

A. Model of computation

Intuitively, the model of computation resembles a 1-way transducer [27], [28] with multiple (read-only) input tapes (one tape for each instruction stream). This allows us to execute programs of unbounded length with a finite control state.³

1) *Model definition*: An operational model is parameterized by cores Cores , stages Stages , and a history parameter $h \in \mathbb{N} \cup \{\infty\}$ which bounds the length of tape to the left of the head. It is a tuple $(\mathcal{Q}, \Delta, q_{\text{init}}, q_{\text{final}})$:

- \mathcal{Q} is a finite set of control states

³A Kripke structure-based formalism is insufficient since we want to execute unbounded programs with distinguished instructions without explicitly modelling control logic.

- $\Delta \subseteq \mathcal{Q} \times (\mathbb{I} \cup \{\neg\})^{|\text{Cores}|} \times \mathcal{Q} \times \text{Act}$ is the transition relation where Act is the set of actions
- $q_{\text{init}} \in \mathcal{Q}$ is the initial state
- $q_{\text{final}} \in \mathcal{Q}$ is the final state which must be absorbing (i.e. it has a self-loop)

A model is finite-state if \mathcal{Q} is finite, and it has bounded-history if $h \in \mathbb{N}$. For the end goal of effective verification, we are interested in finite-state, bounded-history models since it is precisely such models that can be compiled to finite-state systems.

2) *Model semantics*: A configuration is a triple $\gamma = (U, q, V)$ where $U : \text{Cores} \rightarrow \mathbb{I}^*$, $V : \text{Cores} \rightarrow \mathbb{I}^*$ and $q \in \mathcal{Q}$. Intuitively U (V) represent, for each instruction stream, the contents of the input tape to the left (right) of the head respectively. For a bounded history machine, a configuration is *allowed* only if $|U(c)| \leq h$ for all $c \in \text{Cores}$. For unbounded history all configurations are allowed.

The set of actions is

$$\begin{aligned}
\text{Act} = & \{\text{right}(c) \mid c \in \text{Cores}\} \cup \\
& \{\text{stay}\} \cup \\
& \{\text{sched}(c, i, \text{st}) \mid c \in \text{Cores}, \text{st} \in \text{Stages}, i \in [h]\} \cup \\
& \{\text{drop}(c, i) \mid c \in \text{Cores}, i \in [h]\}
\end{aligned}$$

Intuitively, these represent in order: motion of the tape head for c to the right, silent (no-effect), generation of an event, and removing the i^{th} instruction from the left of the head. We provide full semantics in the supplementary material [17].

For word $w \in \mathbb{I}^*$, let $\text{fst}(w)$ denote its first element if $w \neq \epsilon$ and \perp otherwise. Transitions are enabled based on the control state and the instructions that the tape-heads point to: transition $(q_1, (i_1, \dots, i_{|\text{Cores}|}), q_2, _) \in \Delta$ is enabled in configuration $\gamma = (U, q, V)$ if $q_1 = q$ and $\text{fst}(V(c)) = i_c$ for each $c \in \text{Cores}$.

3) *Runs*: The initial configuration is given by $\gamma_{\text{init}}(\mathcal{P}) = (U_{\text{init}}, q_{\text{init}}, V_{\text{init}})$ where $U_{\text{init}} = \lambda c. \epsilon$ and $V_{\text{init}} = \lambda c. \mathcal{I}_c$, i.e. for each core, the left of the tape head is empty, and the right of the tape head consists of the instruction stream for that core. Starting from $\gamma_{\text{init}}(\mathcal{P})$, the machine transitions according to the transition rules. Such a sequence of configurations $\gamma_{\text{init}}(\mathcal{P}) = \gamma_0 \xrightarrow{e_1} \gamma_1 \cdots \xrightarrow{e_m} \gamma_m$, where all γ_i are allowed is called a run. A run is called accepting if it ends in the state q_{final} .

4) *Traces*: The sequence of event labels $\sigma = e_1 \cdots e_m$ annotating a run is the *trace* corresponding to the run. Each label is an event from \mathbb{E} and hence $\sigma \in \mathbb{E}^*$. We view σ as a (linear) μhb execution graph $e_1 \xrightarrow{hb} e_2 \cdots \xrightarrow{hb} e_m$, and hence define $\sigma \models \mathcal{A}$ in the usual way. Accordingly, we will sometimes refer to σ as an execution of a program \mathcal{P} . The set

of traces corresponding to accepting runs of an operational model \mathcal{M} on a program \mathcal{P} are denoted as $\text{traces}_{\mathcal{M}}(\mathcal{P}) \subseteq \mathbb{E}^*$.

IV. SOUNDNESS, COMPLETENESS, AND EQUIVALENCE


We proceed to formalize the notion of equivalence that relates axiomatic and operational models. In literature [29], [2], ISA-level behaviours of programs have been annotated by the read values of `load` operations. Hence, one notion of equivalence might be to require that identical read values be possible between the models. While this may be reasonable for ISA-level behaviours, it can hide microarchitectural features: different microarchitectural executions can have identical architectural results. Given that μspec models executions at the granularity of microarchitectural events, we adopt a stronger notion of equivalence. For soundness, we require that the operational semantics generates linearizations of μhb graphs that are valid under the axiomatic semantics. Formally:

Definition 3 (Soundness). *An operational model \mathcal{M} is sound w.r.t. \mathcal{A} if for any program \mathcal{P} , each trace in $\text{traces}_{\mathcal{M}}(\mathcal{P})$ is a linearization of some μhb graph that is valid under \mathcal{A} .*

Before defining completeness, we need to address a subtlety. Since operational executions are viewed as μhb graphs by interpreting trace-ordering as the `hb` ordering, the operational model always generates linearized μhb graphs. However, in general, linearizations of valid μhb graphs could end up being invalid w.r.t the axioms. Consider Example 2.

Example 2 (Non-refinable axiom). For the following axiom with `Stages = {S}`, the graph (a) is a valid execution. However, both of its linearizations (b) and (c) are invalid. Thus, all of the (totally-ordered) traces generated by our operational models will be deemed invalid under the axiomatic semantics. This renders a direct comparison between operational and axiomatic executions infeasible.

$$\forall i_1, i_2. (\neg \text{hb}(i_1.S, i_2.S) \wedge \neg \text{hb}(i_2.S, i_1.S))$$



To address this issue, we develop the notion of refinability. For two μhb graphs $G = (V, E)$ and $G' = (V', E')$, we say that G' refines G , denoted $G \sqsubseteq G'$ if (1) $V = V'$ and (2) $(e_1, e_2) \in E^+ \implies (e_1, e_2) \in E'^+$.

