

# TransForm: Formally Specifying Transistency Models and Synthesizing Enhanced Litmus Tests

**Naorin Hossain**

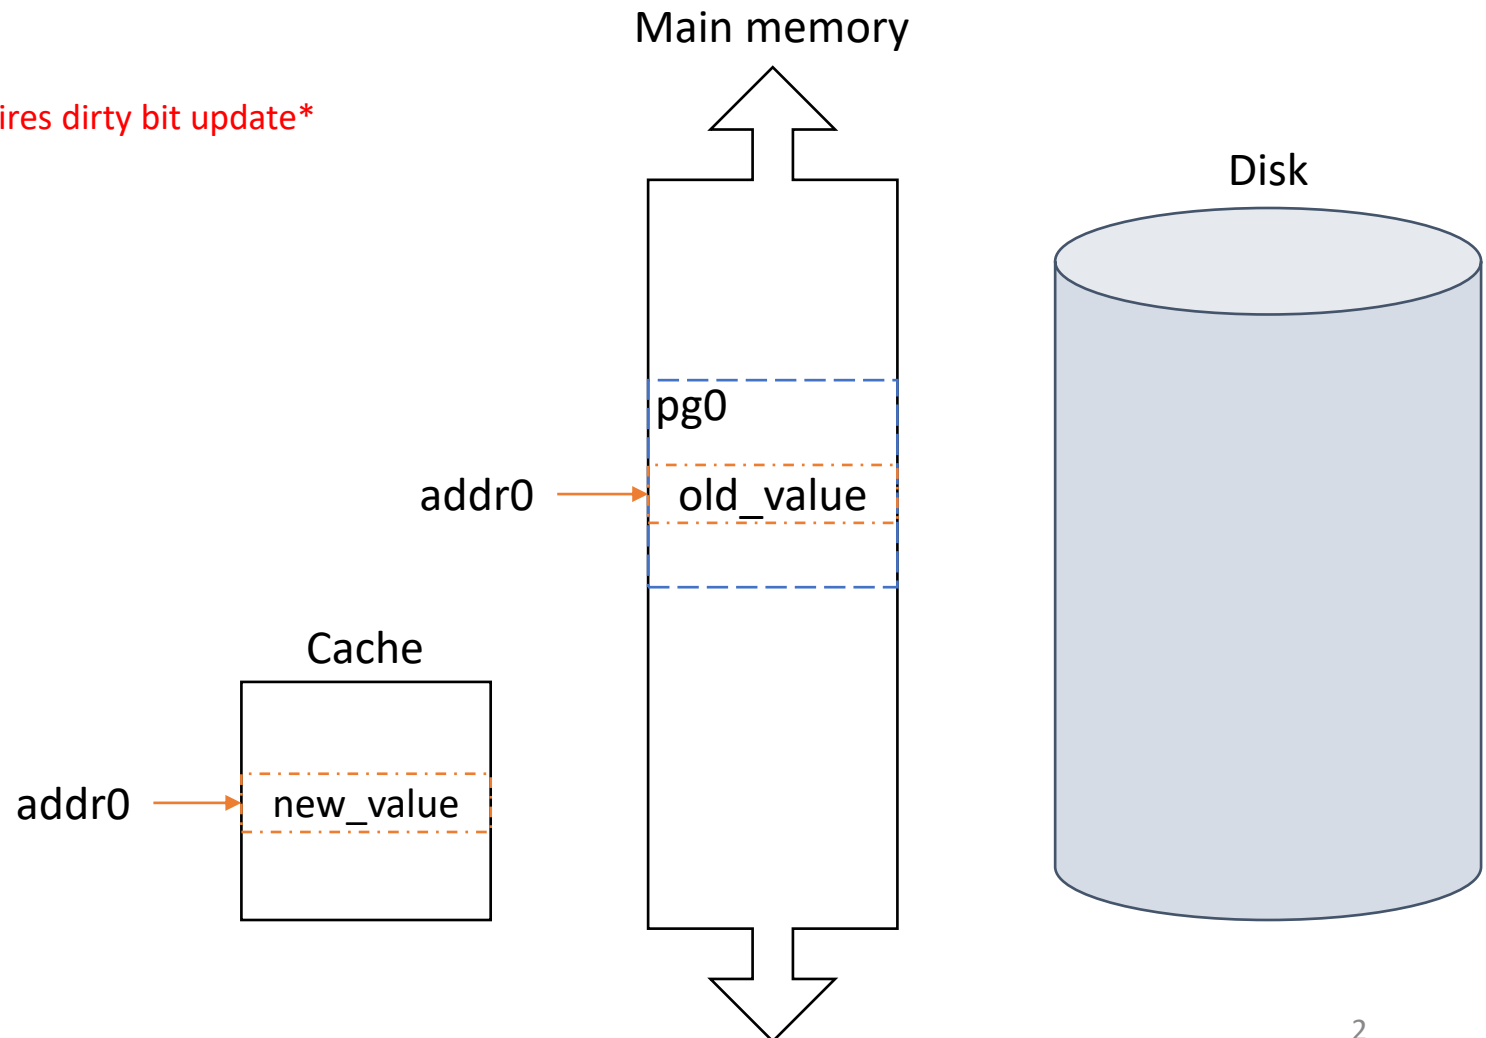
