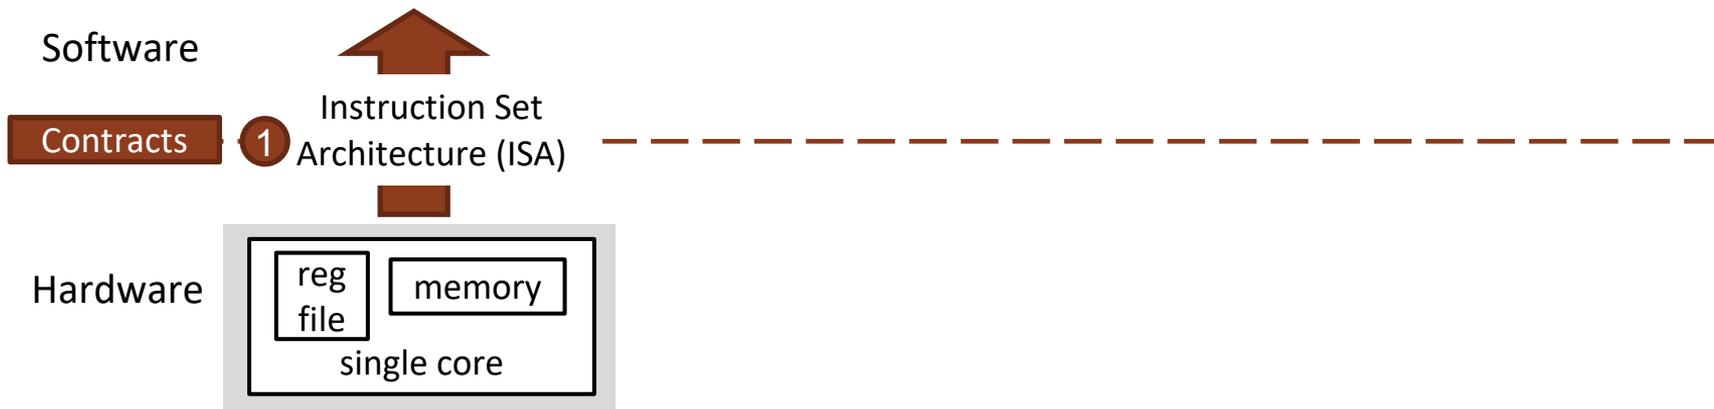


Scalable Assurance via Verifiable Hardware-Software Contracts

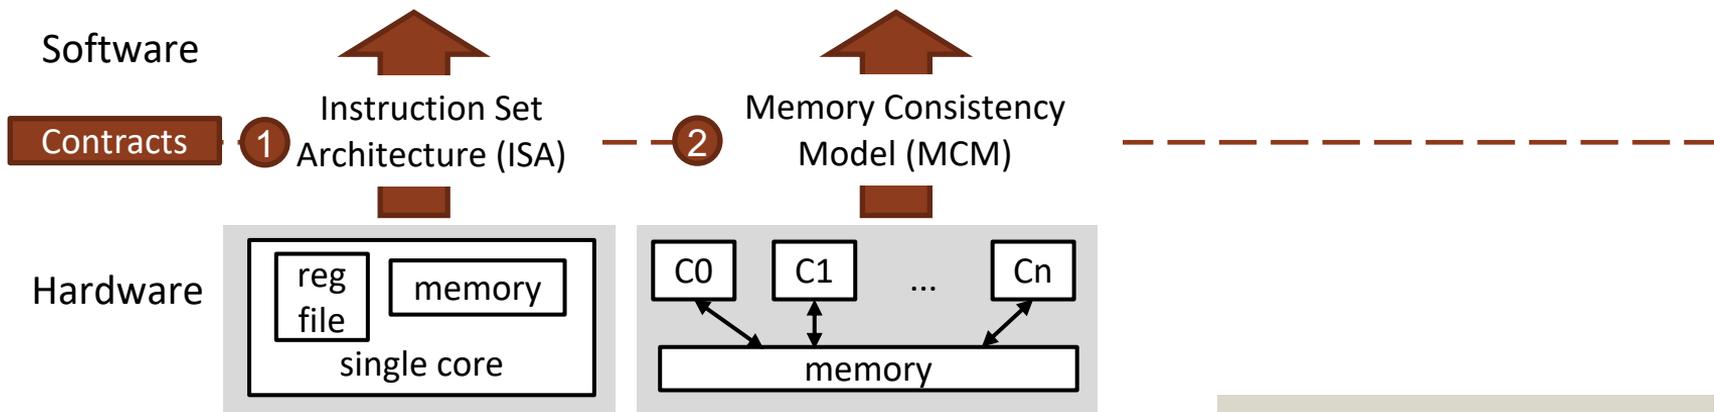
Yao Hsiao¹, Dominic P. Mulligan², Nikos Nikoleris²,
Gustavo Petri², Caroline Trippel¹

¹Stanford University, ²ARM Research

Hardware-software contracts expose hardware correctness and security guarantees to software



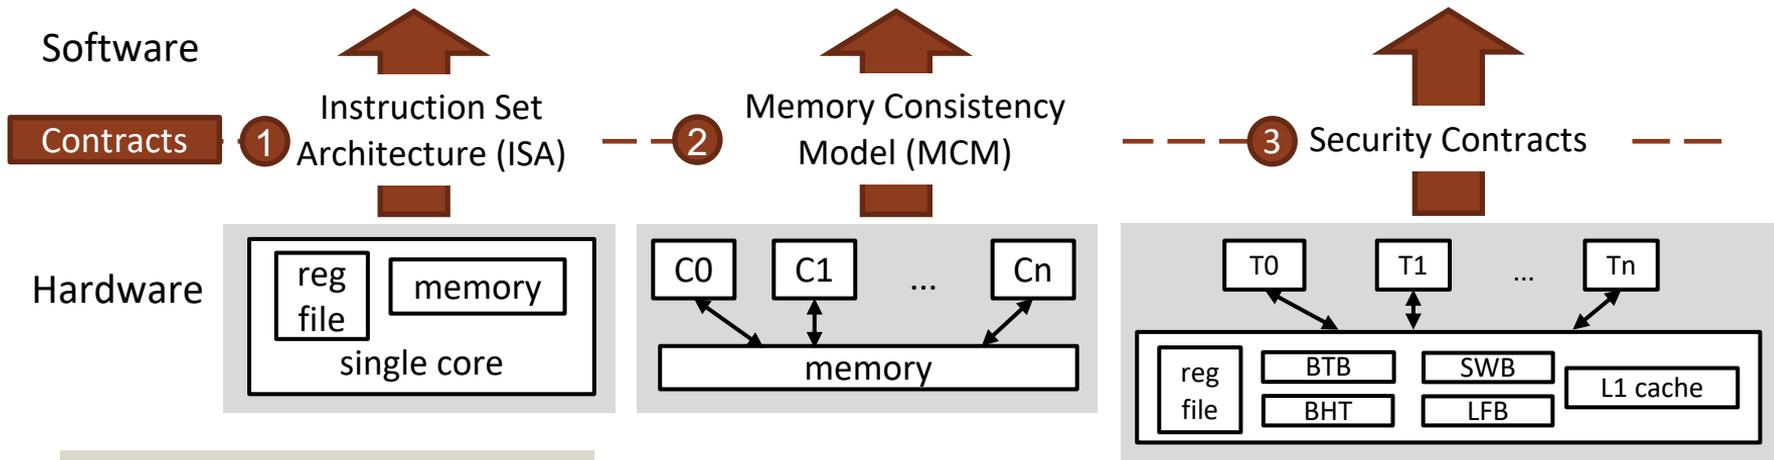
Hardware-software contracts expose hardware correctness and security guarantees to software



x86-tso: Forbids reordering writes
 $W x \rightarrow W y$ in program order

C11 → Power, PLDI'12
C11 → x86, POPL'11
C11 → ARMv7, Sewell+ '16
C11 → ARMv8, Sewell+ '16
Java → Power ECOOP'15
Java → x86, ECOOP'15

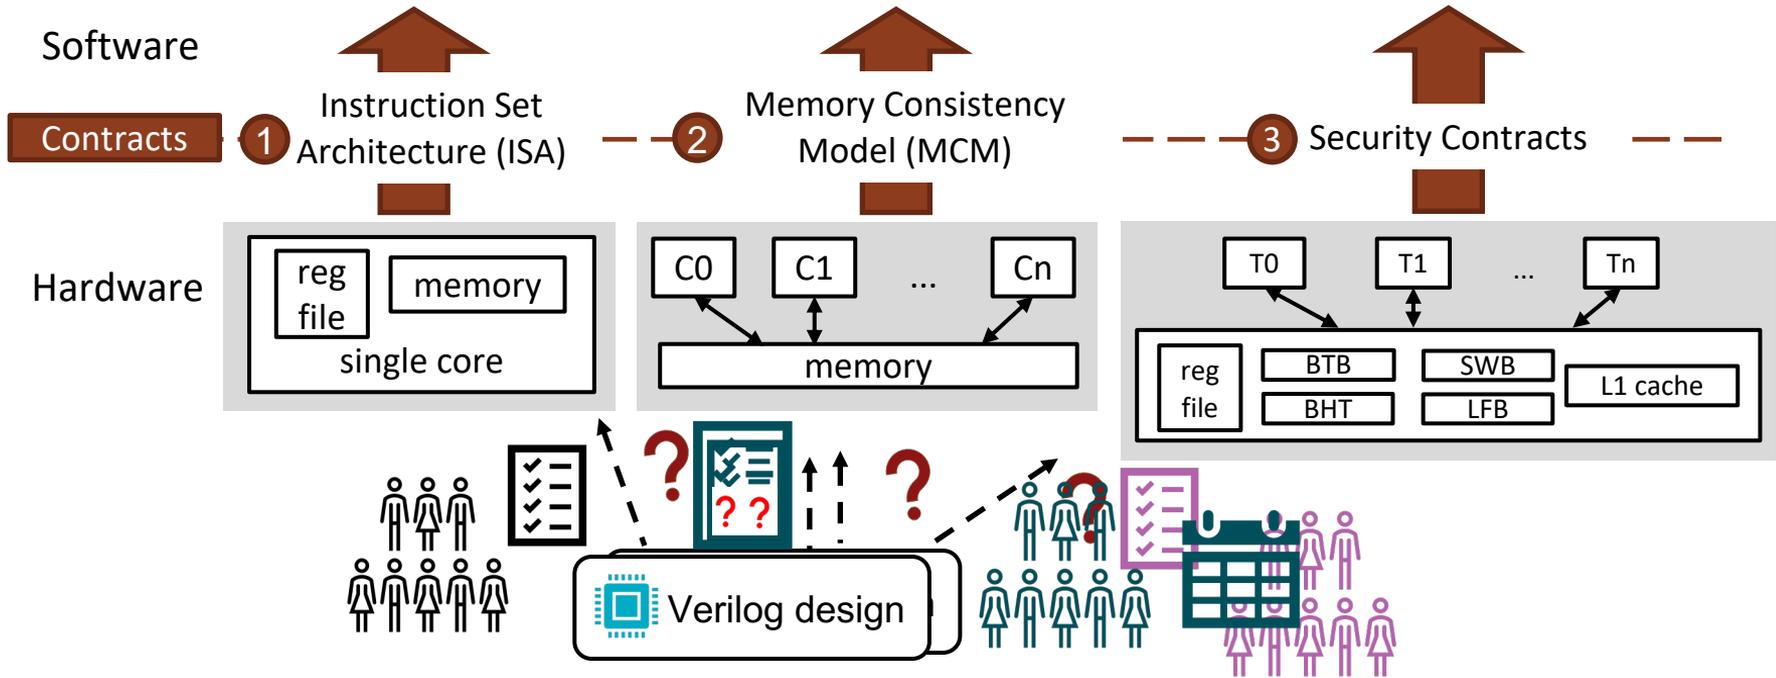
Hardware-software contracts expose hardware correctness and security guarantees to software