Definition 4 (Refinable `hb`). *An axiomatic semantics \mathcal{A} is refinable if for any program \mathcal{P} , and μhb graph G s.t. $G \models_{\mathcal{P}} \mathcal{A}$, we have $G' \models_{\mathcal{P}} \mathcal{A}$ for all linear graphs G' satisfying $G \sqsubseteq G'$.*

Refinability says that all linearizations of a valid graph are valid. While executions under axiomatic semantics are given by (partially-ordered) μhb graphs, our class of operational models generate totally-ordered traces. Refinability bridges this gap by relating valid μhb graphs to valid traces. Interestingly, we can check whether a universal axiomatic semantics satisfies refinability, which at a high level, we show via a small model property (Lemma 1).

Lemma 1. *Given a universal axiomatic semantics we can decide whether the semantics is refinable.*

Refinability is especially important for completeness. For non-refinable semantics, validity of linearizations cannot be checked based on the axioms, as all linearizations may be invalid (Example 2).

We assume that the axiomatic semantics satisfies refinability.

We define completeness and our formal problem statement.

Definition 5 (Completeness). *An operational model \mathcal{M} is complete, if for any program \mathcal{P} and valid μhb graph $G \models_{\mathcal{P}} \mathcal{A}$, $\text{traces}_{\mathcal{M}}(\mathcal{P})$ contains all linearizations of G .*

Formal Problem Statement Given an axiomatic semantics \mathcal{A} , a set of cores `Cores` and stages `Stages`, generate a *finite state, bounded history* model, $\mathcal{M} = (\mathcal{Q}, \Delta, q_{\text{init}}, q_{\text{final}})$, which satisfies soundness and completeness (Defns. 3 and 5).

V. ENABLING SYNTHESIS BY BOUNDING REORDERINGS

In this section, we develop some theoretical results for the synthesis of operational models. First, we show that synthesis of sound and complete (viz. Defn. 3 and 5) finite-state operational models is not possible. Then we provide an underapproximation for the completeness requirement, called *t-completeness*, that enables the synthesis of finite-state models. This still does not allow for bounded-history models as future events can influence past orderings (Example 3). In §VI we add *extensibility* thus enabling our original goal of finite-state and bounded-history models.

A. An impossibility result

We show that it is in fact impossible to develop a finite-state transition system \mathcal{M} that satisfies the requirements prescribed in Defns. 3 and 5. Figure 6 gives an axiomatic semantics $\mathcal{A}^\#$ (with `Stages = {S, T}`) such that for all possible finite-state models, there is some program such that either soundness or completeness is violated. In words, the axioms in Fig. 6 state

$$\begin{aligned} \mathbf{ax0}: & \forall i_1. \quad \text{hb}(i_1.S, i_1.T) \\ \mathbf{ax1}: & \forall i_1, i_2. \quad \text{hb}(i_1.S, i_2.S) \implies \text{hb}(i_1.T, i_2.T) \end{aligned}$$

Fig. 6: Semantics $\mathcal{A}^\#$ that does not allow bounded synthesis

the following constraints: **ax0** says that for each instruction, the `S` stage event happens before the `T` stage, and **ax1** enforces that for any two instructions, the ordering between their `S` stage events implies an identical ordering between their `T` stage events. We have the following:

Theorem 1. *For a single-core program \mathcal{P} with an instruction stream of $|\mathcal{I}_{c_1}| = m$ instructions, there is no model $\mathcal{M} = (\mathcal{Q}, \Delta, q_{\text{init}}, q_{\text{final}})$ that is sound and complete w.r.t. $\mathcal{A}^\#$ and \mathcal{P} , and s.t. $|\mathcal{Q}| < \mathcal{O}(2^m/m)$, even with $h = \infty$.*

We provide an intuitive explanation, deferring details to the supplement [17]. In valid executions of $\mathcal{A}^\#$, `S` stage events

can be ordered arbitrarily, while **T** stage events must maintain the same ordering as that of corresponding **S** stages. Hence the machine must remember the **S** orderings in its finite control. However, the number of such orderings grows (exponentially) with the number of instructions m , implying that existence of a finite-state model that works for all programs is not possible.

Corollary 1. *There does not exist a finite state operational model (even with $h = \infty$) which is sound and complete with respect to the $\mathcal{A}^\#$ axioms.*

B. An underapproximation result

Given the results of the previous section, we must relax some constraint imposed on the operationalization: we choose to relax completeness. To do so, we define an underapproximation called *t-reordering bounded traces*. Intuitively, this imposes two constraints: (a) it bounds the depth of reorderings between instructions on each core, (b) it bounds the number of instructions executed on all other cores, while a core is executing a single instruction.

We observe that (a) is a reasonable assumption since most microarchitectures bound reordering depth, often due to finite reorder buffers. On the other hand, (b) can be thought of as a fairness/starvation-freedom property.

For two instructions i_1, i_2 on the same core, let $\text{diff}_r(i_1, i_2) = \lambda(i_2) - \lambda(i_1)$ (recall that $\lambda(i)$ is the instruction index of i). Consider a trace σ of program \mathcal{P} . For $i \in \text{instrsOf}(\mathcal{P})$, we define the starting index of i , denoted as $\text{start}(i)$, as the index of the first event of instruction i in σ . Similarly we define the ending index, $\text{end}(i)$ as the largest index for some event of i in σ . Let the *prefix-closed end index* of i be the max of end over instructions that are $\leq_r i$: $\text{pfxend}(i) = \max\{\text{end}(i') \mid i' \leq_r i\}$. Two instructions i_1 and i_2 are coupled in a trace (denoted as $\text{coup}(i_1, i_2)$) if the intervals $[\text{start}(i_1), \text{pfxend}(i_1)]$, $[\text{start}(i_2), \text{pfxend}(i_2)]$ overlap.

Definition 6 (*t-reordering bounded traces*). *A trace is t-reordering bounded if, for any pair of instructions i_1, i_2 with $c(i_1) = c(i_2)$, (1) if $i_2.st_2 \xrightarrow{\text{hb}} i_1.st_1$ then $\text{diff}_r(i_1, i_2) < t$ and (2) if $\text{coup}(i_1, i)$, $\text{coup}(i, i_2)$ for some i then $|\text{diff}_r(i_1, i_2)| < t$.*

Intuitively, (1) says that an instruction cannot be reordered with another that precedes it by $\geq t$ indices, while (2) says that instructions on a core cannot be *stalled* while more than t instructions are executed on another. Note that *t-reordering boundedness* is a property of traces, and not of axioms. We now relax completeness (and hence equivalence) to require that the operational model at least generate all *t-reordering bounded linearizations* (instead of all linearizations).