Princeton University

*January 22, 2021*

# Page replacement needed when memory is full and data from disk is being accessed

```
*addr0 = new_value  *requires dirty bit update*  
*  
*  
*  
data = *addr1
```

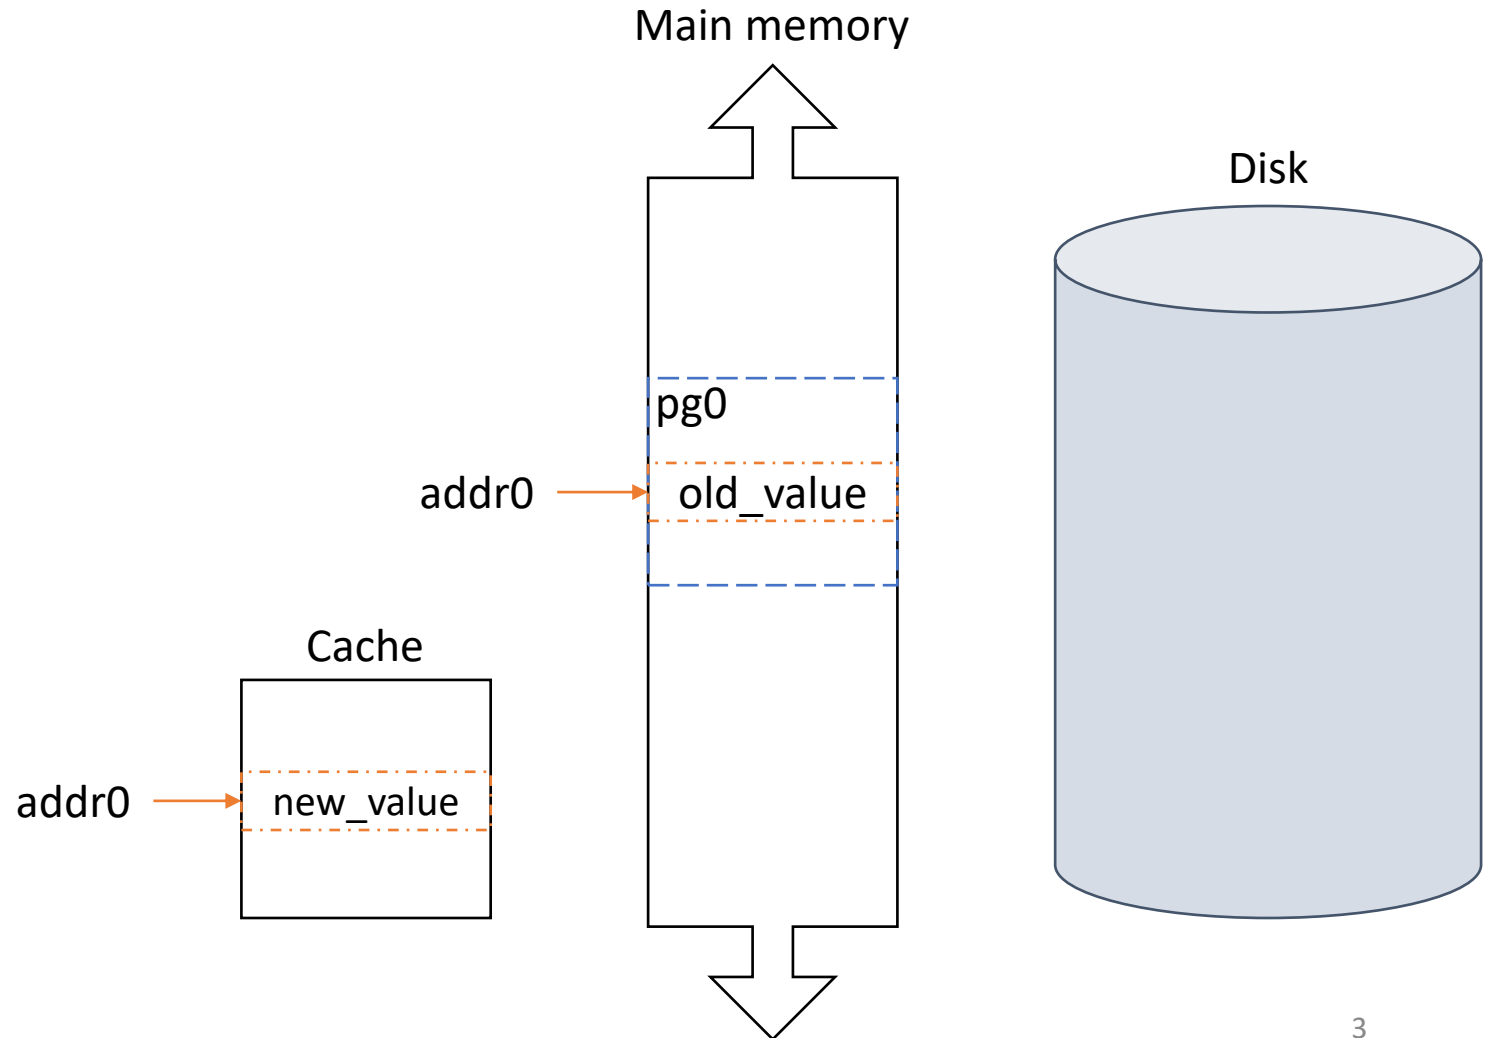
Page Table				
A	D	R	W	Physical page
addr0	0	<b>1</b>	1	pg0
addr1	0	0	1	disk