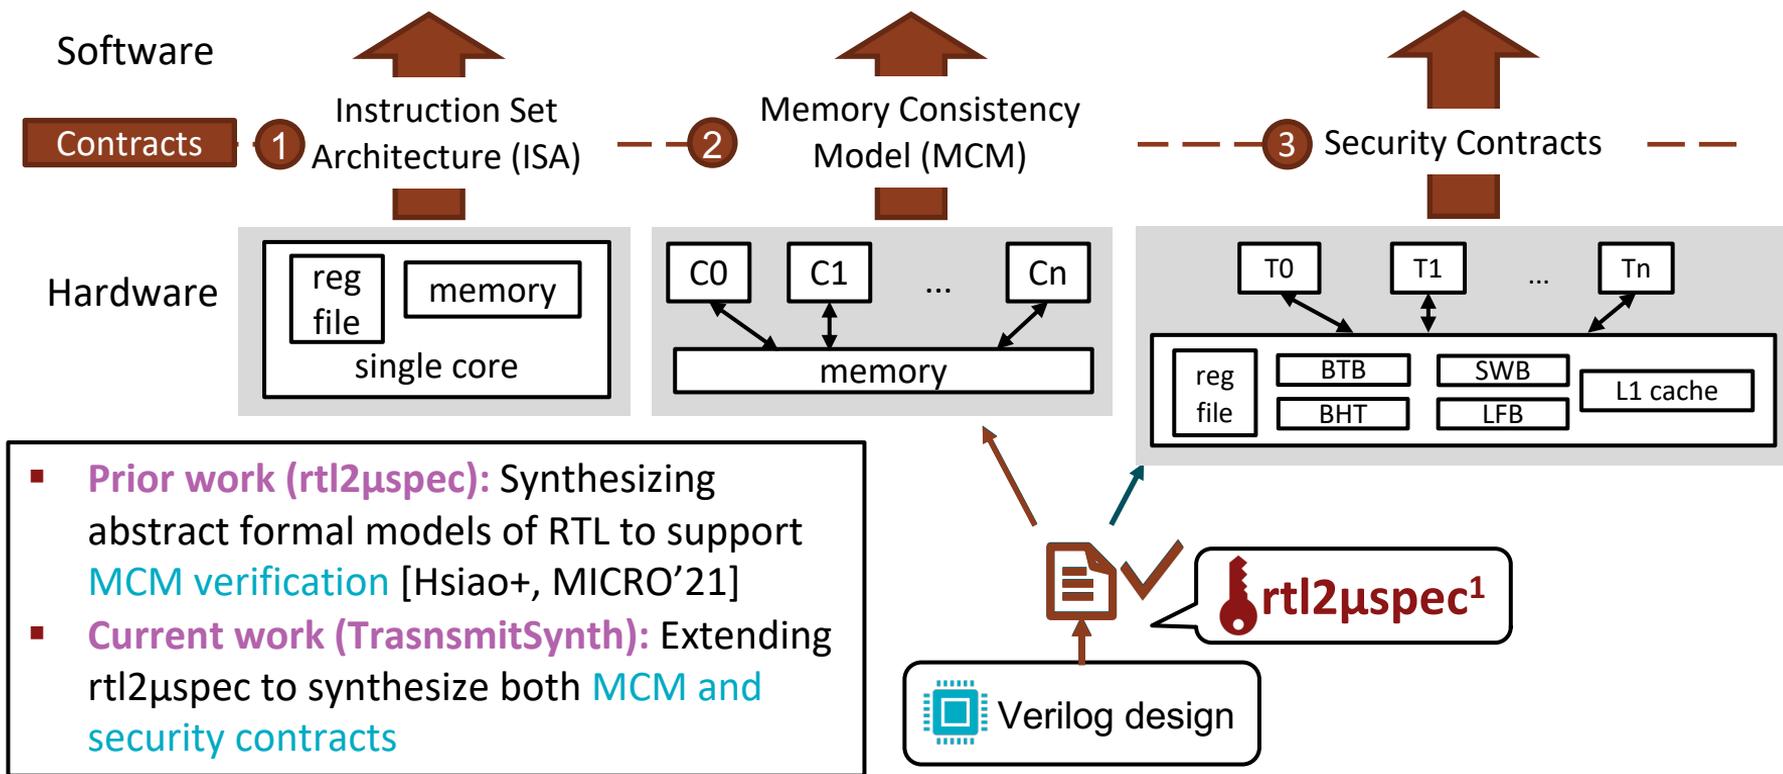
Spetector, S&P'20
Pitchfork, PLDI'20
Blade, POPL'21
Bisec/Haunted, NDSS'21
Clou, ISCA'22

Constant-time programming: Arithmetic instructions
except division are constant time

Hardware-assurance challenge: A gap exists between hardware-software contracts and the RTL they abstract

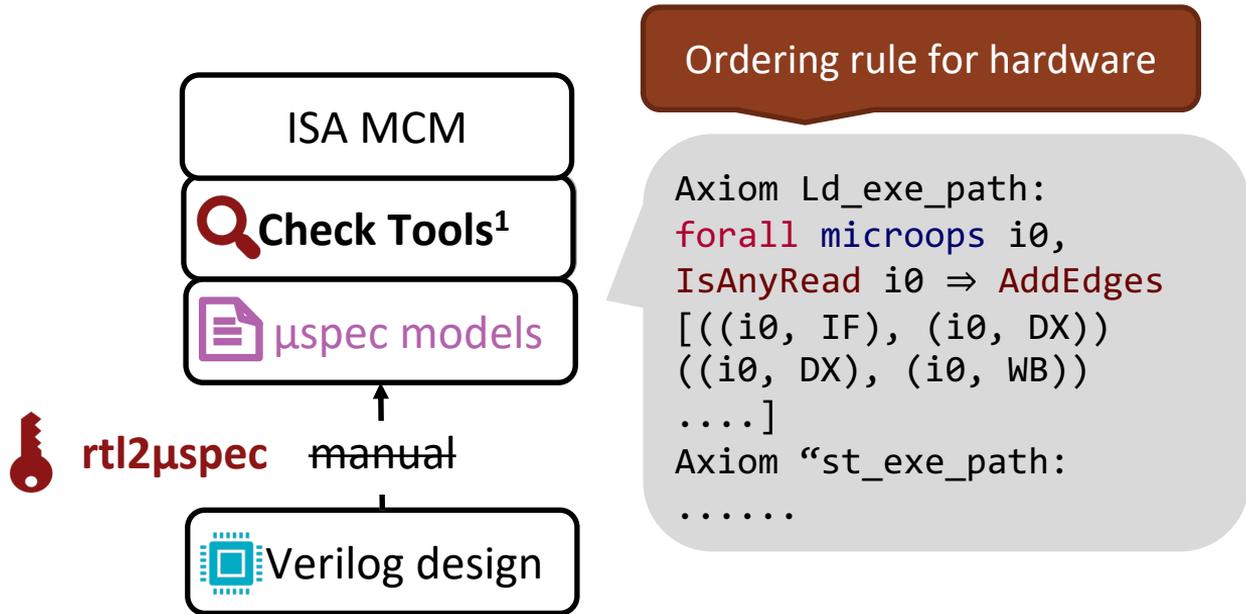


Can we **synthesize** formal hardware-software contracts from RTL?



¹Hsiao et al. "Synthesizing Formal Models of Hardware from RTL for Efficient Verification of Memory Model Implementations." MICRO'21

Check Tools: Automated tools for conducting formal verification of hardware memory model implementations



¹<http://check.cs.princeton.edu>

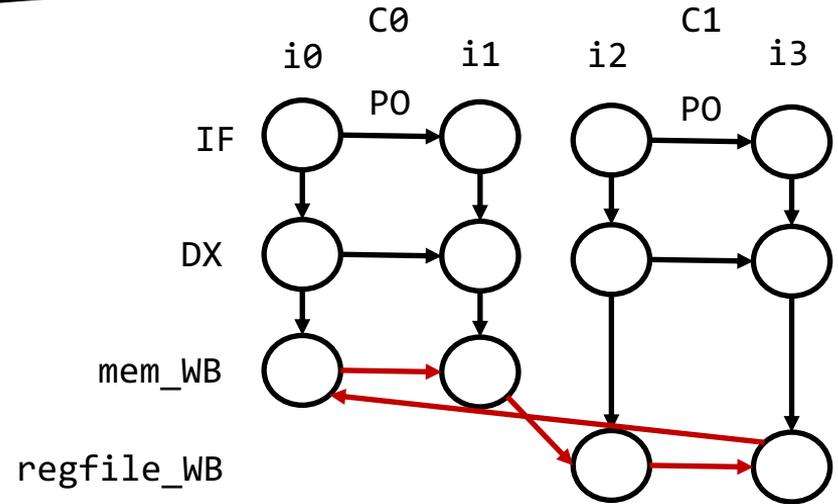
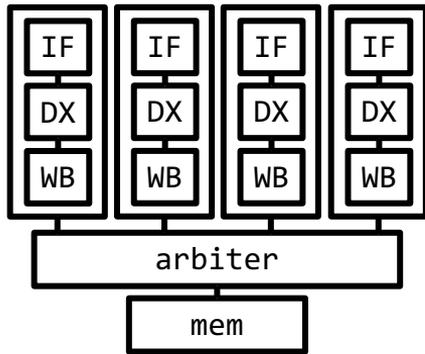
Microarchitectural happens-before (μ hb) analysis

Core 0 Core 1

(i0) W[x] = 1; (i2) R[y] = 1;
(i1) W[y] = 1; (i3) R[x] = 0;

Forbidden on Sequential Consistency (SC)

RISC-V multi-V-scale¹



¹Albert Magyar. 2016. A Verilog implementation of the RISC-V Z-scale microprocessor. <https://github.com/ucb-bar/vscale>.