Definition 5* (*t-completeness*). *An operational model \mathcal{M} is t-complete w.r.t. an axiomatic model \mathcal{A} , if for each program \mathcal{P} and $G \models_{\mathcal{P}} \mathcal{A}$, $\text{traces}_{\mathcal{M}}(\mathcal{P})$ contains all t-reordering bounded linearizations of G .*

Replacing Defn. 5 with its *t*-bounded relaxation (Defn. 5*) addresses the issue of having to keep track of an unbounded number of orderings. However, to allow for finite implemen-




Fig. 7: \mathcal{P} has instruction streams $i_0^1 \cdot i_1^1 \cdot i_2^1 \cdot i_3^1$, $i_0^2 \cdot i_1^2 \cdot i_2^2 \cdot i_3^2$, and $i_0^3 \cdot i_1^3 \cdot i_2^3 \cdot i_3^3$. Blue instructions form the prefix \mathcal{P}' (i.e. $\mathcal{P}' \preceq \mathcal{P}$) and red its residual $\mathcal{P}'' = \mathcal{P} \oslash \mathcal{P}'$. The figure shows executions G' of \mathcal{P}' , G'' of \mathcal{P}'' , and their composition $G = G' \triangleright G''$.

tations in practice, in addition to finite-state, we also require bounded-history ($h \in \mathbb{N}$). This is addressed in the next section.

VI. ADDING EXTENSIBILITY

As illustrated by the following example, the *t-reordering bounded underapproximation* is insufficient to achieve bounded-history operational model synthesis on its own.

Example 3 (Need for extensibility). Consider a single stage axiomatic semantics: $\text{Stages} = \{\mathbf{S}\}$, and predicate $\mathbb{P} = \{\mathbf{P}\}$.

$$\forall \mathbf{i0}, \mathbf{i1}, \mathbf{i2}. (\mathbf{P}(\mathbf{i0}, \mathbf{i1}, \mathbf{i2}) \wedge \mathbf{i0} <_r \mathbf{i1}) \implies \neg \text{hb}(\mathbf{i1.S}, \mathbf{i0.S})$$

There cannot be a sound, *t*-complete, and bounded-history (for bound h) model for this axiom (for some $t > 1$). To see this, consider a (single-core) program \mathcal{P} , with instructions $i_0 \cdot i_1 \cdots i_{h+1}$. Depending on the instructions in \mathcal{P} , the interpretation $\mathcal{P}^{\mathcal{A}}$ of \mathcal{P} can either be (a) $\mathcal{P}^{\mathcal{A}} = \{(i_0, i_1, i_{h+1})\}$ or (b) $\mathcal{P}^{\mathcal{A}} = \{\}$. In the former case, the ordering $i_1.S \xrightarrow{\text{hb}} i_0.S$ is invalid while in the latter it is valid. Since we only allow a h -sized history, $i_0.S$ must be scheduled before the tape-head reaches i_{h+1} , i.e. before the machine can determine which of (a)/(b) hold. Since the machine cannot determine whether events $i_0.S$, $i_1.S$ can be reordered, this leads either to a model which is unsound (always reorders) or incomplete (never reorders).

Thus, we need an additional restriction to enable generation of operational models with a finite history parameter h . We propose extensibility, which intuitively states that partial executions of program \mathcal{P} that have not violated any axioms can be composed with valid executions of the residual program to generate valid complete executions of \mathcal{P} . To do this, we extend the notion of validity to partial executions through *prefix programs*.

A program \mathcal{P} can be split into a prefix \mathcal{P}' (blue) and the residual suffix \mathcal{P}'' (red) (Fig. 7). Formally, \mathcal{P}' is a *prefix* of program \mathcal{P} , if \mathcal{P}' has instruction streams $\{\mathcal{I}'_i\}$, each of which is a prefix of the instr. streams $\{\mathcal{I}_i\}$ of \mathcal{P} . We denote that \mathcal{P}' is a prefix of \mathcal{P} by $\mathcal{P}' \preceq \mathcal{P}$. For programs \mathcal{P} , \mathcal{P}' such that $\mathcal{P}' \preceq \mathcal{P}$ we denote the *residual* of \mathcal{P} w.r.t. \mathcal{P}' as $\mathcal{P}'' = \mathcal{P} \oslash \mathcal{P}'$. \mathcal{P}' has instr. streams \mathcal{I}'_c : for each core c , $\mathcal{I}_c = \mathcal{I}'_c \cdot \mathcal{I}''_c$.

In Fig. 7, for example, the first instruction stream of \mathcal{P} is $i_0 \cdot i_1 \cdot i_2 \cdot i_3$. The prefix program \mathcal{P}' has (the prefix) $i_0 \cdot i_1 \cdot i_2$ as its first instr. stream. On the other hand, the residual program, $\mathcal{P}'' = \mathcal{P} \circ \mathcal{P}'$, has the suffix i_3 as its instruction stream.

For graphs $G' = (V', E')$ and $G'' = (V'', E'')$, with $V' \cap V'' = \emptyset$ we define $G' \triangleright G''$ as the graph $G = (V, E)$ where, (1) $V = V' \cup V''$, and (2) $E = E' \cup E'' \cup \{(e', e'') \mid e' \in \text{sink}(E'), e'' \in \text{source}(E'')\}$. The example in Fig. 7 illustrates such a composition: we have $G = G' \triangleright G''$.

Definition 7 (Extensibility). *An axiom AX satisfies extensibility if for any programs \mathcal{P} and \mathcal{P}' s.t. $\mathcal{P}' \preceq \mathcal{P}$, and $\mathcal{P}'' = \mathcal{P} \circ \mathcal{P}'$ if $G' \models_{\mathcal{P}'}$ AX and $G'' \models_{\mathcal{P}''}$ AX then $G' \triangleright G'' \models_{\mathcal{P}}$ AX. An axiomatic semantics \mathcal{A} satisfies extensibility if all axioms $AX \in \mathcal{A}$ satisfy extensibility.*

We require that the axiomatic model satisfies extensibility. We define μspecRE (RE stands for Refinable, Extensible) as the subset of μspec in which all axioms are refinable and extensible. Finite-state, bounded-history synthesis is feasible for universal axioms in μspecRE , as we discuss in the next section. Like refinability, we can check whether an axiom satisfies extensibility (Lemma 2).

Lemma 2. *Given a universal axiom we can decide whether it satisfies extensibility.*

VII. CONVERTING TO OPERATIONAL MODELS USING AXIOM AUTOMATA

In this section, we describe our approach that converts an axiomatic model into an equivalent operational model \mathcal{M} . In §VII-A we develop *axiom automata*, which are the building blocks of our operationalization: they are automata that check for axiom compliance as the operational model executes. In §VII-B we describe how these automata can be instantiated to ensure validity for bounded programs with arbitrary μspec axioms. §VII-C holds our main result: we describe how axiom automata can be instantiated to get a finite-state bounded-history model for universal axioms in μspecRE .

