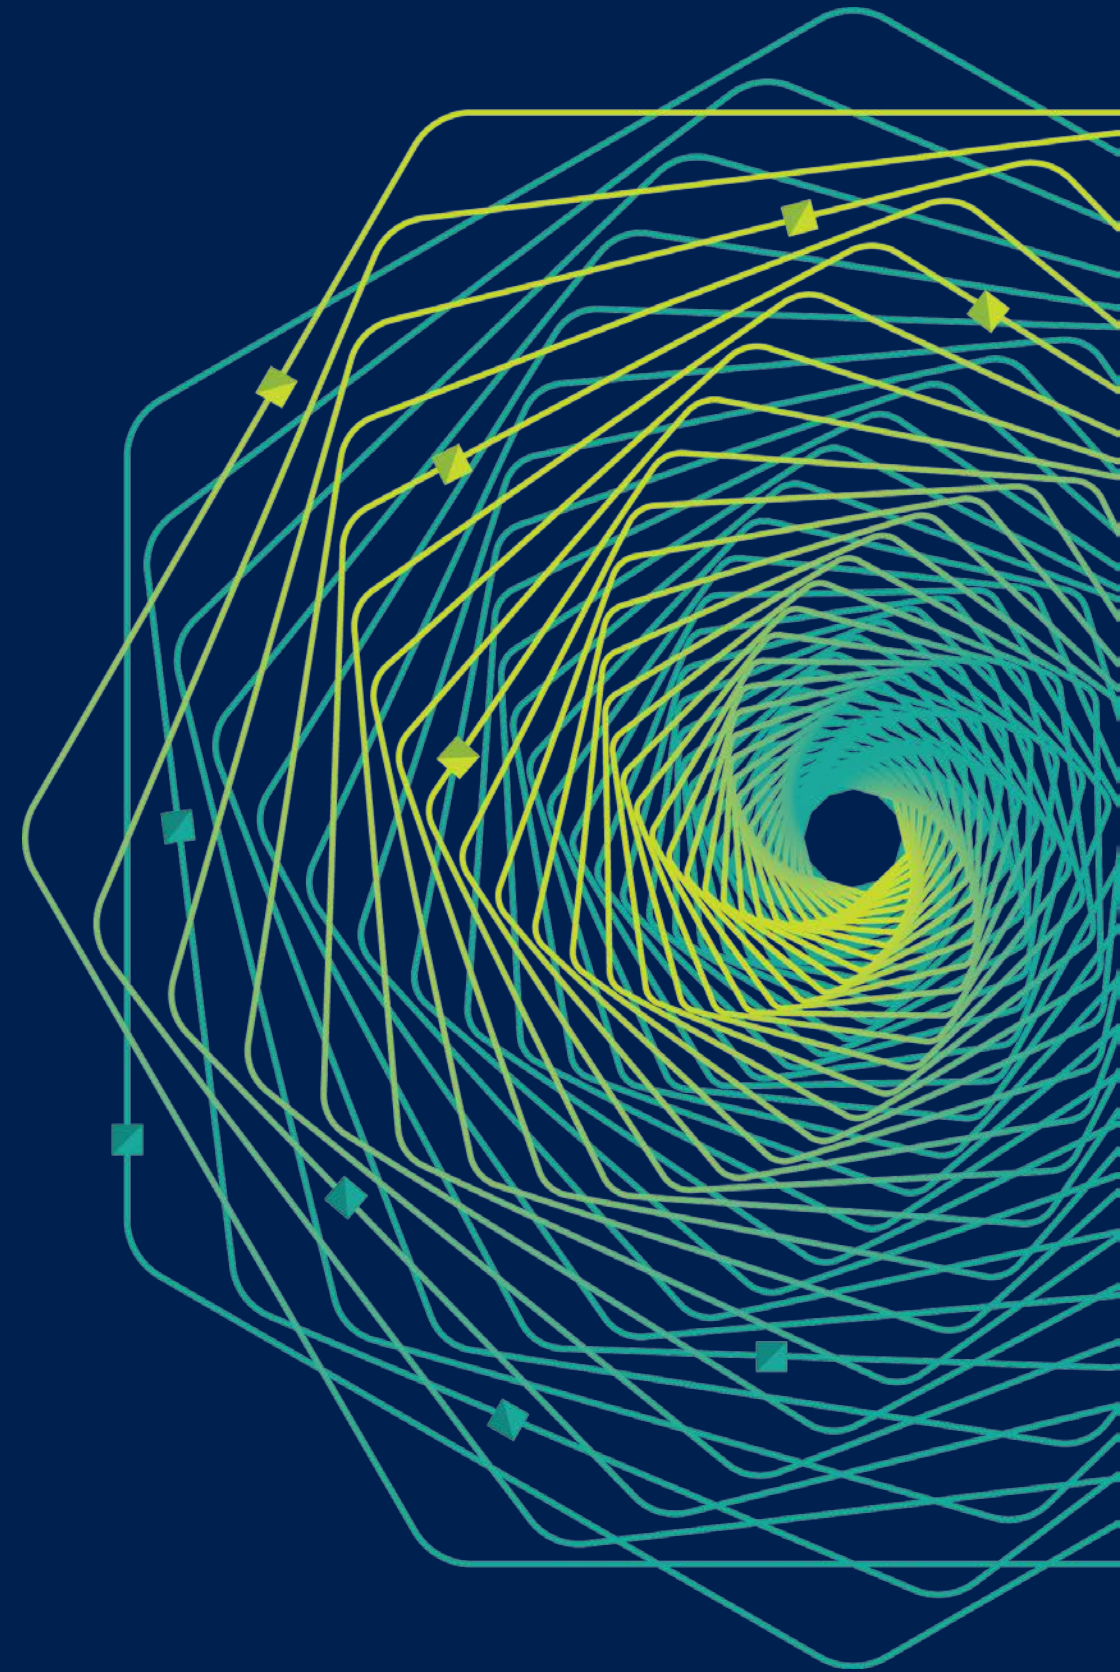


# Research Faculty Summit 2018

Systems | Fueling future disruptions





# Hardware-Aware Security Verification and Synthesis

Margaret Martonosi

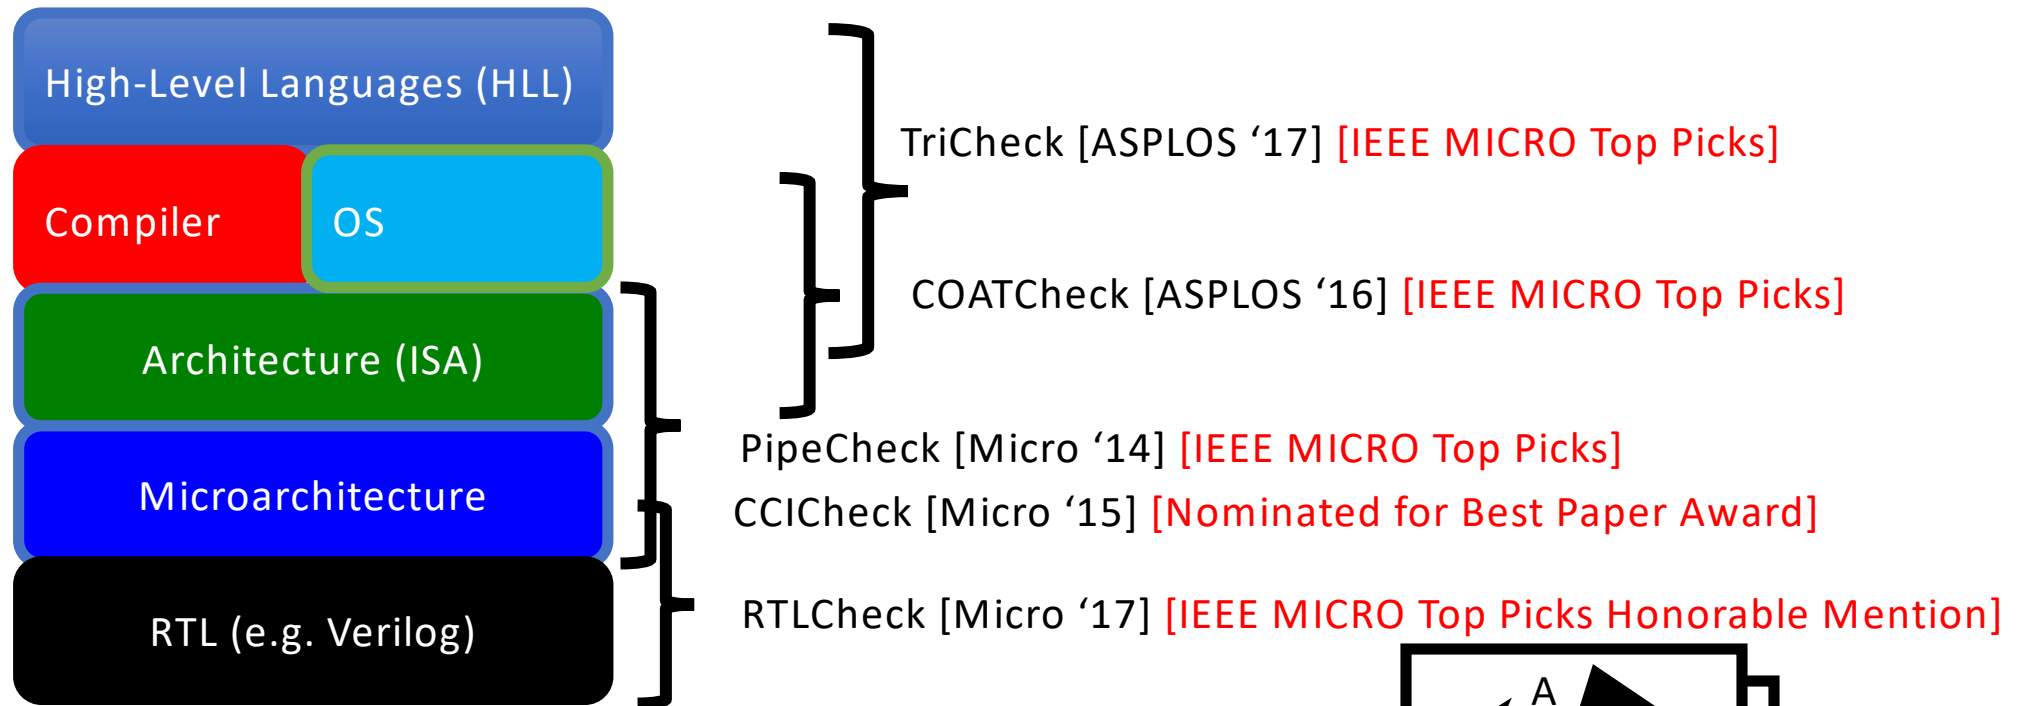
H. T. Adams '35 Professor

Dept. of Computer Science

Princeton University


*Joint work with Caroline Trippel, Princeton CS PhD student  
and Dr. Daniel Lustig, NVIDIA*

# The Check Suite: An Ecosystem of Tools For Verifying Memory Consistency Model Implementations



## Our Approach

- Axiomatic specifications -> Happens-before graphs
- Check **Happens-Before Graphs** via **Efficient SMT solvers**
  - Cyclic => A->B->C->A... **Can't happen**
  - Acyclic => Scenario is **observable**




# Check: Formal, Axiomatic Models and Interfaces

```
Axiom "PO_Fetch":  
forall microops "i1",  
forall microops "i2",  
SameCore i1 i2 /\ ProgramOrder i1 i2 =>  
  AddEdge ((i1, Fetch), (i2, Fetch), "P0").
```

```
Axiom "Execute_stage_is_in_order":  
forall microops "i1",  
forall microops "i2",  
SameCore i1 i2 /\  
  EdgeExists ((i1, Fetch), (i2, Fetch)) =>  
  AddEdge ((i1, Execute), (i2, Execute), "PPO").
```