µhb analysis: Hardware locations

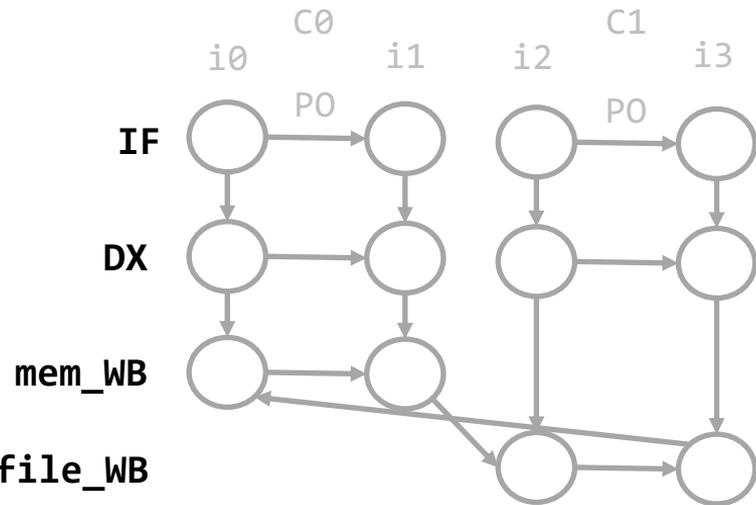
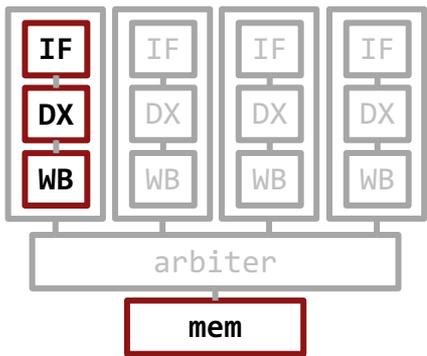
Core 0 Core 1

(i0) W[x] = 1; (i2) R[y] = 1;

(i1) W[y] = 1; (i3) R[x] = 0;

Forbidden on SC

RISC-V multi-V-scale



hardware state elements

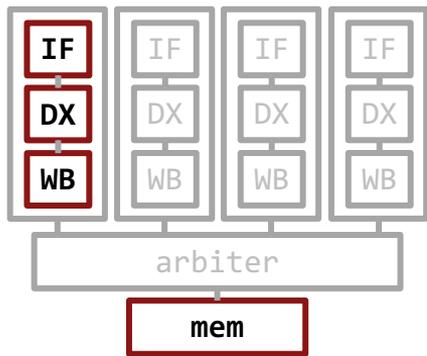
μhb analysis: Instruction execution paths

Core 0 Core 1

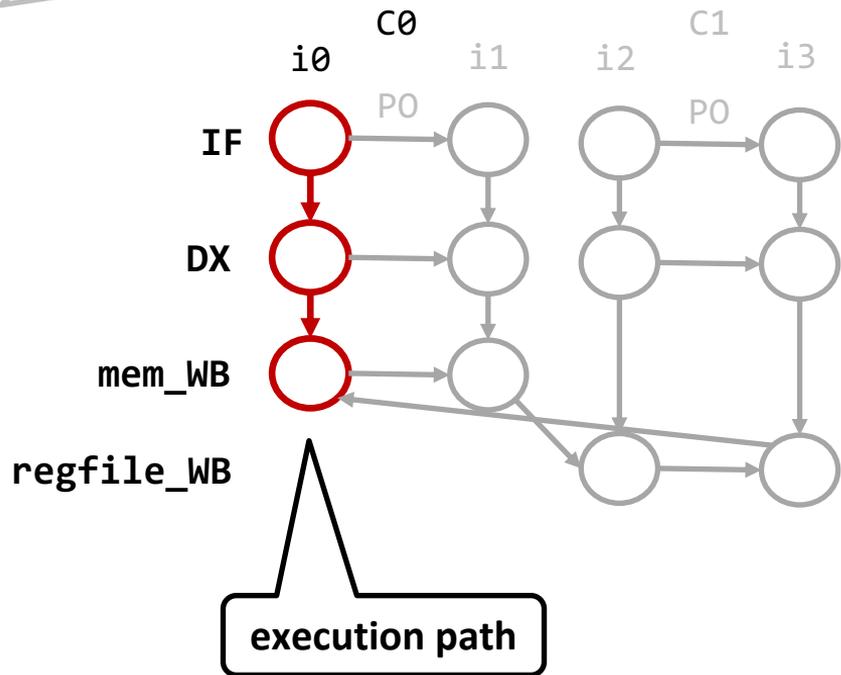
(i0) W[x] = 1; (i2) R[y] = 1;

(i1) W[y] = 1; (i3) R[x] = 0;

RISC-V multi-V-scale



Forbidden on SC



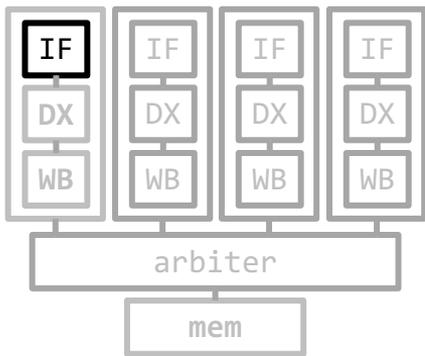
μhb analysis: μhb nodes a.k.a. hardware events

Core 0
(i0) W[x] = 1;
(i1) W[y] = 1;

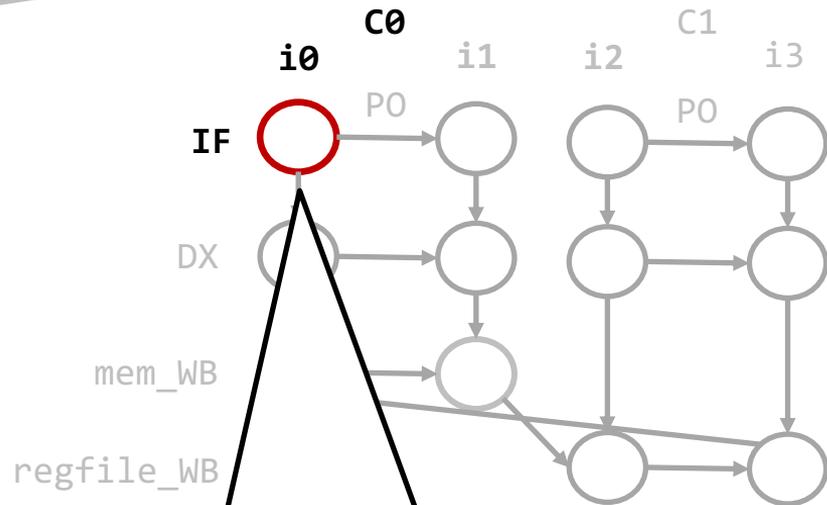
Core 1
(i2) R[y] = 1;
(i3) R[x] = 0;

Forbidden on SC

RISC-V multi-V-scale



Forbidden on SC



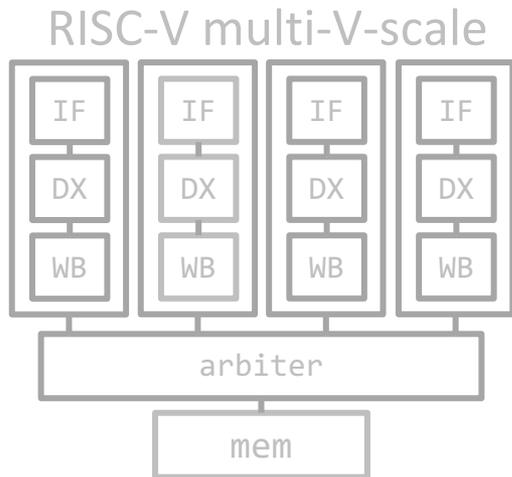
μhb nodes: microarchitectural events,
 <instruction, state element(s)> pair

Reasoning about **execution observability** with μ hb analysis

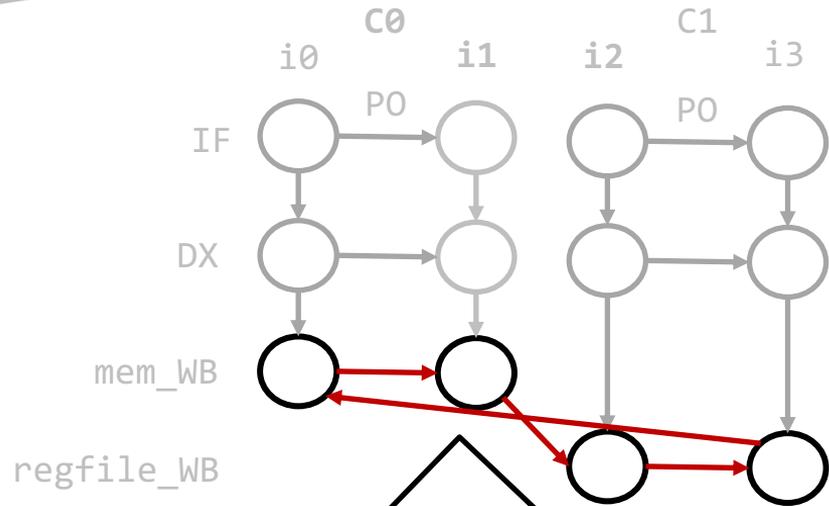
Core 0 Core 1

(i0) W[x] = 1; (i2) R[y] = 1;

(i1) W[y] = 1; (i3) R[x] = 0;



Forbidden on SC



Cyclic μ hb graph \leftrightarrow Execution is **not observable** on the target hardware

Microarchitectural specifications a.k.a. μ spec models

```

Core 0           Core 1
(i0) W[x] = 1;   (i2) R[y] = 1;
(i1) W[y] = 1;   (i3) R[x] = 0;
    
```

Check Tools: search all ways that a program may execute on hardware, according to μ spec model.