We focus on a single universal axiom $\forall i_1, \dots, i_k \phi$, but this can be easily extended to a set of axioms.

A. Axiom Automata

In what follows, we fix a (universal) axiom $AX = \forall i_1 \dots \forall i_k \phi(i_1, \dots, i_k)$, and let $I(AX) = \{i_1, \dots, i_k\}$, $E(AX) = \{i.st \mid i \in I(AX), st \in \text{Stages}\}$. This axiom enforces that $\phi(\cdot)$ holds for all k -tuples of instructions in the given program. An axiom automaton is a finite state automaton that monitors whether $\phi(\cdot)$ holds for a single k -tuple of instructions. Our operational model is composed of several such automata - thereby allowing us to check all k -tuples. We now define axiom automata, starting with some auxiliary definitions.

Let $\text{nonhb}(AX)$ denote the non-hb atoms in ϕ , i.e. instruction predicate applications and $<_r$ orderings. A *context* is an assignment (of true/false) to each atom in $\text{nonhb}(AX)$; $\text{cxt} : \text{nonhb}(AX) \rightarrow \mathbb{B}$. Each variable assignment $s : I(AX) \rightarrow \mathbb{I}$ fixes the valuation of all $\text{nonhb}(AX)$ atoms (following the

semantics in §II). Hence each assignment s leads to a unique context, which we denote as $\text{cxt}(s)$.

We extend assignments to events and words over events. For $e = i.st$, we define $s(e) = s(i).st$ and for $w \in E(AX)^*$,

$$s(w) = s(w[0]) \cdots s(w[|w| - 1]) \in \mathbb{E}^*$$

As mentioned in §III-A4, we interpret $s(w) \in \mathbb{E}^*$ as the μhb graph $w[0] \xrightarrow{\text{hb}} w[1] \cdots \xrightarrow{\text{hb}} w[|w| - 1]$.

Observe that once we fix the context, the validity of $\phi(\cdot)$ only depends on the value of the hb atoms in ϕ . Hence for two assignments s_1, s_2 with the same context: $\text{cxt}(s_1) = \text{cxt}(s_2)$, s_1 and s_2 share the same set of valid executions: $s_1(w)$ satisfies ϕ if and only if $s_2(w)$ does. This implies that across different assignments s , there are only finitely many valid sets of executions over events in $s(E(AX))$ - one for each context. Intuitively, contexts divide the set of all possible assignments into classes which admit similar orderings.

As a consequence of the above, for each AX and context cxt, we can construct a finite state automaton that recognizes acceptable orderings of $E(AX)$ (Lemma 3). The main observation behind Lemma 3 is that once the context (i.e. interpretation of the nonhb(AX) atoms) is fixed, the allowed orderings can be represented as a language over the symbolic events $E(AX)$.

Lemma 3 (Axiom-Automata). *Given an axiom AX and context cxt, there exists a finite-state automaton $\text{aa}(AX[\text{cxt}])$ over alphabet $E(AX)$ with language $\{w \mid w \in E(AX)^{\text{perm}}, s(w) \models \phi(i_1, \dots, i_k)[s] \text{ for all } s \text{ that agree with cxt}\}$.*

B. Deploying axiom automata

1) *Concretization of an axiom automaton:* The automaton $\text{aa}(AX[\text{cxt}])$ mentioned in Lemma 3 recognizes orderings over the symbolic alphabet $E(AX)$ that lead to ϕ being satisfied. Our end goal, however, is identifying acceptable orderings over the (non-symbolic) events \mathbb{E} . This requires us to generate concrete instances of axiom automata, one for each assignment $s : I(AX) \rightarrow \mathbb{I}$, which we now do.

Given an assignment $s : I(AX) \rightarrow \mathbb{I}$, we denote the (*concretized*) automaton for s w.r.t AX as $\text{aa}(AX, s)$. The automaton $\text{aa}(AX, s)$ is identical to $\text{aa}(AX[\text{cxt}(s)])$, except that the symbolic alphabet $E(AX)$ replaced by its image $s(E(AX))$ under s . Intuitively (by §VII-A), the set of valid orderings of events in $s(E(AX))$ is characterized by the context of s , $\text{cxt}(s)$. This means that the acceptable orderings of events in $s(E(AX))$ is identical to the set of words (orderings) accepted by $\text{aa}(AX[\text{cxt}(s)])$, except that the symbolic events $E(AX)$ should be replaced by their concrete counterparts, $s(E(AX))$. This justifies the definition of $\text{aa}(AX, s)$.

We extend the notation $\text{aa}(AX, s)$ from a single assignment to a set of assignments. For $I \subseteq \mathbb{I}$, we denote by $\text{aa}(AX, I)$ the set of axiom automata over I : $\{\text{aa}(AX, s) \mid s : I(AX) \rightarrow I\}$.

2) *A basic operationalization:* Lemma 3 and the concretization defined in §VII-B1 suggest an operationalization for AX. For a program \mathcal{P} , if a trace σ is accepted by *all* (concrete) automata $\text{aa}(AX, \text{instrsOf}(\mathcal{P}))$ then $\sigma \models \phi[s]$ holds for each assignment s , thus satisfying AX. The number of




Fig. 8: Completed prefix (pCM), in-progress (IP) and not-fetched postfix (pNF) of instructions during execution.

these automata is $|\text{aa}(\text{AX}, \text{instrsOf}(\mathcal{P}))| \sim |\text{instrsOf}(\mathcal{P})|^k$ for an axiom with k universally quantified variables. Since this increases with \mathcal{P} , the model is not finite state. Even so, this enables us to construct operational models for a given bound on $|\text{instrsOf}(\mathcal{P})|$. We can do this even for non-universal axioms by converting existential quantifiers into finite disjunctions over $\text{instrsOf}(\mathcal{P})$. We demonstrate an application of this in §VIII, where we check that a processor satisfies an axiom ensuring correctness of read values.

C. Bounding the number of active instructions

As the discussion from §VII-B2 concludes, generating all concrete automata (statically) for arbitrary μspec specifications does not give us a finite state model. We need to bound the number of automata maintained at any point in the trace. In order to do this, for each index in the trace, we identify *active* instructions: an active instruction is one for which we need to maintain ordering information at that index. We observe that under the t -bounded reordering under-approximation, only a bounded number of instructions are active. This, in turn implies that we only need to maintain a bounded number of axiom automata. We now formalize these concepts.