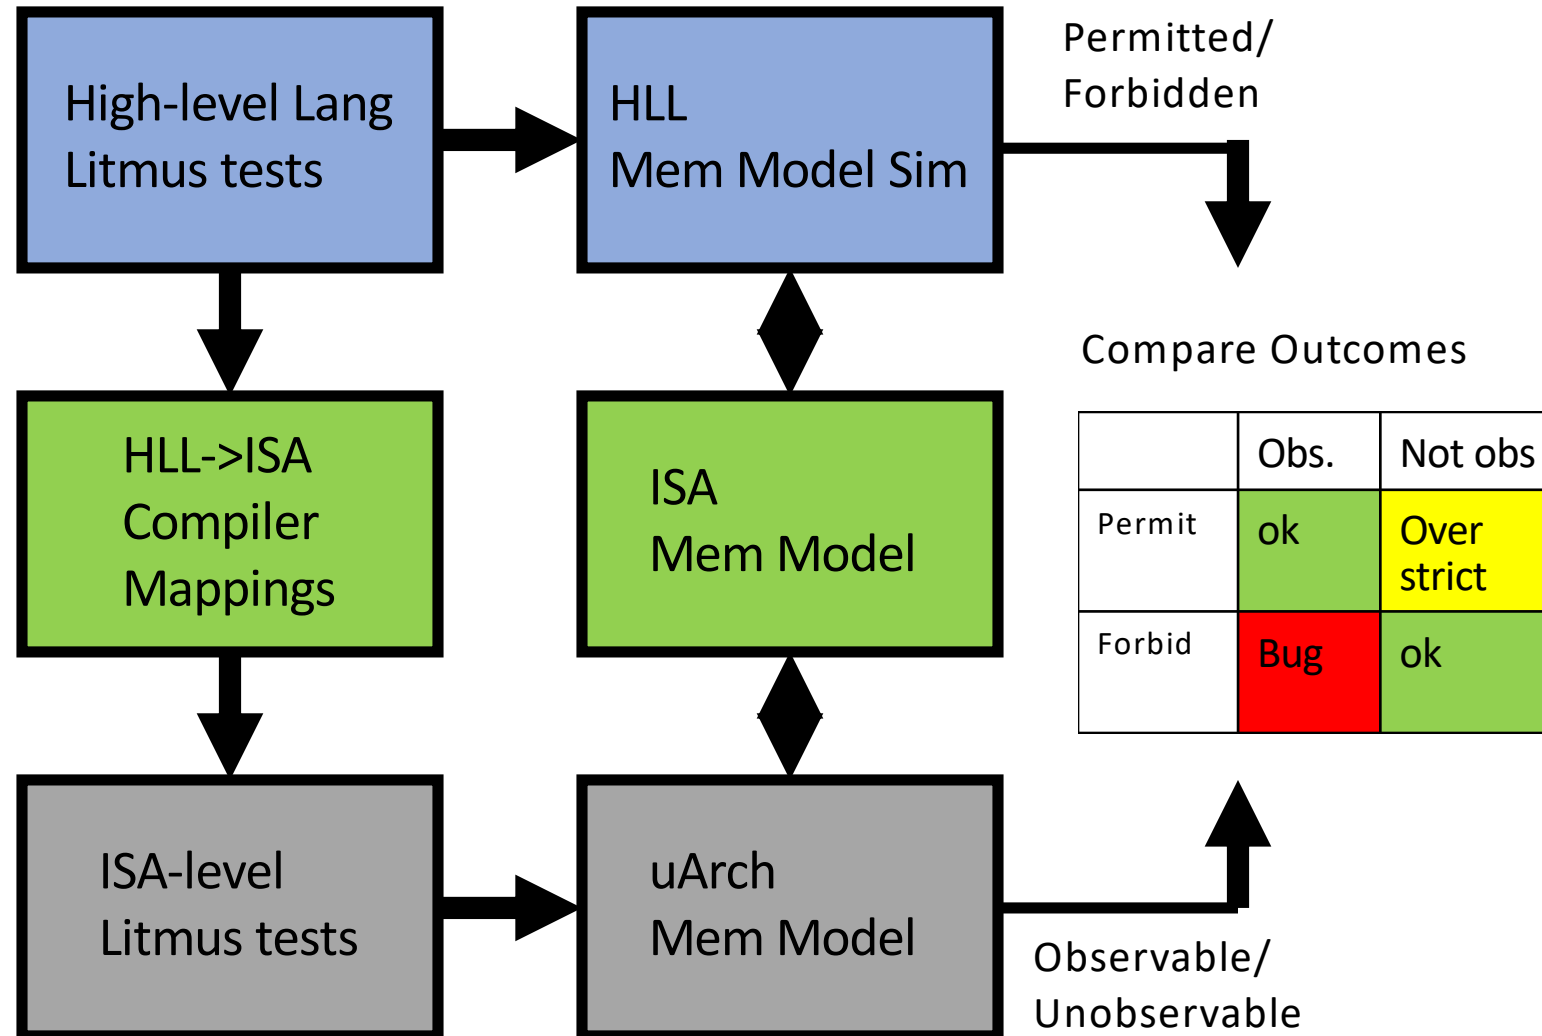
**Microarchitecture Specification in  
*μSpec DSL***