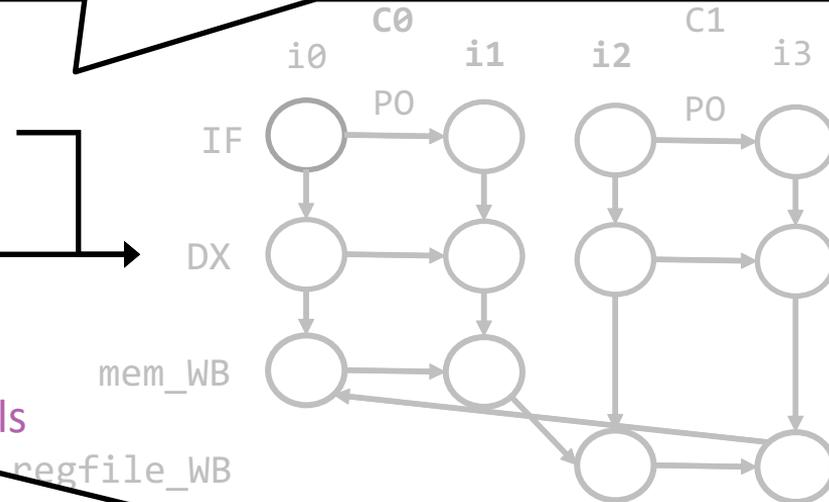
Axi

Manually written ☹️

```

forall microops i0,
IsAnyRead i0 ⇒ AddEdges [
((i0, IF), (i0, DX)),
((i0, DX), (i0, WB)),
....]
....
    
```

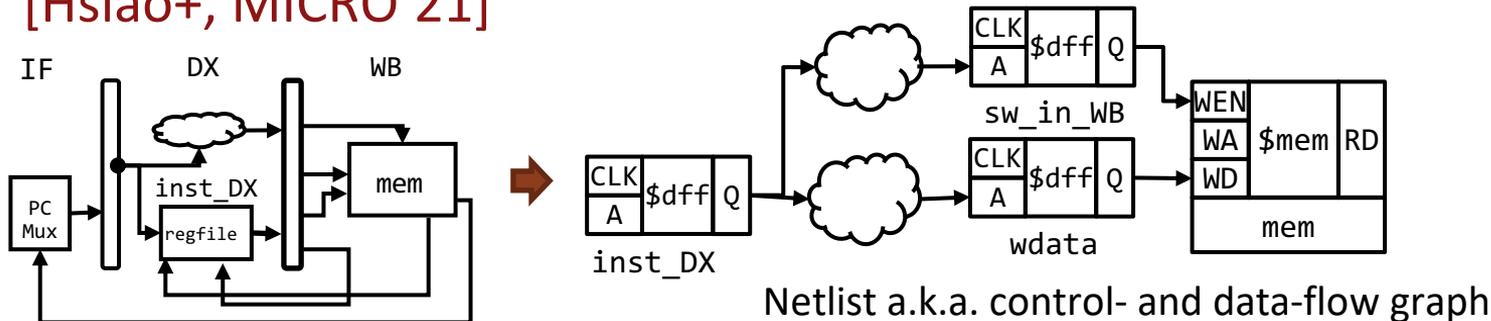
μ spec models



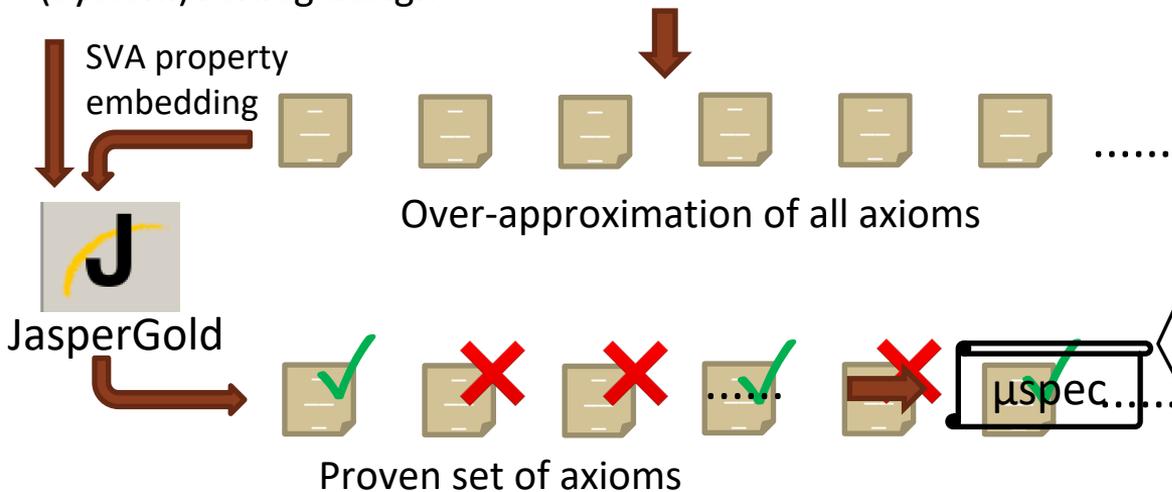
μ spec model of a microarchitecture specifies the space of all possible μ hb graphs—i.e., the space of all possible hardware-specific program executions.



Automated synthesis of μ spec models from RTL with rtl2 μ spec [Hsiao+, MICRO'21]



(System)Verilog design



Open-source RISC-V multi-V-scale case study: 6.84 mins serial proof time w/120 SVA properties evaluated (> 780x performance improvement over prior work [Manerakr+ MICRO'17])

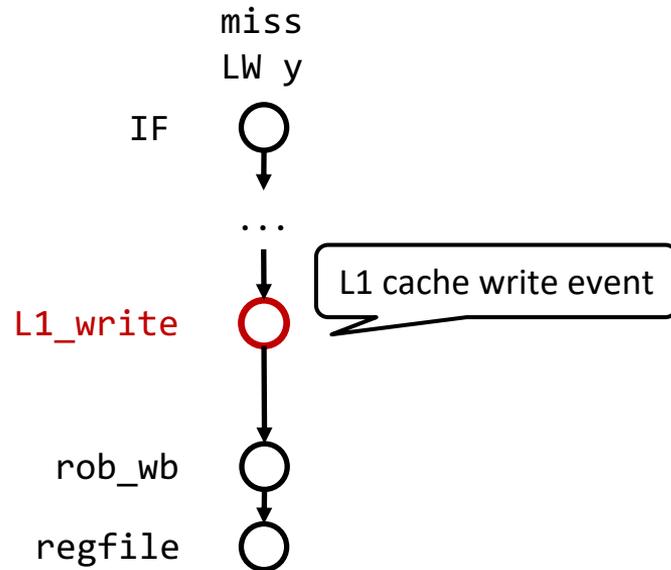
Hardware security verification with μ spec models [Trippel+, MICRO'18]

① LW [base+secret]

② LW [y]

1. Load cache miss path
($y \neq \text{base} + \text{secret}$)

| tag | vld | data |
|-------------|-----|------|
| base+secret | 1 | ... |
| y | 1 | |



A **transmit instruction** (or transmitter) is an instruction which **modulates hardware resources differently** as a function of its **operands, results, or data-at-rest** in a microarchitecture, resulting in **leakage**.

Hardware security verification with μ spec models [Trippel+, MICRO'18]

① LW [base+secret]

② LW [y]

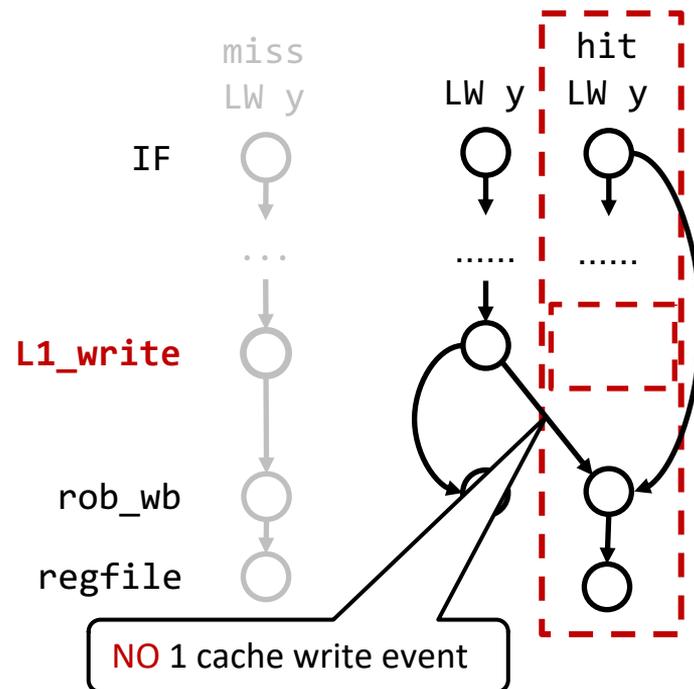
1. Load cache miss path
($y \neq \text{base} + \text{secret}$)

2. Load cache hit path
($y == \text{base} + \text{secret}$)

| tag | vld | data |
|-------------|-----|------|
| base+secret | 1 | ... |
| y | 1 | |

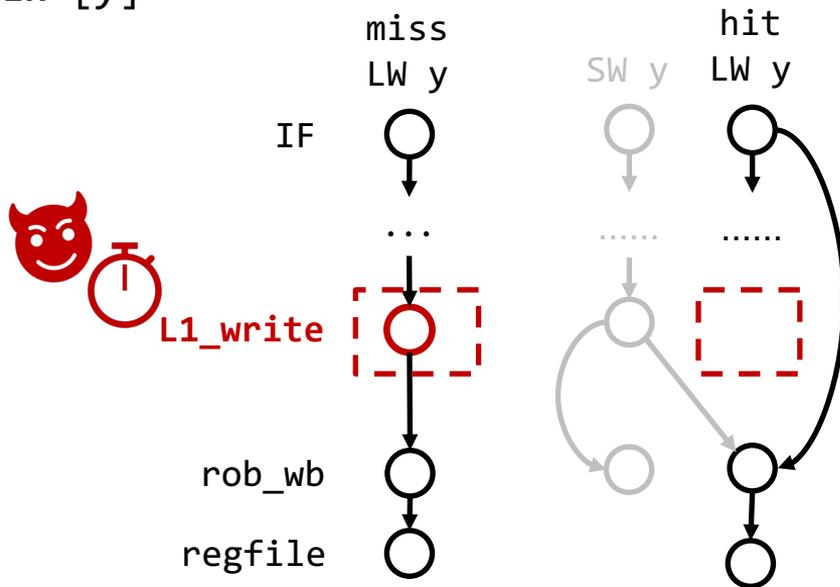
| tag | vld | data |
|-------------|-----|------|
| base+secret | 1 | ... |
| - | - | - |

A **transmit instruction** (or transmitter) is an instruction which modulates hardware resources differently as a function of its operands, results, or data-at-rest in a microarchitecture, resulting in **leakage**.