For a t -reordering bounded trace σ of a program \mathcal{P} and a trace index $0 \leq j \leq |\sigma|$, let $\text{CM}(j)$ and $\text{NF}(j)$ be instructions which have executed all and none of their events at $\sigma[j]$ respectively. We define the following auxiliary terms:

$$\begin{aligned} \text{pCM}(j) &= \{i \mid \forall i'. i' \leq_r i \implies i' \in \text{CM}(j)\} \\ \text{pNF}(j) &= \{i \mid \forall i'. i \leq_r i' \implies i' \in \text{NF}(j)\} \\ \text{IP}(j) &= \text{instrsOf}(\mathcal{P}) \setminus (\text{pCM}(j) \cup \text{pNF}(j)) \end{aligned}$$

Intuitively $\text{pCM}(j)$ represents the prefix-closed set of *completed* instructions, $\text{pNF}(j)$ represents the postfix-closed set of *not-fetched* instructions, and $\text{IP}(j)$ are the rest - the *in-progress* instructions (see Fig. 8). By the first condition of t -reordering boundedness, in-progress (IP) instructions on each core are bounded by t for all j (Lemma 4):

Lemma 4. For any t -reordering bounded trace σ , for all $0 \leq j \leq |\sigma|$, we have, $|\text{IP}(j)| \leq |\text{Cores}| \cdot t$.

Active instructions Two instructions i, i' are k -coupled in a trace σ if they form a coupling chain of length k : i.e. there exist instructions i_1, \dots, i_{k-1} such that $\text{coup}(i, i_1), \text{coup}(i_1, i_2), \dots, \text{coup}(i_{k-1}, i')$. For trace σ , $0 \leq j \leq |\sigma|$ and $k \in \mathbb{N}$, we define k -active instructions at j , $\text{AC}_k(j)$, as instructions from $\text{pCM}(j) \cup \text{IP}(j)$ which are k -coupled with some instruction from $\text{IP}(j)$.

Intuitively, for a μspecRE axiom with k universally quantified variables, the execution of two instructions affect each




Fig. 9: Experimental setup.

other only if they are k -coupled. In particular, maintaining ordering information is important for instructions which are k -coupled with the in-progress instructions. As Lemma 5 shows, these *active* instructions - $\text{AC}_k(j)$ - are bounded at any given point in the trace.

Lemma 5. For each k , there is a (program-independent) bound b_k , s.t. for any t -reordering bounded trace σ , for all $0 \leq j \leq |\sigma|$, we have $|\text{AC}_k(j)| \leq b_k$.

The operational model Our operational model maintains the in-progress instructions (IP) on its tape. At each step it schedules an event from these instructions. The validity of event scheduling is ensured by maintaining orderings between events corresponding to the active instructions. Lemmas 4, 5 imply that at all points in the trace, (1) the set IP is bounded and (2) the active instructions - AC_k - are bounded (as a function of b_k). Consequently, this results in a model which has finite state (used to maintain orderings between events of AC_k) and bounded history (owing to (1)). This gives us the main result - a finite state, bounded history operational model.

Theorem 2. For a (refinable) universal axiomatic semantics that satisfies extensibility, synthesis of finite-state, bounded-history operational models satisfying Def. 3 and 5* is feasible.

VIII. CASE STUDIES

In this section, we demonstrate applications of operationalization. We discuss three case studies: (1) `multi_vscale` [30] is a multi-core extension of the 3-stage in-order `vscale` [31] processor, (2) `tomasulo` is an OoO processor based on [32], and (3) `sdram_ctrl` is an SDRAM-controller [33].

For each case, we instrument the hardware designs by exposing ports that signal the execution of events (e.g. the PC ports in Fig. 9). We convert axioms into an operational model \mathcal{M} based on the approach discussed in §VII. \mathcal{M} is compiled to RTL and is synchronously composed with the hardware design, where it transitions on the exposed event signals. Thus, any violating behaviour of the hardware will lead \mathcal{M} into a non-accepting (`bad`) state. Hence by specifying `!bad` as a safety property, we can perform verification of the RTL design w.r.t. the axioms. The operationalization approach enables us to perform both bounded and unbounded verification using off-the-shelf hardware model checkers. We highlight that this would not have been possible without operationalization.

We use the Yosys-based [34] SymbiYosys as the model-checker, with boolector [35] and abc [36] as backend solvers for BMC and PDR proof strategies respectively. Experiments

Instructions	PDR	BMC ($d = 20$)
ALU-R	1m46s	14m30s
ALU-I	2m11s	11m31s
Load+Store	2m18s	13m35s

Fig. 10: Proof runtimes for (**ax1** \wedge **ax2**).

are performed on an Intel Core i7 machine with 16GB of RAM. We use our algorithm to automatically generate axiom automata. The compilation of the generated automata to RTL and their instrumentation with the design is done manually. However, in the future this could be automated following the procedure developed in §VII. The experimental designs are available at <https://github.com/adwait/axiomatic-operational-examples>.

Highlights. We demonstrate how the operationalization framework enables us to leverage off-the-shelf model checking tools implementing bounded and (especially) unbounded proof techniques such as IC3/PDR. This would not have been possible directly with axiomatic models. Even when Thm. 2 does not apply (e.g. non-universal/non-extensible axioms), following §VII-B2 we can fall back on a BMC-based check over all possible programs under a bound on $|\text{instrsOf}(\mathcal{P})|$.

A. The *multi_vscale* processor

a) *Pipeline axioms on a single core:* We begin with the single-core variant of *multi_vscale*. We are interested in verifying the pipeline axioms for this core. The first axiom states that pipeline stages must be in **Fet-DX-WB** order and the second enforces in-order fetch.

$$\begin{aligned} \mathbf{ax1}: & \forall i1. (\mathbf{hb}(i1.\mathbf{Fet}, i1.\mathbf{DX}) \wedge \\ & \quad \mathbf{hb}(i1.\mathbf{DX}, i1.\mathbf{WB})) \\ \mathbf{ax2}: & \forall i1, i2. i1 <_r i2 \Rightarrow \mathbf{hb}(i1.\mathbf{Fet}, i2.\mathbf{Fet}) \end{aligned}$$

The setup schematic is given in Figure 9: \mathcal{M} is the operational model implemented in RTL (note that we could do this only because the model is finite state and requires a finite history h). Given that it is a 3-stage in-order processor, at any given point each core has at most 3 instructions in its pipeline and we can safely choose a history parameter of $h = 3$, and \mathcal{M} is complete for a reordering bound of $t = 3$. We replace the *imem_hrdata* (instruction data) connection to the core by an input signal that we can symbolically constrain. Using this input signal, we can control the program (instruction stream) executed by the core.