**Exhaustive consideration of  
all possible executions**

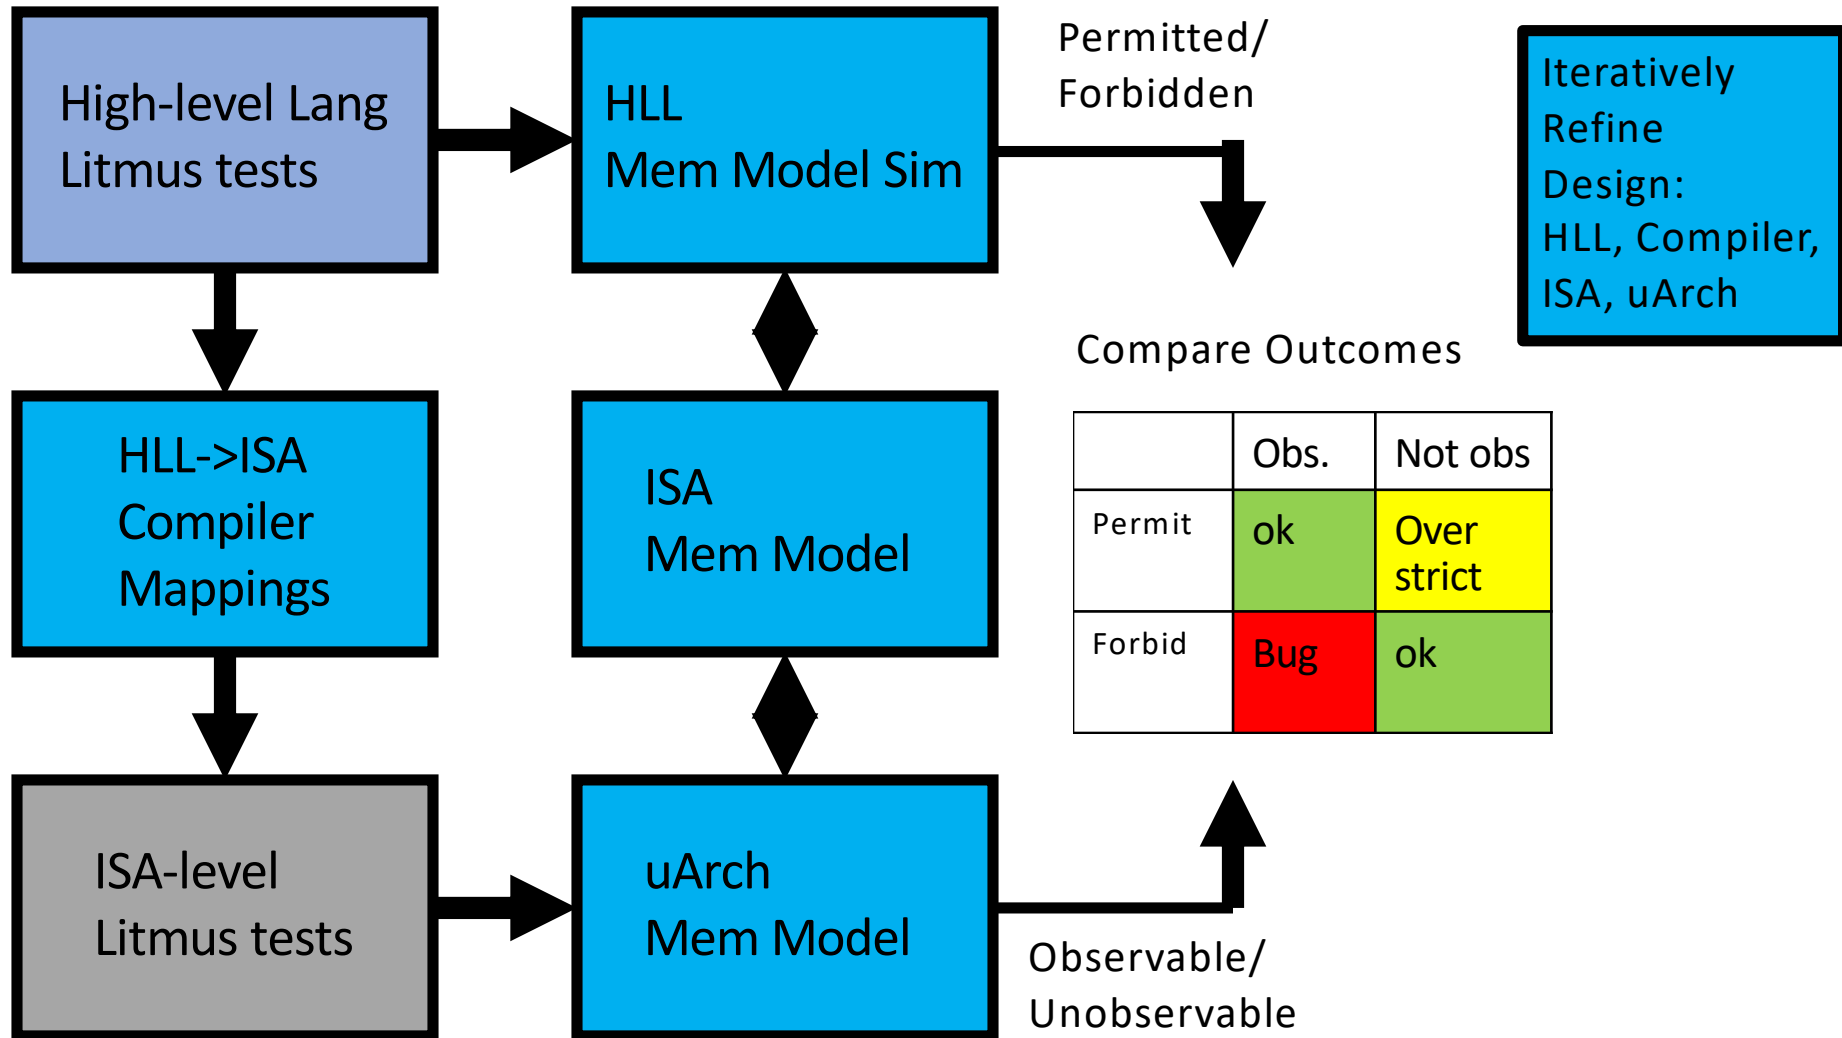


**Microarchitectural happens-before ( $\mu hb$ ) graphs**

# TriCheck Framework: Verifying Memory Event Ordering from Languages to Hardware



# TriCheck Framework: Verifying Memory Event Ordering from Languages to Hardware



# TriCheck Framework: RISC-V Case Study

1701 C11 Programs

High-level Lang  
Litmus tests

HLL  
Mem Model Sim

Permitted/  
Forbidden

Base RISC-V ISA:  
144 buggy outcomes  
Base+Atomics:  
221 buggy outcomes

Conclusion: Draft RISC-V spec could not serve as a legal C11 compiler target.

Status: RISC-V Memory Model Working Group formed to address these issues. Just voted to ratify new, improved RISC-V memory model.

are Outcomes

Obs.	Not obs
ok	Over strict
Bug	ok

ISA-level  
Litmus tests

uArch  
Mem Model

Observable/

7 Distinct RISC-V Implementations (All abide by RISC-V specifications, but vary in reordering / performance)

# CheckMate: From Memory Consistency Models to Security

January 2018:  
Spectre & Meltdown

Well-known cache  
side-channel attack

{ Flush+Reload

+

Widely-used  
hardware feature

{ Speculation

New exploit

=

2 new attacks



SHARE



**TRIPLE MELTDOWN: HOW SO MANY RESEARCHERS FOUND A 20-YEAR-OLD CHIP FLAW AT THE SAME TIME**



# Attack Discovery & Synthesis: What We Would Like

1. Specify  
system to study

Formal interface and specification of  
given system implementation

2. Specify attack  
pattern

E.g. Subtle event sequences during  
program's execution

3. Synthesis

Either output synthesized attacks. Or  
determine that none are possible

# Attack Discovery & Synthesis:

## CheckMate TL;DR

1. Specify system to study

2. Specify attack pattern

3. Synthesis

- **What we did:** Developed a tool to do this, based on the uHB graphs from previous sections.
- **Results:** Automatically synthesized Spectre and Meltdown, as well as two new distinct exploits and many variants.