Key insight: If an instruction is a transmitter, it can instantiate more than one execution path in a μ hb graph

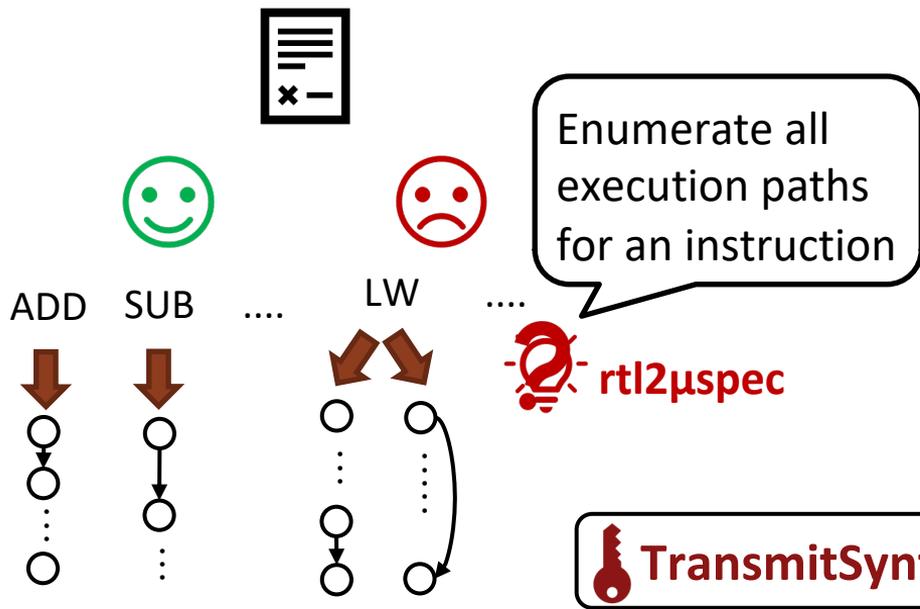
- ① LW [base+secret]
- ② LW [y]



```

Axiom Ld_exe_path:
forall microops i0,
IsAnyRead i0  $\Rightarrow$  AddEdges [
((i0, IF), (i0, DX)),
((i0, L1_write), (i0, rob_wb)),
....] // miss_path
 $\oplus$ 
AddEdges [
((i0, IF), (i0, DX)),
((i0, IF), (i0, rob_wb)),
....] // hit path
    
```

TransmitSynth: Automated synthesis of security contracts from RTL via identification of transmit instructions

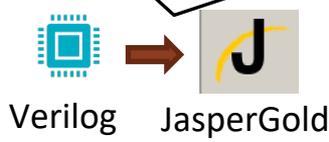


Key limitation of rtl2µspec: the **single execution path assumption**.

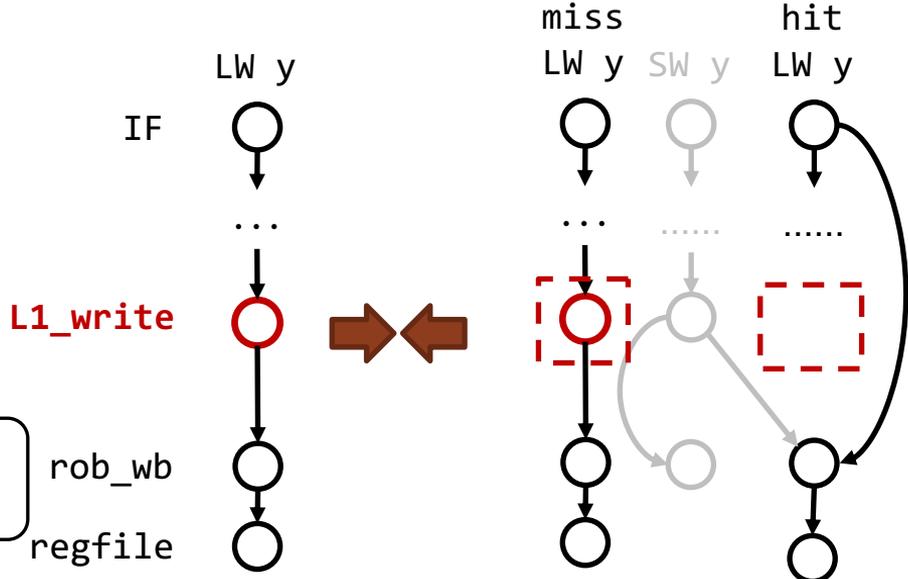
- Each instruction type updates the **same state elements exactly once** every time it executes.
- rtl2µspec cannot recognize:
 - > “Sometimes-updated” flops
 - > “Repeatedly-updated” flops

Resolving the single execution path assumption: recognizing sometimes-updated flops with write-enable analysis

SVA property: Does there **exist** an execution of **LW** that induces a **value** change on **L1_line[0].data** ✓



yes
Conclude: LW always updates L1_line[0].data



Resolving the single execution path assumption: recognizing sometimes-updated flops with write-enable analysis

Challenge: updating a flip-flop does not imply a value change.

SVA property: Does there **exist** an execution of **LW** that induces a **value change** on **L1_line[0].data**?

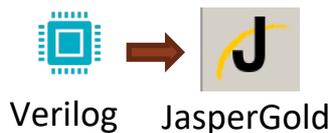


Verilog JasperGold

Solution: conduct static analysis of RTL netlist to derive write enables for flip-flops.

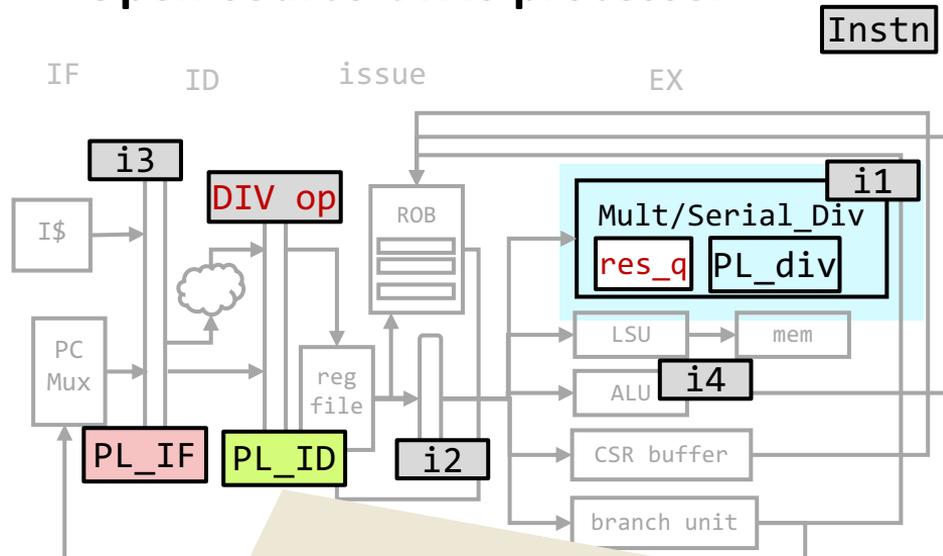
SVA property: Does the execution of **LW** **always/ever** cause **L1_line[0].data**'s **write-enable** to be asserted?

Resolving the single execution path assumption: recognizing sometimes-updated flops with write-enable analysis



SVA property: Does the execution of **DIV** always/ever cause `res_q`'s write enable to be asserted when DIV is at PL_DIV?

Open-source CVA6 processor¹



Performing Location: A design region which consists of an instruction identifier, optional instruction tracking logic, and a set of datapath registers. For an instruction to update a particular datapath register it must be residing in its associated performing location.

¹Zaruba et al. <https://github.com/openhwgroup/cva6>

Resolving the single execution path assumption: recognizing repeatedly-updated flops with performing location (PL)

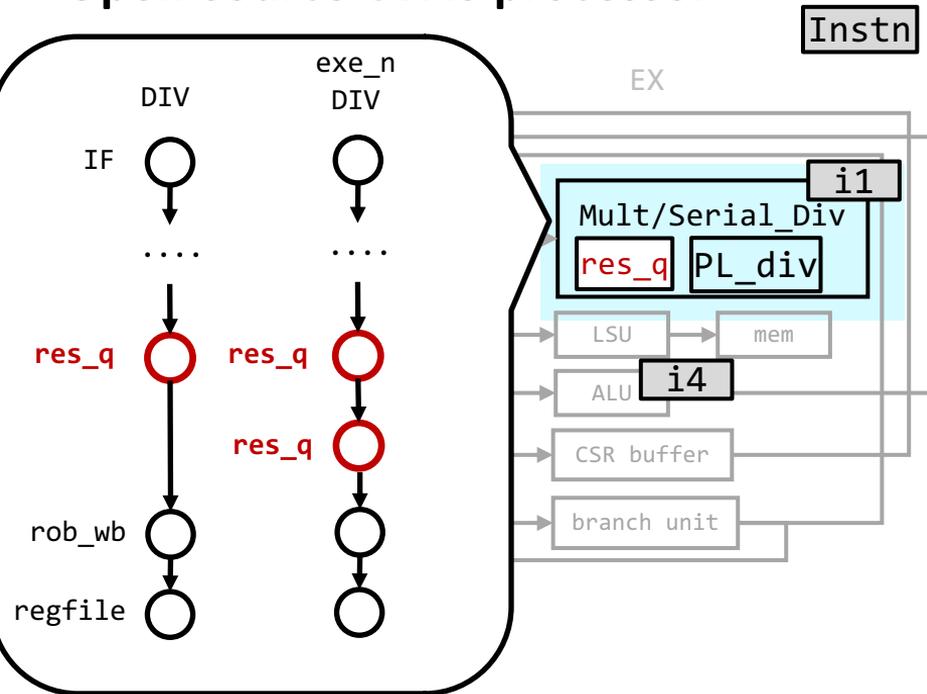
SVA property: Does the execution of DIV always/ever cause res_q's write enable to be asserted?