Verification is performed with a PDR based proof using the *abc pdr* backend. We experiment with various choices of instructions fed to the processor (by symbolically constraining *imem_hrdata*). In Fig. 10, we show the constraint and its PDR proof runtime, with BMC runtime (depth = 20) for comparison. These examples demonstrate our ability to prove unbounded correctness.

b) *Memory ordering on multi-core:* We now configure the design with 2 cores: c_0, c_1 , both initialized with symbolic load and store operations. We then perform verification w.r.t. the **ReadValues** (RV) axiom shown below. This axiom says

$ I $	$ AA $	BMC d	Time
4	16	12	3m10s
6	36	16	15m48s
8	64	20	1h58m

Fig. 11: Proof runtimes for the Read-Values axiom for different instruction counts ($|I|$).

that for any read instruction (**i1**), the value read should be the same as the most recent write instruction (**i2**) on the same address, or it should be the initial value.

$$\begin{aligned} \mathbf{RV}: & \forall i1, \exists i2, \forall i3. \mathbf{IsRead}(i1) \Rightarrow \\ & (\mathbf{DataInit}(i1) \vee (\mathbf{IsWrite}(i2) \wedge \\ & \quad \mathbf{SameAddr}(i1, i2) \wedge \mathbf{hb}(i2.\mathbf{DX}, i1.\mathbf{DX}) \\ & \quad \wedge \mathbf{ValEq}(i1, i2) \wedge ((\mathbf{IsWrite}(i3) \wedge \\ & \quad \mathbf{SameAddr}(i1, i3)) \Rightarrow \\ & \quad (\mathbf{hb}(i3.\mathbf{DX}, i2.\mathbf{DX}) \vee \mathbf{hb}(i1.\mathbf{DX}, i3.\mathbf{DX})))))) \end{aligned}$$

This not a universal axiom, and hence Thm. 2 does not apply. However, for bounded programs we can construct $|\text{instrsOf}(\mathcal{P})|^2$ concrete automata (since there are two universally quantified variables: **i1**, **i3**) as discussed in §VII-B2. We convert the existential quantifier over **i2** into a finite disjunction over $\text{instrsOf}(\mathcal{P})$. We perform BMC queries for programs with $|I| = |\text{instrsOf}(\mathcal{P})| = 4, 6, 8$.

By keeping instructions symbolic, we effectively prove correctness for *all* programs within our bound $|I|$. The table alongside shows the instruction bound, $|I|$, the number of axiom automata $|AA|$, BMC depth d , and proof runtime. Though our theoretical results apply to universal axioms, this shows how an axiom automata-based operationalization can be applied to arbitrary axioms by bounding $|\text{instrsOf}(\mathcal{P})|$.

B. An OoO processor: *tomasulo*

Our second design is an out-of-order processor (based on [32]) that implements Tomasulo’s algorithm. The processor has stages: **F** (fetch), **D** (dispatch), **I** (issue), **E** (execute), **WB** (writeback), and **C** (commit). We verify in-order-commit, program-order fetch, and pipeline order axioms for this processor. A BMC proof (with $d = 20$) takes $\sim 2\text{m}$.

The axiom **axDep** given below is crucial for correct execution in an OoO processor. It enforces that execute (**E**) stages for consecutive instructions should be in program order if the destination of the first instruction is same as the source of the second, i.e. dependent instructions are executed in order.

$$\mathbf{axDep}: \forall i1, i2. (i1 <_r i2 \wedge \mathbf{Cons}(i1, i2) \wedge \mathbf{DepOn}(i1, i2)) \Rightarrow \mathbf{hb}(i1.\mathbf{E}, i2.\mathbf{E})$$

We add a program counter (*pc*) to instructions and define $\mathbf{Cons}(i_1, i_2) \equiv \text{pc}(i_1) + 4 = \text{pc}(i_2)$ and $\mathbf{DepOn}(i_1, i_2) \equiv \text{dest}(i_1) = \text{src1}(i_2) \vee \text{dest}(i_1) = \text{src2}(i_2)$.

As before, we compose the operational model \mathcal{M} corresponding to this axiom with the RTL design. We symbolically constrain the processor to execute a sequence of symbolic (add and sub) instructions and assert **!bad**. A BMC query

($d = 20$) results in an assertion violation. We manually identified the bug as being caused by the incorrect reset of entries in the Register Alias Table (RAT) in the `Com` stage. When committing instruction i_0 , the entry $\text{RAT}(\text{dest}(i_0))$ is reset, while some instruction i_1 with $\text{dest}(i_0) = \text{dest}(i_1)$ is issued at the same cycle. A third instruction i_2 with $\text{src1}(i_2) = \text{dest}(i_0)$ then reads the result of i_0 instead of i_1 , violating the axiom. We fix this bug and perform a BMC proof ($d = 20$), which takes $\sim 6\text{m}30\text{s}$. This demonstrates how our technique can be used to identify a bug, correct it and check the fixed design.

C. A memory controller: `sdram_ctrl`

To demonstrate the versatility of our approach, we experiment with an SDRAM controller [33], which interfaces a processor host with an SDRAM device, with a ready-valid interface for read/write requests. All intricacies related to interfacing with the SDRAM are handled by maintaining appropriate control state in the controller. In the following, we once again convert axioms into an operational model by our technique, and compose the generated model with the design.

First we verify pipeline-stage axioms for `sdram_ctrl` for write (4-stages) and read (5-stages) operations executed by the host. A PDR-based (unbounded) proof for the pipeline axioms requires $\sim 8\text{m}$. Next we verify properties related to SDRAMs refresh operation [37]. The controller ensures that the host-level behaviour is not affected by refreshes by creating an illusion of atomicity for writes and reads. This results in the axiom that once a write or read operation is underway, no refresh stage should execute before it is completed. We once again prove this property with PDR, which takes $\sim 1\text{m}30\text{s}$.

IX. RELATED WORK

There has been much work on developing axiomatic (declarative) models for memory consistency in parallel systems, at the ISA level [2], [38], [39], the microarchitectural level [12], [16], [11], and the programming language level [20], [40], [41], [42], [43]. There has also been work on constructing equivalent operationalizations for these models, e.g., for Power [2], ARMv8 [10], RA[8], C++ [7], and TSO [19], [9]. These constructions are accompanied by hand-written/theorem-prover based proofs, demonstrating equivalence with the axiomatic model. In principle, our work is related to these, however we enable *automatic* generation of equivalent operational models from axiomatic ones, eliminating most of the manual effort.