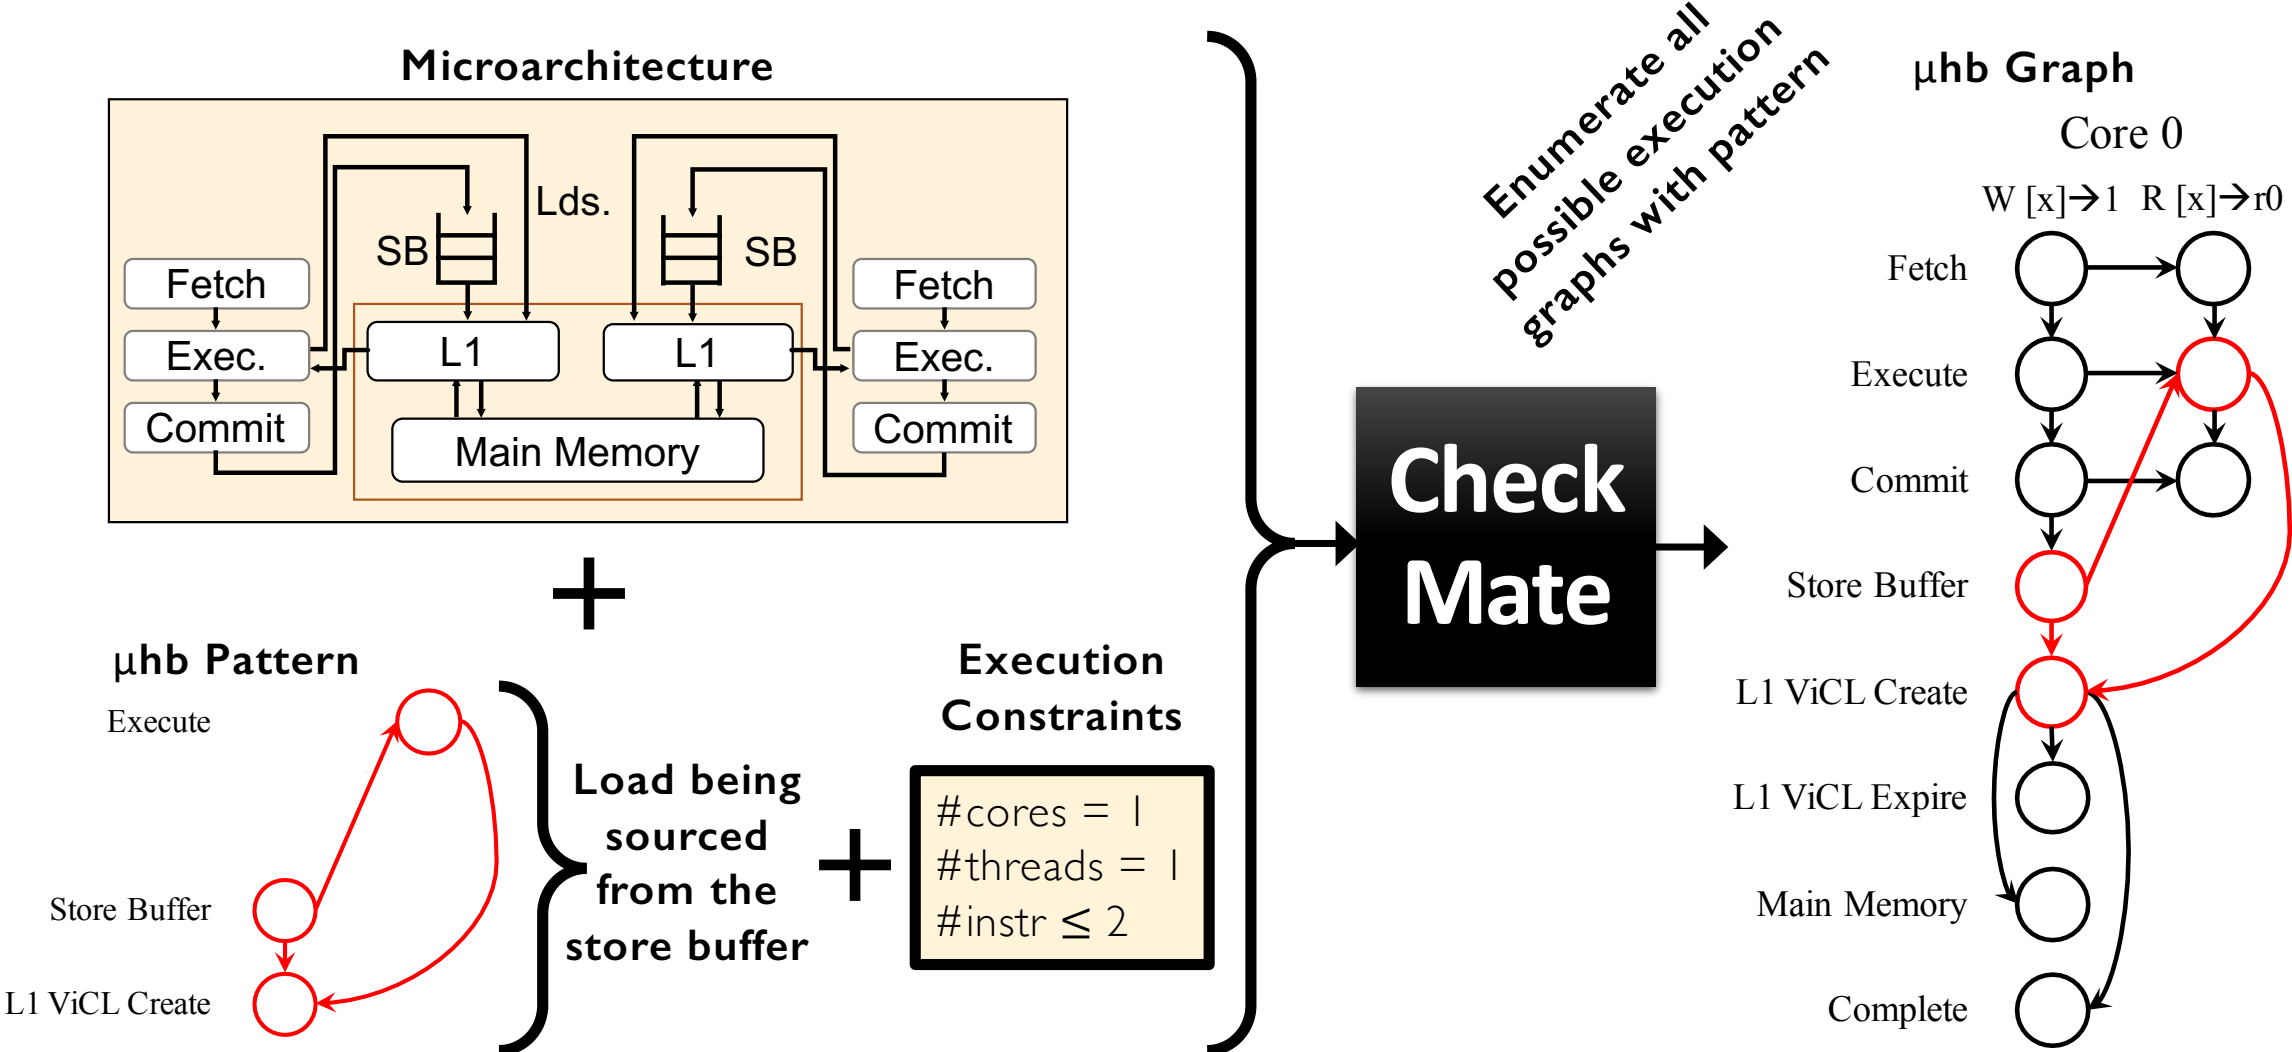
In more  
detail...