Verilog JasperGold

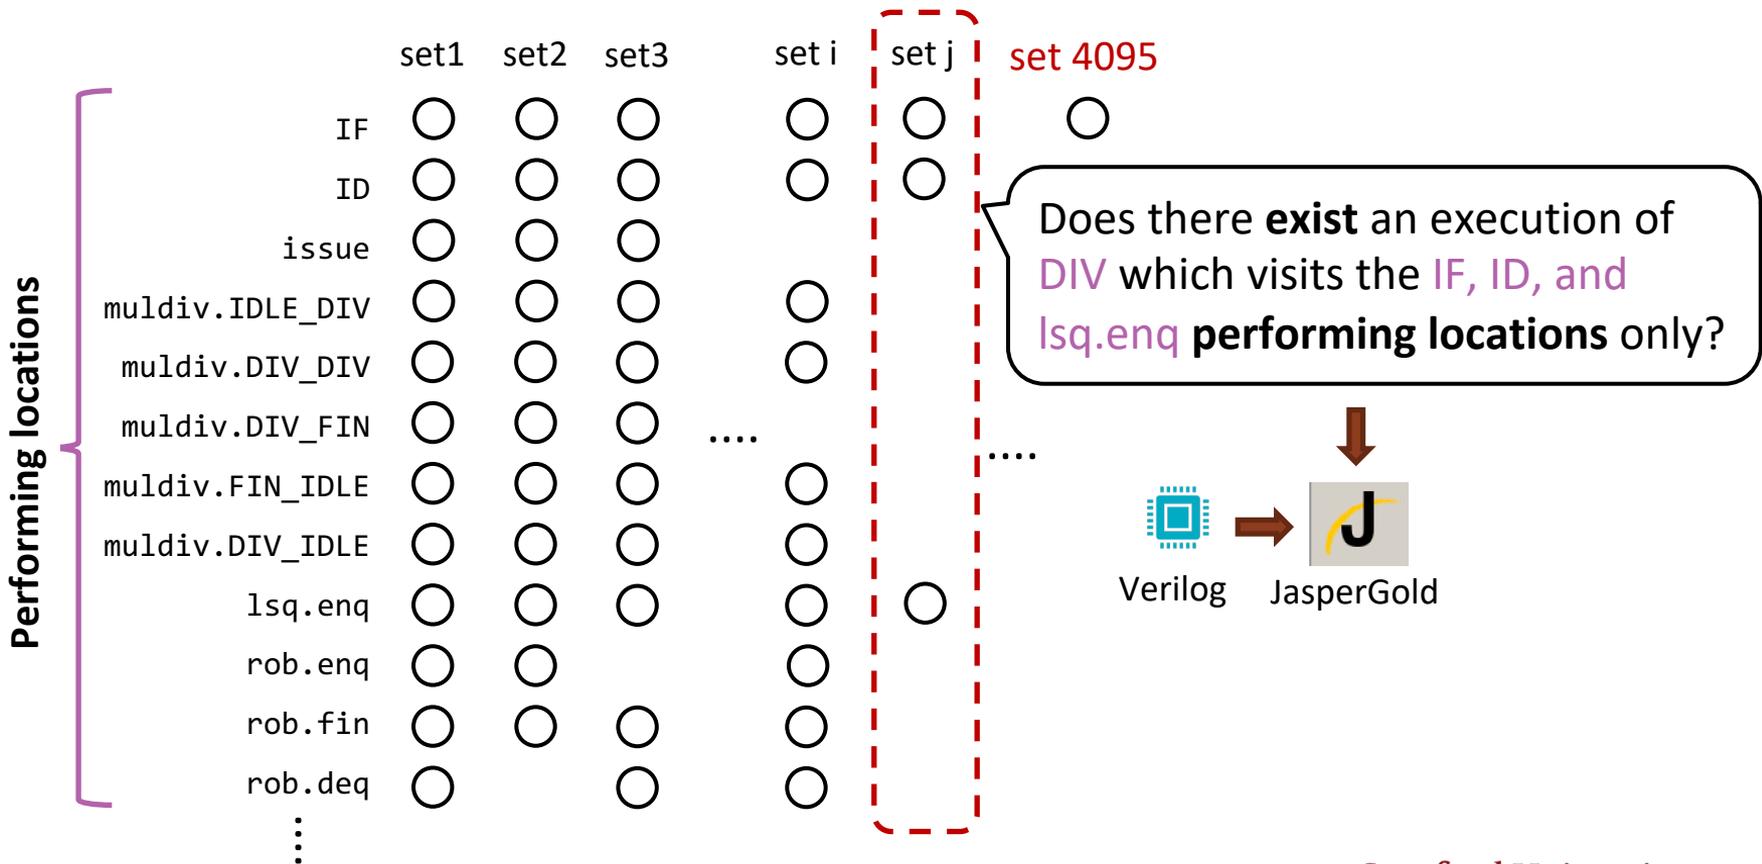
SVA property: Does the execution of DIV always/ever cause res_q's write enable to be asserted when DIV is at PL_DIV ?