At an abstract level, we have been inspired by classic works that have developed connections between logics and automata [44], [45]. There is a large body of work on synthesis of operational implementations as well as monitors from temporal specifications (e.g. [46], [47], [48]), most commonly those written in Linear Temporal Logic (LTL) [49] and its variants (e.g. [50]). In this paper we perform a similar conversion but for a very different logic: μspec specifies constraints over partial orders while LTL does so over totally ordered traces. Additionally, the elements over which constraints are enforced

is also different: μspec constrains orderings of a known set of events, while LTL does so over traces with potentially differing sets of events (atoms). These differences make a direct comparison with the previously mentioned works ineffectual, and have required us to develop novel concepts in this work.

In terms of the application to proving properties, the work closest to ours is RTLCheck [13], which compiles constraints from μspec to SystemVerilog assertions. These assertions are checked on a per-program basis. On the other hand, we demonstrate the ability to prove unbounded correctness. Additionally, for axioms that are not generally operationalizable (for unbounded programs), we demonstrate the ability to generate an operational model for some apriori known bound on the program size. In this case, we can verify correctness for *all* programs of size upto that bound, as opposed to on a per-program basis as RTLCheck does. RTL2 μspec [51] aims to perform the reverse conversion: from RTL to μspec axioms.

X. CONCLUSION

In this paper we make strides towards enabling greater interoperability between operational and axiomatic models, both through theoretical results and case studies. We derive μspecRE , a restricted subset of the μspec domain-specific language for axiomatic modelling. We show that the generation of an equivalent finite-state operational model is impossible for general μspec axioms, though it is feasible for universal axioms in μspecRE . From a practical standpoint, we develop an approach based on axiom automata that enables us to automatically generate such equivalent operational models for universally quantified axioms in μspecRE (or for arbitrary μspec axioms if equivalence up to a bound is sufficient).

The challenges we surmount for our conversion (discussed in §I) find parallels in manual operationalization works [7], and we believe that the above concepts can be extended to formalisms such as Cat [2]. Our practical evaluation illustrates the key impact of this work—its ability to enable users of axiomatic models to take advantage of the vast number of techniques that have been developed for operational models in the fields of formal verification and synthesis.

XI. FUTURE WORK

An interesting direction for future work is to enrich μspec semantics (e.g., with quantitative operators) such that valid executions are guaranteed to satisfy t -reordering boundedness. In addition to allowing generation of finite-state operational models, we believe that such axioms would also capture processor executions more precisely.

While some aspects of executions are easier to specify operationally, others (e.g., non-deterministic scheduling) are better suited to axiomatic specifications. Another direction for future work is combining operational and axiomatic modelling, for example using tools such as UCLID5 [52], [53].

ACKNOWLEDGMENTS

This work was supported in part by Intel under the Scalable Assurance program and by DARPA contract FA8750-20-C-0156.