# CheckMate Methodology

1. Frame classes of attacks as patterns of event interleavings?
  - > Essentially a snippet out of a happens-before graph
2. Specify hardware using uSpec axioms
  - > Determine if attack is realizable on a given hardware implementation

# Microarchitecture-Aware Program Synthesis



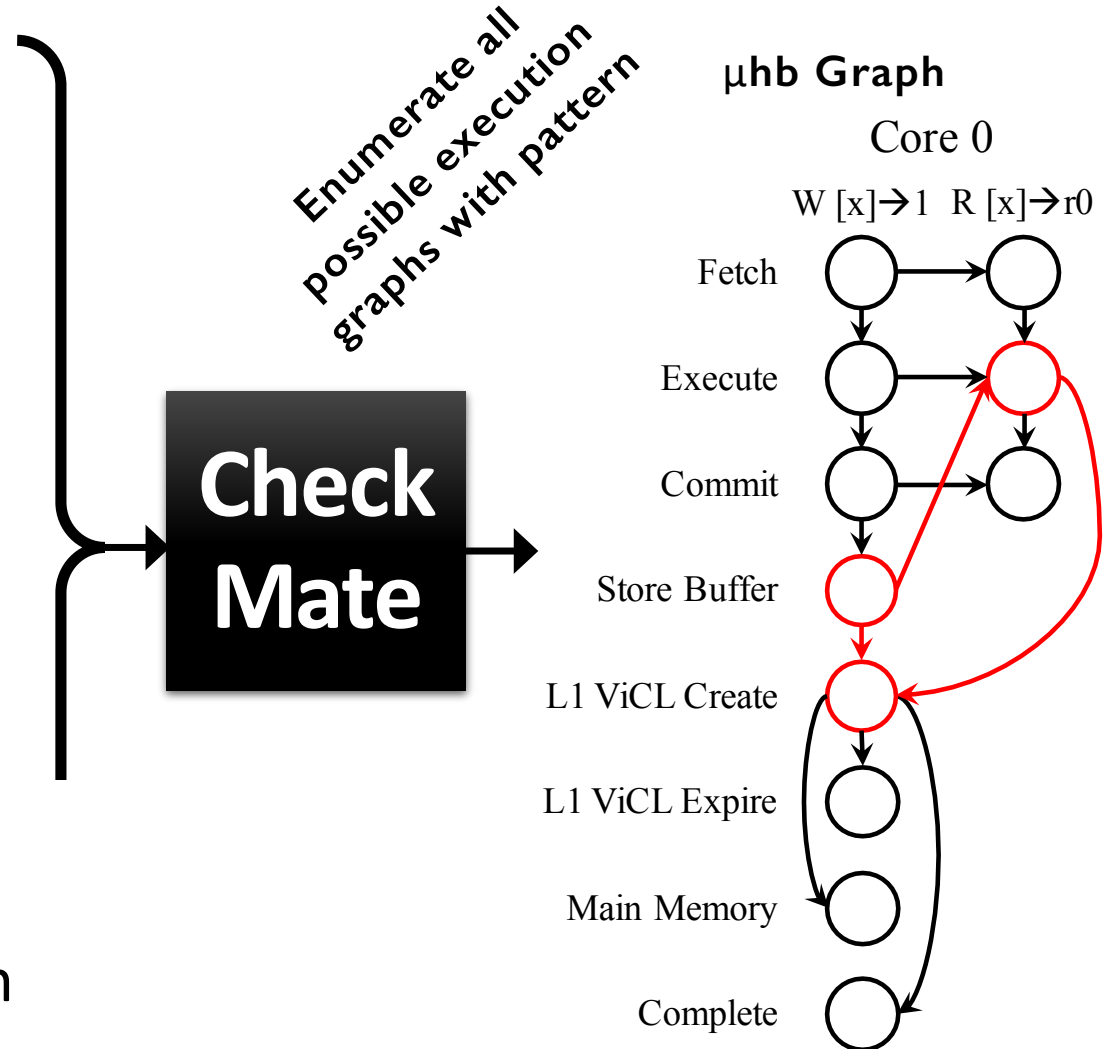
# Microarchitecture-Aware Program Synthesis