Open-source CVA6 processor¹

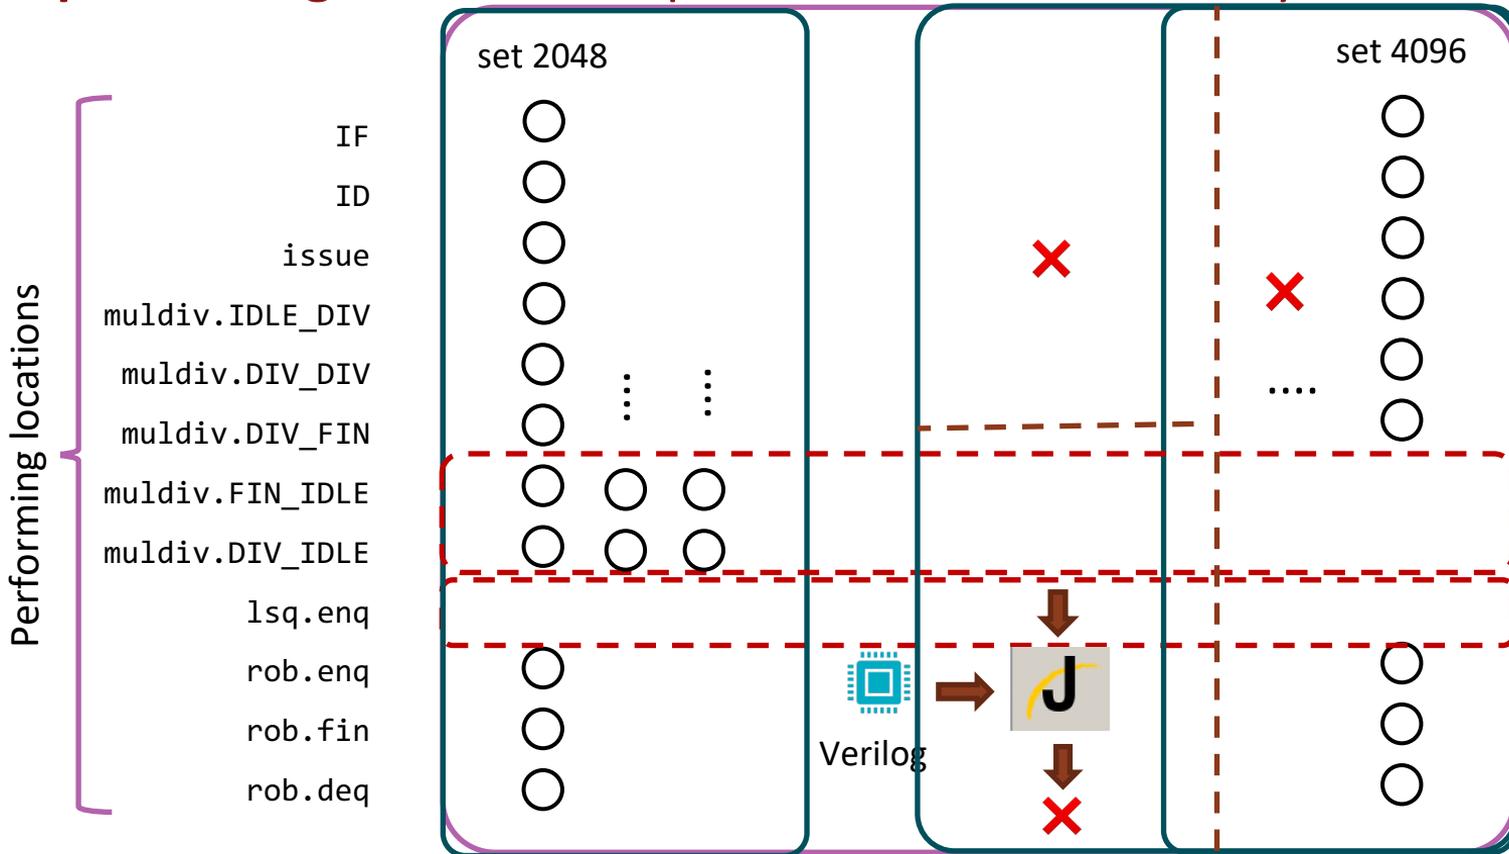


¹Zaruba et al. <https://github.com/openhwgroup/cva6>

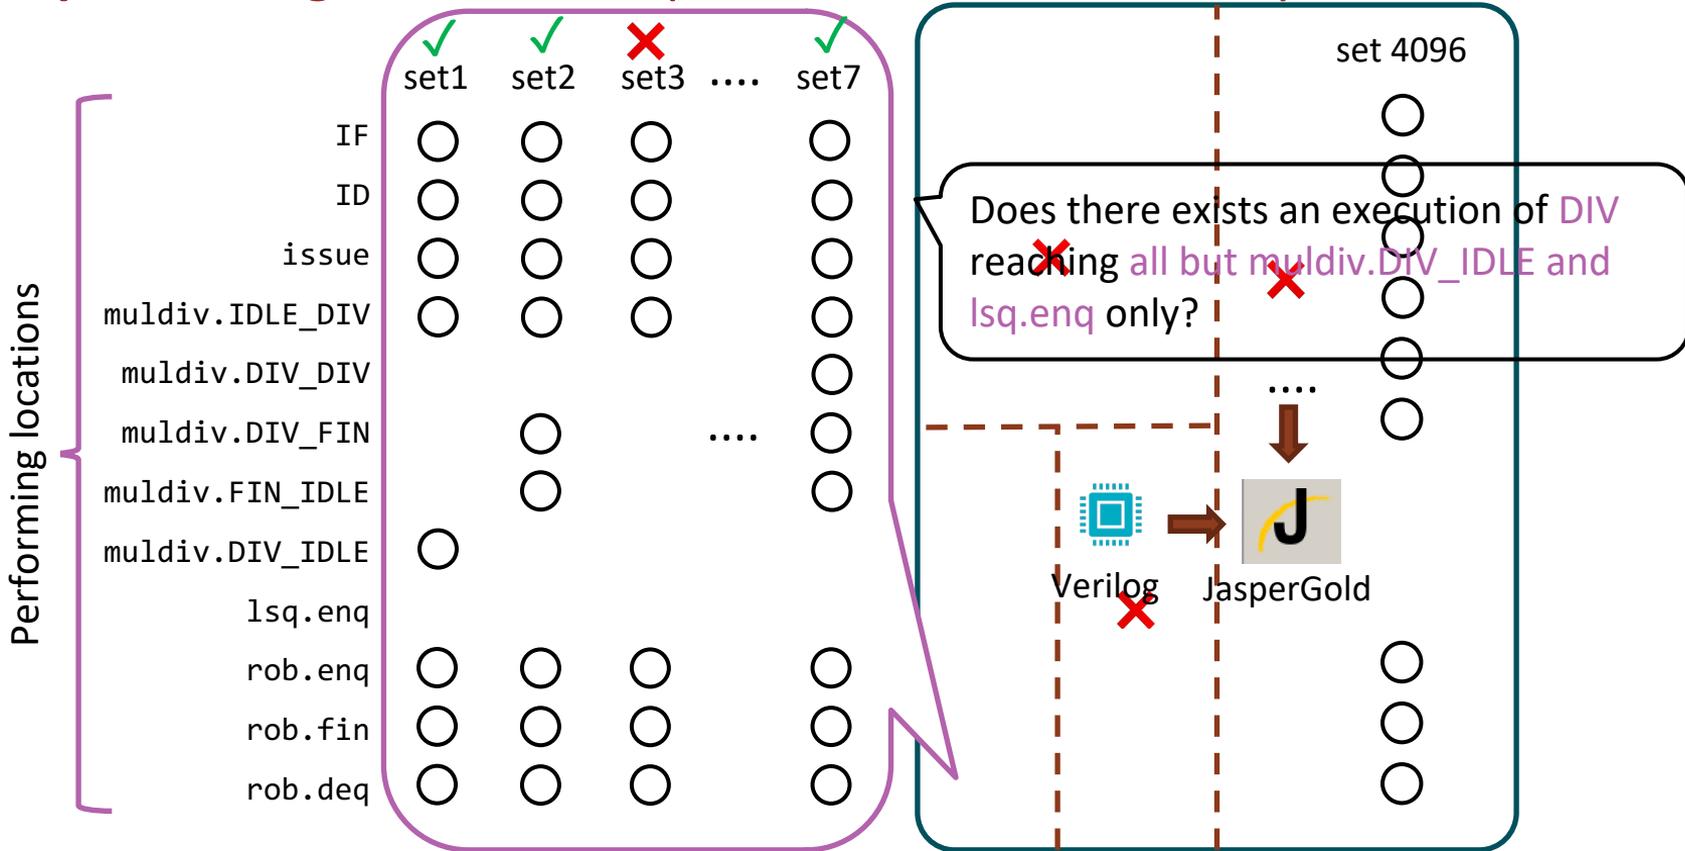
Synthesize coarse-grained execution paths: a set of visited performing locations + a partial order in which they are visited



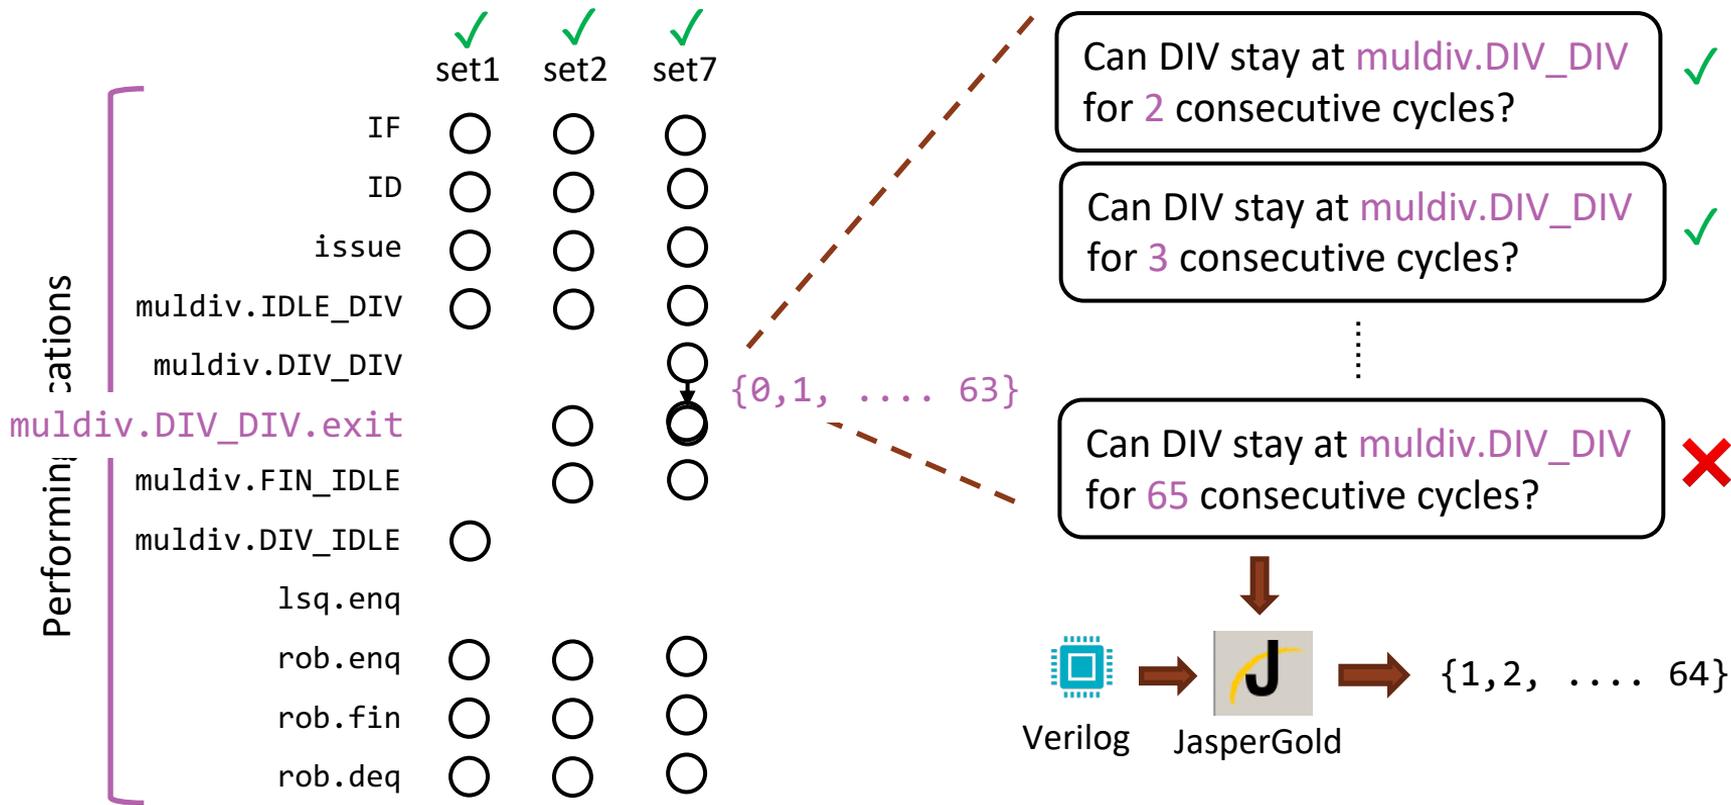
Synthesize coarse-grained execution paths: a set of visited performing locations + a partial order in which they are visited



Synthesize coarse-grained execution paths: a set of visited performing locations + a partial order in which they are visited



Synthesize coarse-grained execution paths: a set of visited performing locations + a partial order in which they are visited



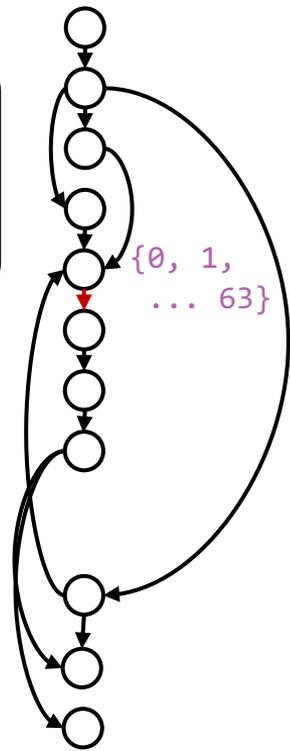
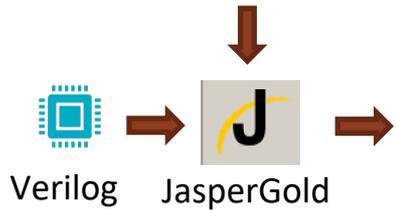
Synthesize coarse-grained execution paths: a set of visited performing locations + a **partial order** in which they are visited

Performing locations

| | set1 | set2 | set7 |
|---------------------|-----------------------|-----------------------|-----------------------|
| IF | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| ID | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| issue | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| muldiv.IDLE_DIV | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| muldiv.DIV_DIV | | | <input type="radio"/> |
| muldiv.DIV_DIV.exit | | | <input type="radio"/> |
| muldiv.DIV_FIN | | <input type="radio"/> | <input type="radio"/> |
| muldiv.FIN_IDLE | | <input type="radio"/> | <input type="radio"/> |
| muldiv.DIV_IDLE | <input type="radio"/> | | |
| lsq.enq | | | |
| rob.enq | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| rob.fin | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| rob.deq | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |

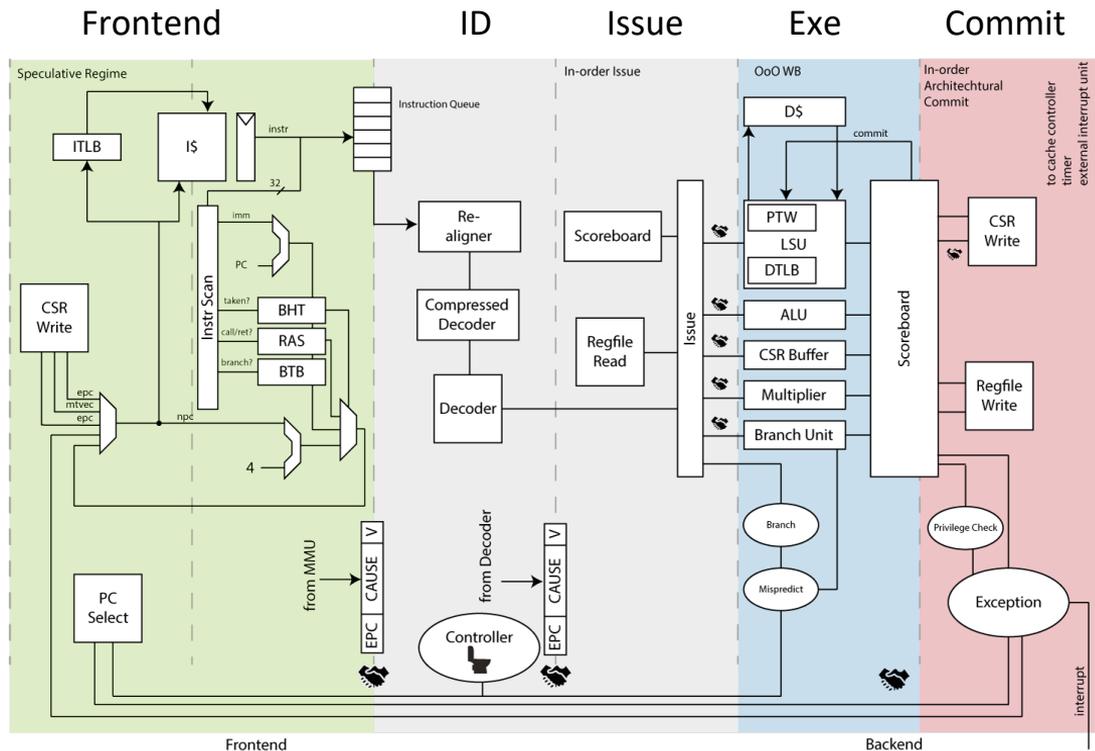
{0, 1, ... 63}

SVA property: For an ordered pair of nodes (u, v), does DIV always reach u before it reaches v?