REFERENCES

- [1] Christel Baier and Joost-Pieter Katoen. Principles of model checking. 2008.
- [2] Jade Alglave, Luc Maranget, and Michael Tautschnig. Herding Cats: Modelling, Simulation, Testing, and Data Mining for Weak Memory. *ACM Trans. Program. Lang. Syst.*, 36(2), July 2014.
- [3] Yatin A. Manerkar. *Progressive Automated Formal Verification of Memory Consistency in Parallel Processors*. PhD thesis, Princeton University, Princeton, NJ, USA, 2020.
- [4] Jerry R. Burch and David L. Dill. Automatic verification of pipelined microprocessor control. In *CAV*, 1994.
- [5] Aaron R. Bradley. SAT-based model checking without unrolling. In *VMCAI*, 2011.
- [6] Niklas Eén, Alan Mishchenko, and Robert K. Brayton. Efficient implementation of property directed reachability. *2011 Formal Methods in Computer-Aided Design (FMCAD)*, pages 125–134, 2011.
- [7] Kyndylan Nienhuis, Kayvan Memarian, and Peter Sewell. An operational semantics for C/C++11 concurrency. In *OOPSLA*, 2016.
- [8] Ori Lahav, Nick Giannarakis, and Viktor Vafeiadis. Taming release-acquire consistency. *Proceedings of the 43rd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, 2016.
- [9] Scott Owens, Susmit Sarkar, and Peter Sewell. A better x86 memory model: x86-TSO. In *TPHOLS*, 2009.
- [10] Christopher Pulte, Shaked Flur, Will Deacon, Jon French, Susmit Sarkar, and Peter Sewell. Simplifying ARM concurrency: multicopy-atomic axiomatic and operational models for ARMv8. *Proceedings of the ACM on Programming Languages*, 2:1 – 29, 2018.
- [11] Daniel Lustig, Geet Sethi, Margaret Martonosi, and Abhishek Bhat-tacharjee. COATCheck: Verifying Memory Ordering at the Hardware-OS Interface. In *21st International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, 2016.
- [12] Daniel Lustig, Michael Pellauer, and Margaret Martonosi. PipeCheck: Specifying and verifying microarchitectural enforcement of memory consistency models. *2014 47th Annual IEEE/ACM International Symposium on Microarchitecture*, pages 635–646, 2014.
- [13] Yatin A. Manerkar, Daniel Lustig, Margaret Martonosi, and Michael Pellauer. RTLCheck: Verifying the memory consistency of RTL designs. *2017 50th Annual IEEE/ACM International Symposium on Microarchitecture (MICRO)*, pages 463–476, 2017.
- [14] Yatin A. Manerkar, Daniel Lustig, Margaret Martonosi, and Aarti Gupta. PipeProof: Automated Memory Consistency Proofs for Microarchitectural Specifications. *2018 51st Annual IEEE/ACM International Symposium on Microarchitecture (MICRO)*, pages 788–801, 2018.
- [15] Caroline Trippel, Daniel Lustig, and Margaret Martonosi. CheckMate : Automated exploit program generation for hardware security verification. 2018.
- [16] Yatin A. Manerkar, Daniel Lustig, Michael Pellauer, and Margaret Martonosi. CCIcheck: Using μ hb graphs to verify the coherence-consistency interface. *2015 48th Annual IEEE/ACM International Symposium on Microarchitecture (MICRO)*, pages 26–37, 2015.
- [17] Adwait Godbole, Yatin A. Manerkar, and Sanjit A. Seshia. Automated Conversion of Axiomatic to Operational Models: Theory and Practice. <https://arxiv.org/abs/2208.06733>, 2022.
- [18] Jeremy Manson. The Java memory model. In *POPL '05*, 2005.
- [19] Peter Sewell, Susmit Sarkar, Scott Owens, Francesco Zappa Nardelli, and Magnus O. Myreen. x86-TSO: A Rigorous and Usable Programmer’s Model for x86 Multiprocessors. *Communications of the ACM*, 53:89 – 97, 2010.
- [20] Mark Batty, Scott Owens, Susmit Sarkar, Peter Sewell, and Tjark Weber. Mathematizing C++ concurrency. In *POPL '11*, 2011.
- [21] Mustaque Ahamad, Gil Neiger, James E. Burns, Prince Kohli, and Phillip W. Hutto. Causal memory: definitions, implementation, and programming. *Distributed Computing*, 9:37–49, 2005.
- [22] Evgenii Moiseenko, Anton Podkopaev, Ori Lahav, Orestis Melkonian, and Viktor Vafeiadis. Reconciling event structures with modern multiprocessors. *ArXiv*, abs/1911.06567, 2020.
- [23] Alan Jeffrey and James Riely. On thin air reads towards an event structures model of relaxed memory. *2016 31st Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*, pages 1–9, 2016.
- [24] Brian Norris and Brian Demsky. CDSchecker: checking concurrent data structures written with C/C++ atomics. *Proceedings of the 2013 ACM SIGPLAN international conference on Object oriented programming systems languages & applications*, 2013.
- [25] Stavros Aronis. Effective techniques for Stateless Model Checking. 2018.
- [26] Michalis Kokologiannakis and Viktor Vafeiadis. HMC: Model checking for hardware memory models. *Proceedings of the Twenty-Fifth International Conference on Architectural Support for Programming Languages and Operating Systems*, 2020.
- [27] Jean Berstel. Transductions and context-free languages. In *Teubner Studienbücher : Informatik*, 1979.
- [28] Jacques Sakarovitch. Elements of automata theory. 2009.
- [29] Luc Maranget, Jade Alglave, Susmit Sarkar, and Peter Sewell. Litmus: Running Tests against Hardware. In *TACAS'11, 17th International Conference on Tools And Algorithms for the Construction and Analysis of Systems*, Saarbrücken, Germany, March 2011.
- [30] Yatin A. Manerkar. multi-vsacle. https://github.com/ymanerka/multi_vsacle/tree/multicore.
- [31] LGTMCU. vscale. <https://github.com/LGTMCU/vscale>. [Online; accessed 11-05-2021].
- [32] Soham-Das-2021. Tomasulo. <https://github.com/Soham-Das-2021/Tomasulo-Machine>. [Online; accessed 11-05-2021].
- [33] Stafford Horne. SDRAM controller. <https://github.com/stffrdhrn/sdram-controller>. [Online; accessed 11-05-2021].
- [34] Clifford Wolf, Johann Glaser, and Johannes Kepler. Yosys-A Free Verilog Synthesis Suite. 2013.
- [35] Aina Niemetz, Mathias Preiner, and Armin Biere. Boolector 2.0. *J. Satisf. Boolean Model. Comput.*, 9(1):53–58, 2014.
- [36] Berkeley Logic Synthesis and Verification Group. ABC: A system for sequential synthesis and verification, release 70930. <http://www.eecs.berkeley.edu/~alanmi/abc/>.
- [37] Bruce Jacob, Spencer W. Ng, and David T. Wang. Memory systems: Cache, DRAM, disk. 2007.
- [38] Dennis Shasha and Marc Snir. Efficient and correct execution of parallel programs that share memory. *ACM Trans. Program. Lang. Syst.*, 10:282–312, 1988.
- [39] RISC-V Foundation. The RISC-V Instruction Set Manual, Volume I: User-Level ISA, Document Version 2.2.
- [40] Mark Batty, Alastair F. Donaldson, and John Wickerson. Overhauling SC atomics in C11 and OpenCL. *Proceedings of the 43rd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, 2015.
- [41] Viktor Vafeiadis, Thibaut Balabonski, Soham Sundar Chakraborty, Robin Morisset, and Francesco Zappa Nardelli. Common compiler optimisations are invalid in the C11 memory model and what we can do about it. *Proceedings of the 42nd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, 2015.
- [42] Conrad Watt, Christopher Pulte, Anton Podkopaev, G. Barbier, Stephen Dolan, Shaked Flur, Jean Pichon-Pharabod, and Shu yu Guo. Repairing and mechanising the JavaScript relaxed memory model. *Proceedings of the 41st ACM SIGPLAN Conference on Programming Language Design and Implementation*, 2020.
- [43] Ori Lahav, Viktor Vafeiadis, Jeehoon Kang, Chung-Kil Hur, and Derek Dreyer. Repairing sequential consistency in C/C++11. *Proceedings of the 38th ACM SIGPLAN Conference on Programming Language Design and Implementation*, 2017.
- [44] J. Richard Büchi. On a decision method in restricted second order arithmetic. 1990.
- [45] Pierre Wolper, Moshe Y. Vardi, and A. Prasad Sistla. Reasoning about infinite computation paths. In *24th Annual Symposium on Foundations of Computer Science (sfcs 1983)*, pages 185–194, 1983.
- [46] Bernd Finkbeiner and Sven Schewe. Bounded synthesis. *International Journal on Software Tools for Technology Transfer*, 15:519–539, 2012.
- [47] Klaus Havelund and Grigore Rosu. Synthesizing monitors for safety properties. In *TACAS*, 2002.
- [48] Klaus Havelund and Grigore Rosu. Efficient monitoring of safety properties. *International Journal on Software Tools for Technology Transfer*, 6:158–173, 2003.
- [49] Zohar Manna and Amir Pnueli. The temporal logic of reactive and concurrent systems. In *Springer New York*, 1992.
- [50] Vasumathi Raman, Alexandre Donzé, Dorsa Sadigh, Richard M. Murray, and Sanjit A. Seshia. Reactive synthesis from signal temporal logic specifications. *Proceedings of the 18th International Conference on Hybrid Systems: Computation and Control*, 2015.
- [51] Yao Hsiao, Dominic P. Mulligan, Nikos Nikolieris, Gustavo Petri, and Caroline Trippel. Synthesizing formal models of hardware from RTL for efficient verification of memory model implementations. *MICRO-54*:

54th Annual IEEE/ACM International Symposium on Microarchitecture, 2021.

- [52] Sanjit A. Seshia and Pramod Subramanyan. UCLID5: Integrating modeling, verification, synthesis and learning. *2018 16th ACM/IEEE International Conference on Formal Methods and Models for System Design (MEMOCODE)*, pages 1–10, 2018.
- [53] Elizabeth Polgreen, Kevin Cheang, Pranav Gaddamadugu, Adwait Godbole, Kevin Laeuffer, Shaokai Lin, Yatin A. Manerkar, Federico Mora, and Sanjit A. Seshia. UCLID5: Multi-modal formal modeling, verification, and synthesis. In Sharon Shoham and Yakir Vizel, editors, *Computer Aided Verification*, pages 538–551, Cham, 2022. Springer International Publishing.