## Microarchitecture Specification

```
Axiom "PO_Fetch":  
forall microops "i1",  
forall microops "i2",  
SameCore i1 i2 /\ ProgramOrder i1 i2 =>  
  AddEdge ((i1, Fetch), (i2, Fetch), "P0").  
  
Axiom "Execute_stage_is_in_order":  
forall microops "i1",  
forall microops "i2",  
SameCore i1 i2 /\  
  EdgeExists ((i1, Fetch), (i2, Fetch)) =>  
  AddEdge ((i1, Execute), (i2, Execute), "PPO").
```

**Prior Check tools work  
addresses many of these issues**

- SW/OS/HW events and locations
- SW/OS/HW ordering details
- Hardware optimizations, e.g., speculation
- Processes and resource-sharing
- Memory hierarchies and cache coherence protocols

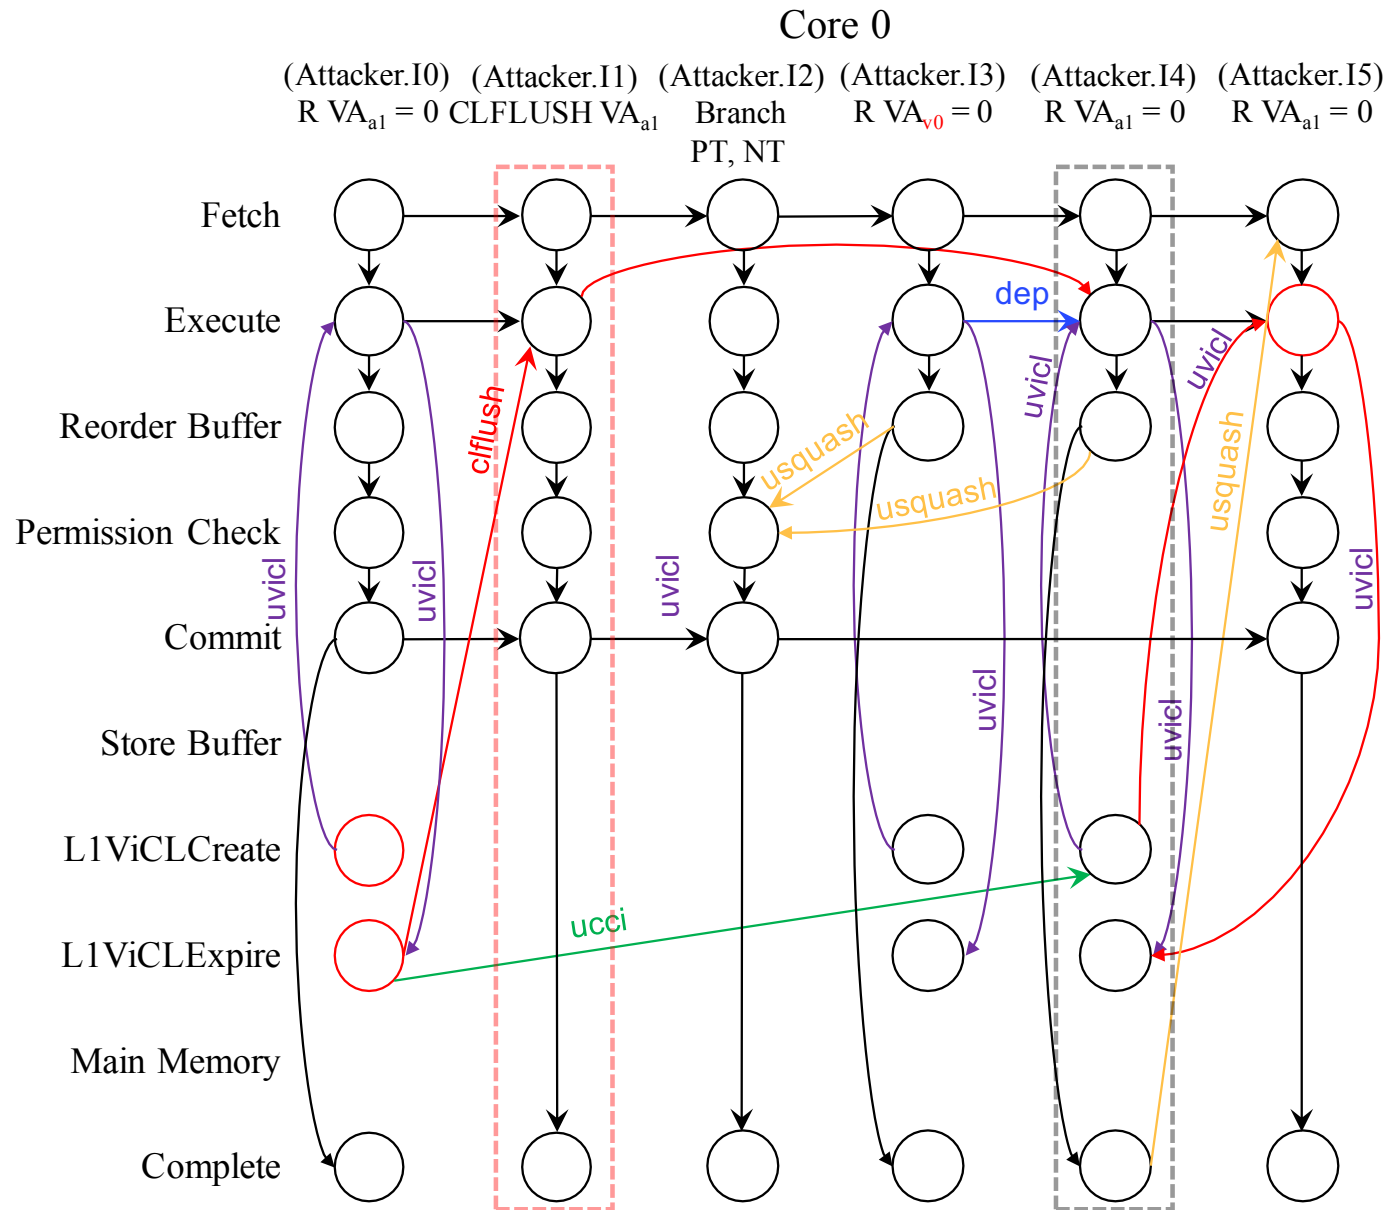


# Relational Model Finding (RMF): A Natural Fit for Security Litmus Test Synthesis

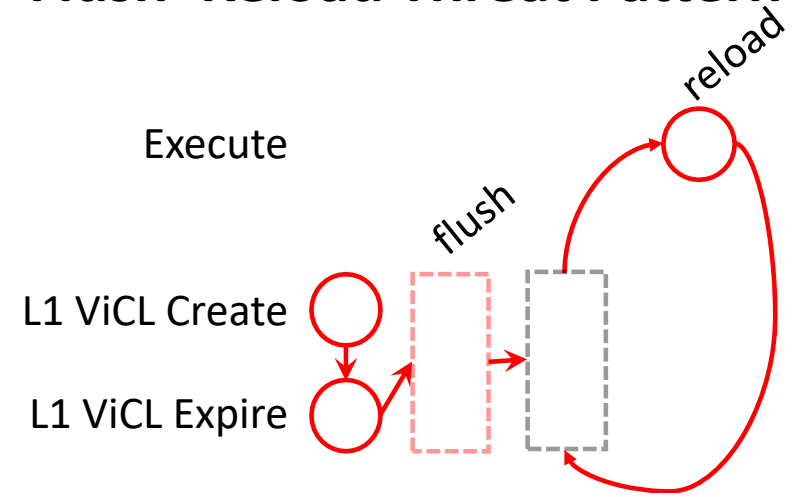
A relational model is a set of constraints on an abstract system (for CheckMate, a  $\mu$ hb graph) of.

- Set of abstract objects (for CheckMate,  $\mu$ hb graph nodes)
- Set of N-dimensional relations (for example., 2D  $\mu$ hb graph edges relation connecting 2 nodes)
- For CheckMate, the constraints are a  $\mu$ hb pattern of interest
- RMF attempts to find and satisfying “instance” (or  $\mu$ hb graph)
- Implementation: Alloy DSL maps RMF problems onto Kodkod model-finder, which in turn uses off-the-shelf SAT solvers
- CheckMate Tool maps  $\mu$ spec HW/OS spec to Alloy

# Spectre (Exploits Speculation)



## Flush+Reload Threat Pattern




## Spectre Security Litmus Test


Initial conditions: [x]=0, [y]=0	
Attacker T0	
R [VA <sub>a1</sub> ] → 0	
CLFLUSH [VA <sub>a1</sub> ]	◀ Flush
Branch → PT,NT	
R [VA <sub>v0</sub> ] → r1	
R [f(r1)=VA <sub>a1</sub> ] → 0	
R [VA <sub>a1</sub> ] → 0	◀ Reload



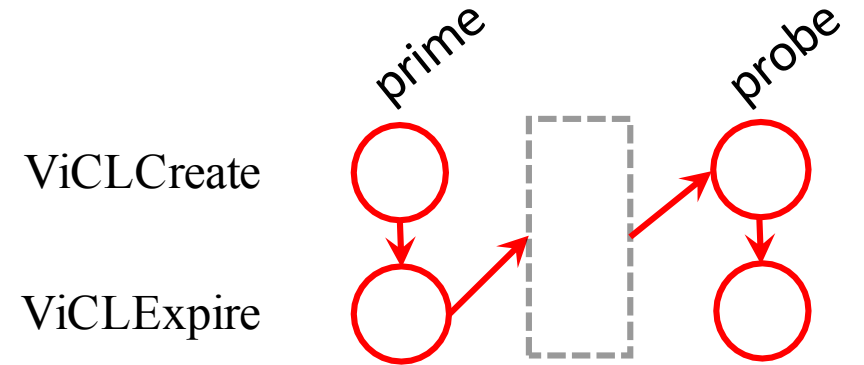
# Prime&Probe Attack Pattern: Synthesizing MeltdownPrime & SpectrePrime



# SpectrePrime uhb Graph



## Prime+Probe Threat Pattern



## Spectre Security Litmus Test

Initial conditions: [x]=0, [y]=0	
