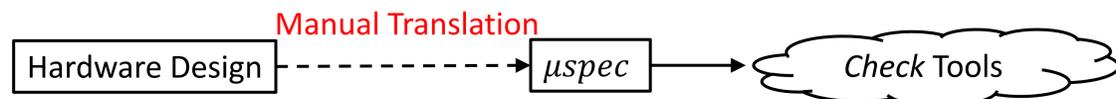


Synthesizing Formal Models of Hardware from RTL for Efficient Verification of Memory Model Implementations

Introduction & Background

- Motivation:** Check tools¹ enable hardware memory consistency model (MCM) verification, but require manually-constructed formal microarchitecture specifications ($\mu spec$ models) as input



- Goal:** Efficient Check-based verification of hardware MCM implementations rooted in RTL
- Key Challenges:** 1) Discerning $\mu spec$ model completeness 2) Resolving the gap between operational RTL and axiomatic $\mu spec$ models

A Taxonomy for Constructing Complete $\mu spec$ Models

- $\mu spec$ models** specify the space of all possible *micro-architectural happens-before* graphs (μhb graphs) using first-order logic axioms that, given a program and a microarchitecture, instantiate

- μhb nodes: Microarchitectural events, i.e., (instruction, location) pairs
- μhb edges: Happens-before (HB) relationships between nodes

- Verilog:** Temporal description of a hardware design (Fig. 1a)
- $\mu spec$ model taxonomy:** Happens-Before Invariants (HBIs)

preserved by an operational Verilog design

- Intra-instruction HBIs:** How each instruction flows through the design (Fig. 1b, \rightarrow in Fig. 2c)

- Inter-instruction HBIs:** Interactions between instructions during program executions

- Structural dependencies: Serialized access to shared state elements (Fig. 1c, \rightarrow in Fig. 2c)
- Dataflow dependencies: Share data through common resources (Fig. 1d, \rightarrow / \rightarrow in Fig. 2c)

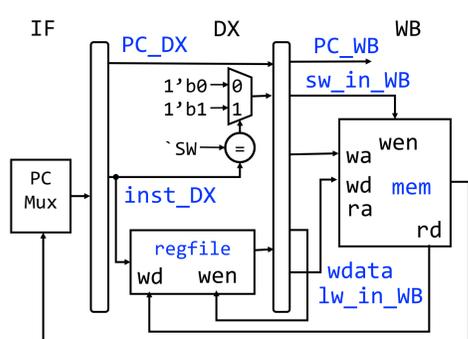


Fig. 1a: Multi-V-Scale

```

    Axiom W_path: forall microops i1,
    IsAnyWrite i1 => AddEdges [
    ((i1, inst_DX), (i2, sw_in_WB));
    (i1, inst_DX), (i1, lw_in_WB));
    (i1, sw_in_WB), (i1, mem)].
  
```

Fig. 1b: Intra-inst. HBIs

```

    Axiom Structural_Temporal:
    forall microops i1, i2,
    IsAnyWrite i1 => IsAnyRead i2 =>
    ProgramOrder i1 i2 => AddEdge(
    (i1, mem), (i2, regfile)).
  
```

Fig. 1c: Inter-inst. Structural HBIs

Core 0	Core 1
(i0) lw x, 0	(i2) lw y, 1
(i1) sw y, 1	

Fig. 1e: Program input. Observable execution in μhb graph of Fig. 2c

```

    Axiom Dataflow:
    forall microops i1, i2,
    IsAnyWrite i1 => IsAnyRead i2 =>
    SamePA i1 i2 => SameData i1 i2 =>
    NoWritesInBetween i1 i2 =>
    AddEdge((i1, mem), (i2, regfile)).
  
```

Fig. 1d: Inter-inst. Dataflow HBIs

Synthesizing a $\mu spec$ Model from RTL with rtl2 $\mu spec$

- A Verilog netlist** contains a subset of the targeted HBIs. A **full-design data-flow graph (DFG)** represents over-approximation of data-flow that can be induced by instruction execution.

- DFG nodes:** instruction-dependent registers
- DFG edges:** over-approximation of information flow one register to another (an HB relation)

- Intra-instruction HBI synthesis:**

- Perform DFS on full-design DFG to generate **intra-instruction HBI hypotheses** as SVAs
- SVA evaluation** outputs full set of **intra-instruction HBIs** and **instruction specific DFGs**

- Inter-instruction HBI synthesis:**

- Pairwise compare instruction-specific DFGs to identify each type of dependency
- Formulate dependencies as **HBI hypotheses** and instantiate them as SVAs
- SVA evaluation outputs a full set of **inter-instruction HBIs**