# Page replacement needed when memory is full and data from disk is being accessed

```
*addr0 = new_value
*
*
*
data = *addr1
```

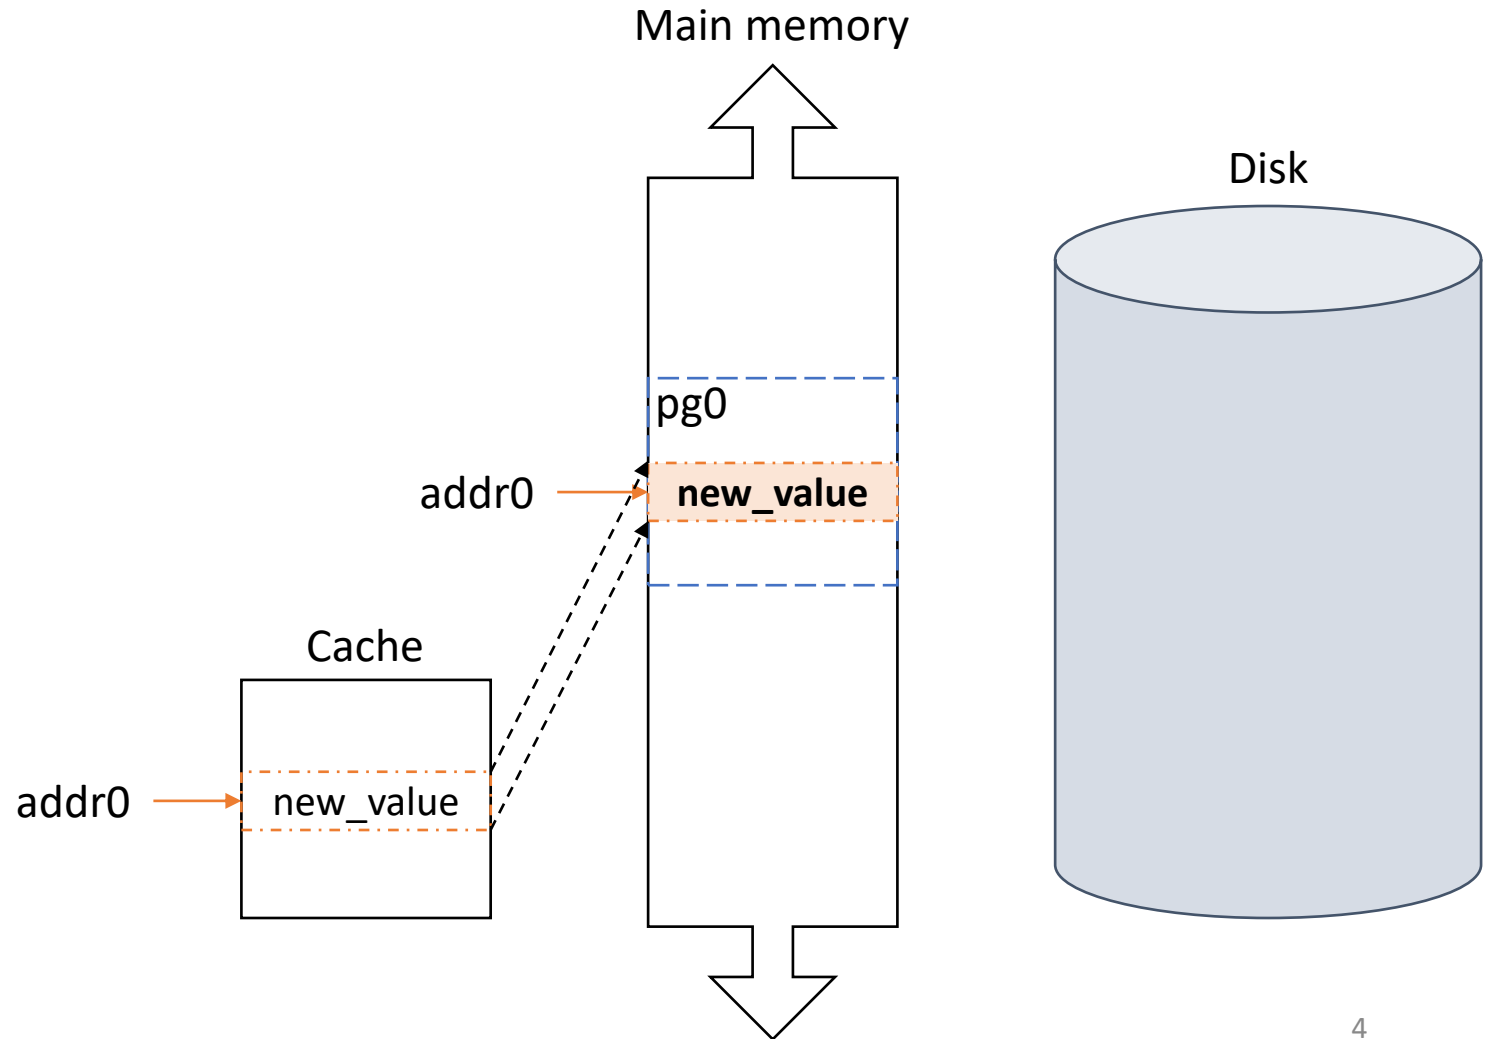
Page Table				
A	D	R	W	Physical page
addr0	0	1	1	pg0
addr1	0	0	1	disk



# Page replacement needed when memory is full and data from disk is being accessed

```
*addr0 = new_value
*
*
*
data = *addr1
```


Page Table				
A	D	R	W	Physical page
addr0	0	<b>0</b>	1	pg0
addr1	0	0	1	disk



# Page replacement needed when memory is full and data from disk is being accessed

```
*addr0 = new_value
*
*
*
data = *addr1
```