Attacker T0	Attacker T0
R [VA <sub>a1</sub> ] → 0	R [VA <sub>a1</sub> ] → 0
Branch → PT, NT	
R [VA <sub>v0</sub> ] → r1	
W [f(r1)=VA <sub>a1</sub> ] → 0	
	R [VA <sub>a1</sub> ] → 0

← Prime (pointing to the first R [VA<sub>a1</sub>] → 0)

← Probe (pointing to the second R [VA<sub>a1</sub>] → 0)

# Overall Results: What exploits get synthesized? And how long does it take?

Exploit Pattern	#Instrs (RMF Bound)	Output Attack	Minutes to synthesize 1 <sup>st</sup> exploit	Minutes to synthesize all exploits	#Exploits Synthesized
Flush +Reload	4	Traditional Flush+Reload	6.7	9.7	70
	5	Meltdown	27.8	59.2	572
	6	Spectre	101.0	198.0	1144
Prime +Probe	3	Traditional Prime+Probe	5.4	6.7	12
	4	MeltdownPrime	17.0	8.2	24
	5	SpectrePrime	71.8	76.7	24

# CheckMate: Takeaways

- New Variants reported: SpectrePrime and MeltdownPrime
  - Speculative cacheline invalidations versus speculative cache pollution
  - Software mitigation is the same as for Meltdown & Spectre
- Key overall philosophy:
  - Event ordering in security exploit patterns aligns strongly with MCM analysis
  - Move from ad hoc analysis to formal automated synthesis.
  - Span software, OS, and hardware for holistic hardware-aware analysis

[Trippel, Lustig, Martonosi. <https://arxiv.org/abs/1802.03802>]

[Trippel, Lustig, Martonosi. MICRO-51. October, 2018. [http://check.cs.princeton.edu/papers/ctrippel\\_MICRO51.pdf](http://check.cs.princeton.edu/papers/ctrippel_MICRO51.pdf)]

# Acknowledgements

- CheckMate: **Caroline Trippel** (Princeton CS PhD student) and Dan Lustig (NVIDIA)
- Funding: NSF, NVIDIA Graduate Fellowship
- Check Tools, additional co-authors: **Yatin Manerkar**, Abhishek Bhattacharjee, Michael Pellauer, Geet Sethi

Me: <http://www.princeton.edu/~mrm>

Group Papers: <http://mrmgroup.cs.princeton.edu>

Check and CheckMate Tools: <http://check.cs.princeton.edu>

Thank you!