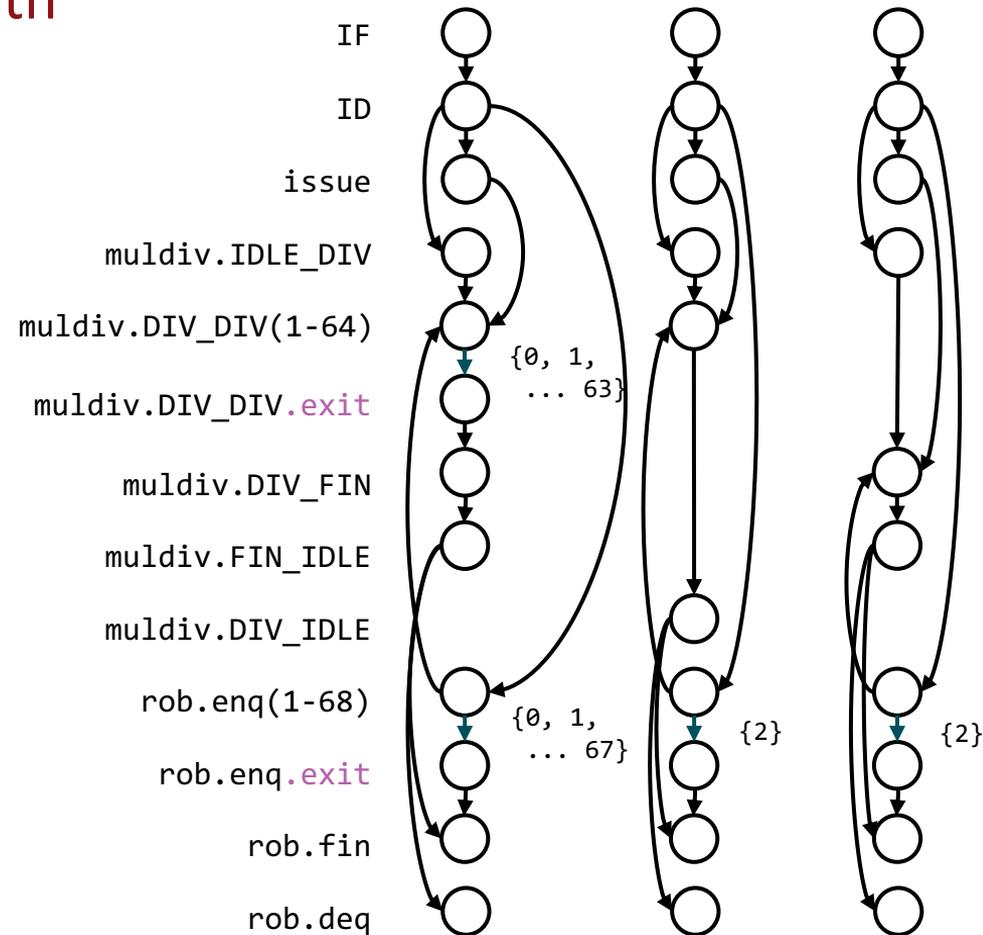
Open-source RISC-V CVA6 processor case study with proof-of-concept TransmitSynth methodology and tool

- RV64I, M, A, and C
- Single in-order issue
- Out-of-order write-backs
- Speculation
- DIV executes variably as function of its operands



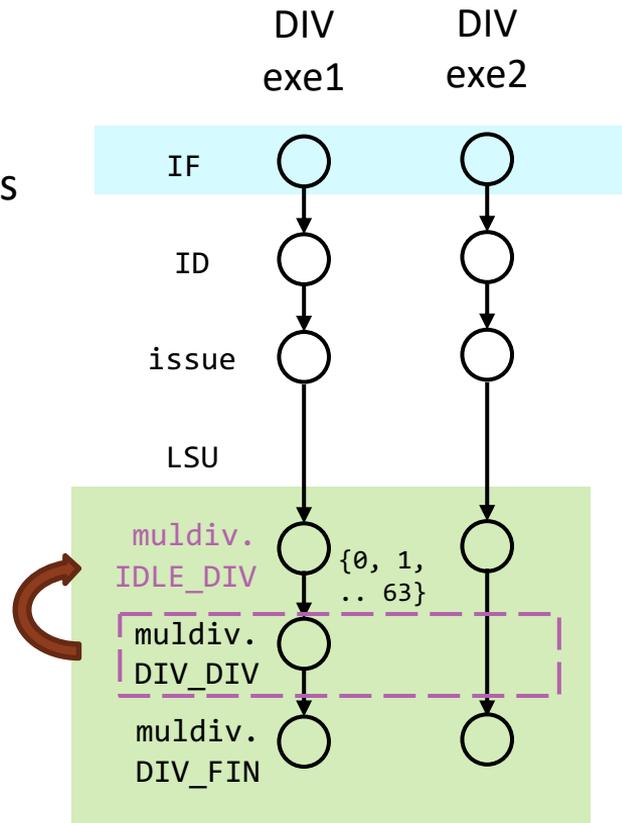
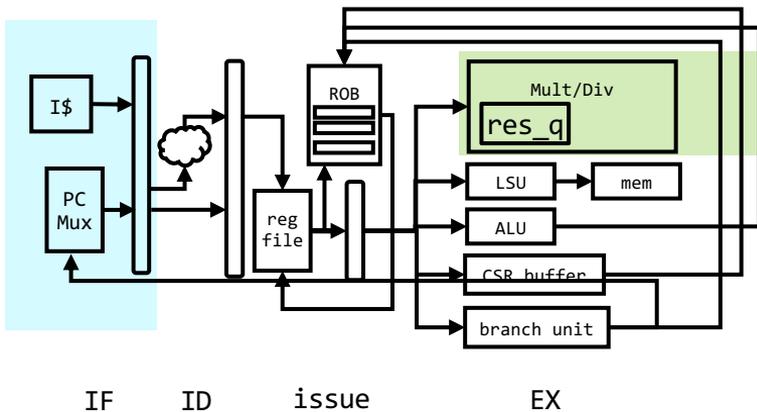
Discovering Transmitters with TransmitSynth

- DIV has a total of **66 execution paths!**
- **Runtime** to synthesize paths with TransmitSynth: **96 minutes** of serial (**parallelizable**) proof time
- Note that CVA6 has **4.54x** more flops and **16.2x** more gates than the multi-V-scale



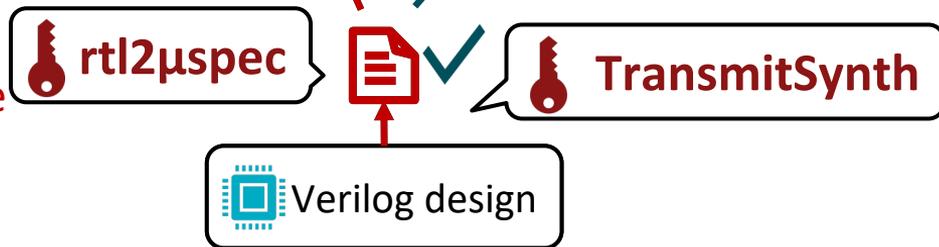
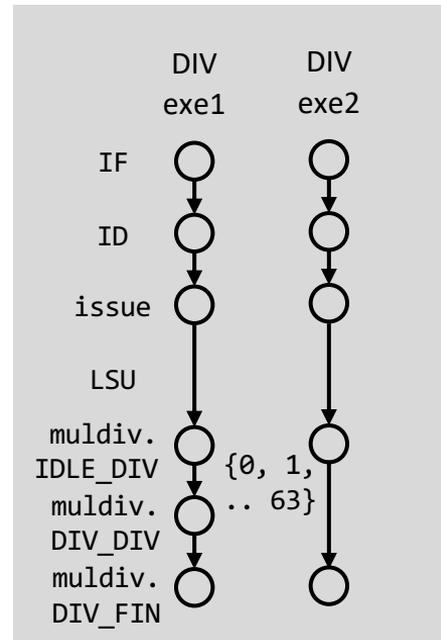
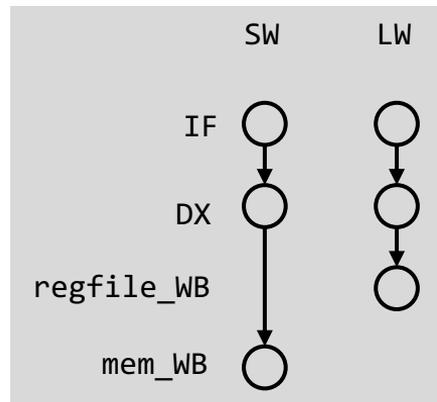
Ongoing work

- What information is leaked?
 - › What causes the different execution
 - › Dataflow behavior on some performing locations
- Finer-grained execution path
 - › What are updated during a performing location



Conclusions: Automated synthesis of hardware-software contracts from RTL is feasible

- **rtl2 μ spec¹**: Automated synthesis of μ spec models from RTL
 - › Synthesized a μ spec model of the **open-source RISC-V multi-V-scale** processor in **6.84 minutes serial proof time**
- **TransmitSynth**: Automated synthesis of security contracts from RTL via identification of transmit instructions
 - › Synthesized **66 paths** for DIV instructions on the **open-source RISC-V CVA6** processor in **96 minutes serial proof time**



¹<https://github.com/yaohsiaopid/rtl2uspec>

Thanks!

yaohsiao@stanford.edu

This presentation and recording belong to the authors.
No distribution is allowed without the authors' permission.