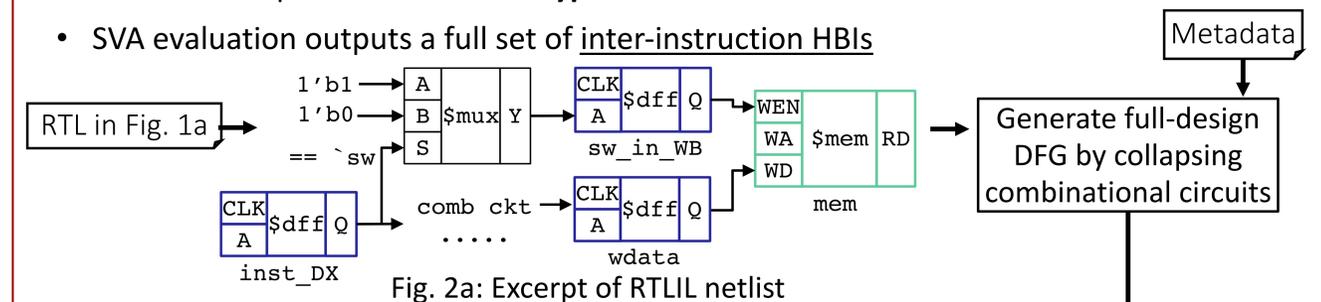


Fig. 2a: Excerpt of RTLIL netlist

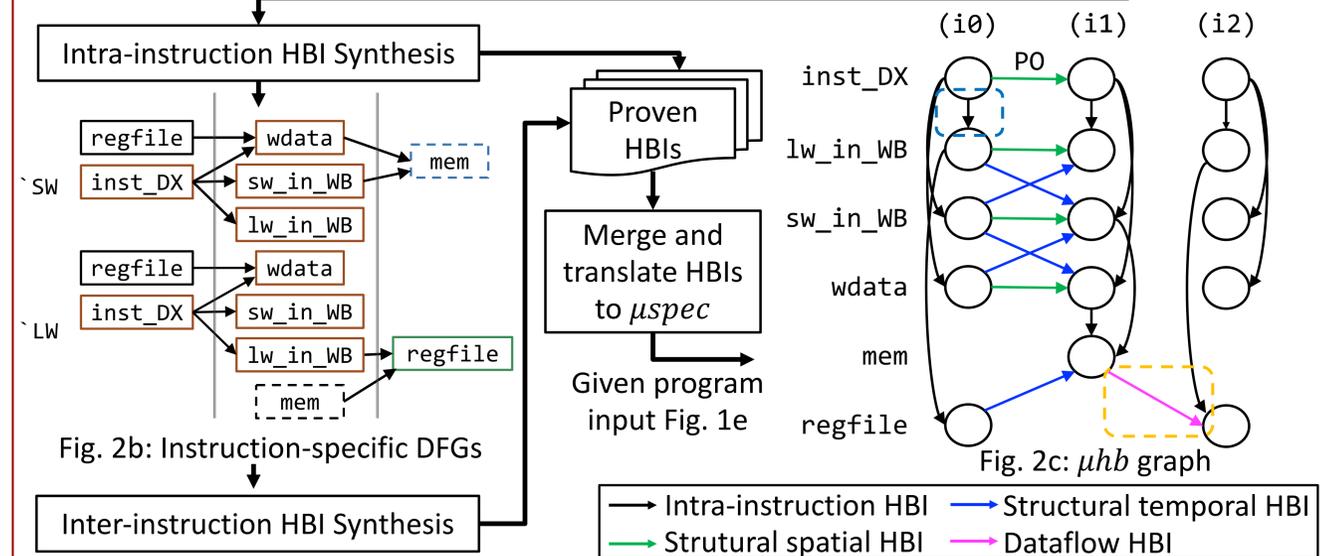
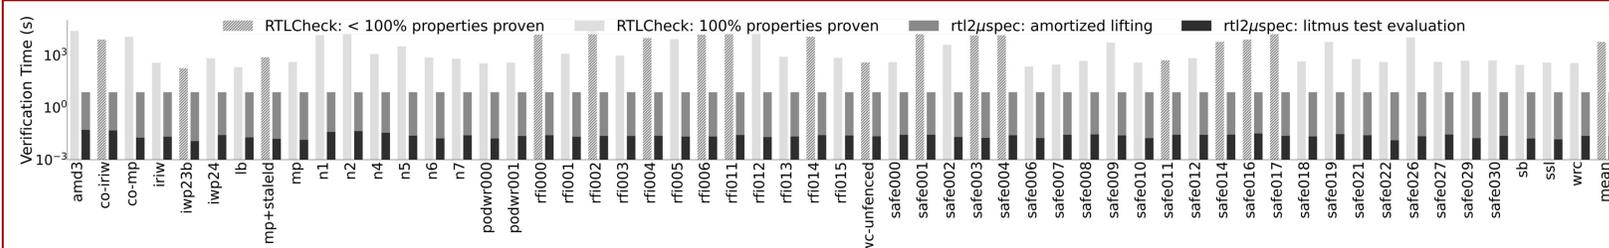


Fig. 2b: Instruction-specific DFGs

Fig. 2c: μhb graph

\rightarrow Intra-instruction HBI
 \rightarrow Structural temporal HBI
 \rightarrow Structural spatial HBI
 \rightarrow Dataflow HBI

Case Study Result



Performance comparison of rtl2 $\mu spec$ -assisted versus RTLCheck [1]-based verification of hardware MCMs (avg: 7.36 sec vs 25 min)

Conclusion & Contribution

- Define what constitutes a complete $\mu spec$ model for an RTL design
- rtl2 $\mu spec$ tool** for synthesizing complete, and proven correct by construction, $\mu spec$ models from RTL with minimal user-intervention
- Verification** of the RISC-V multi-V-Scale MCM implementation: rtl2 $\mu spec$ synthesizes a $\mu spec$ model in 6.84 mins, during which a new bug is found. Subsequent Check-based verification takes less than 1 second per litmus test.

¹ <https://check.cs.princeton.edu/> [1] Manerkar, Yatin A., et al. RTLCheck: Verifying the memory consistency of RTL designs." Proceedings of the 50th Annual IEEE/ACM International Symposium on Microarchitecture. 2017.