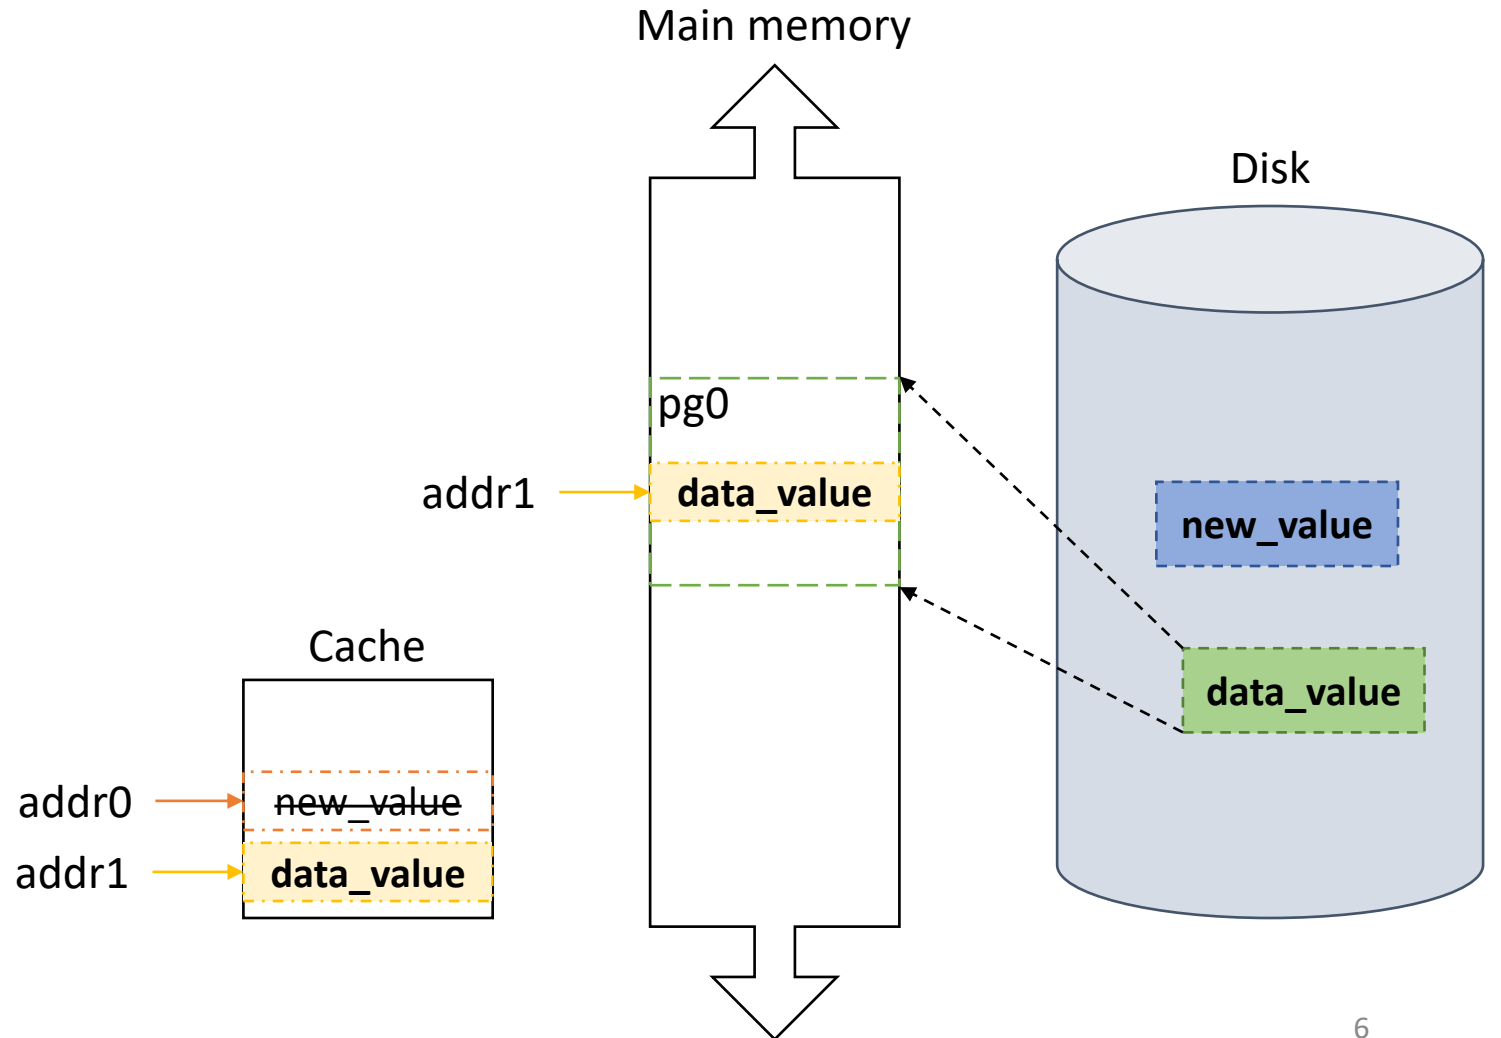
Page Table					
A	D	R	W	Physical page	
addr0	0	0	1	1	disk
addr1	0	0	1	0	disk



# Page replacement needed when memory is full and data from disk is being accessed

```
*addr0 = new_value  
*  
*  
*  
data = *addr1
```

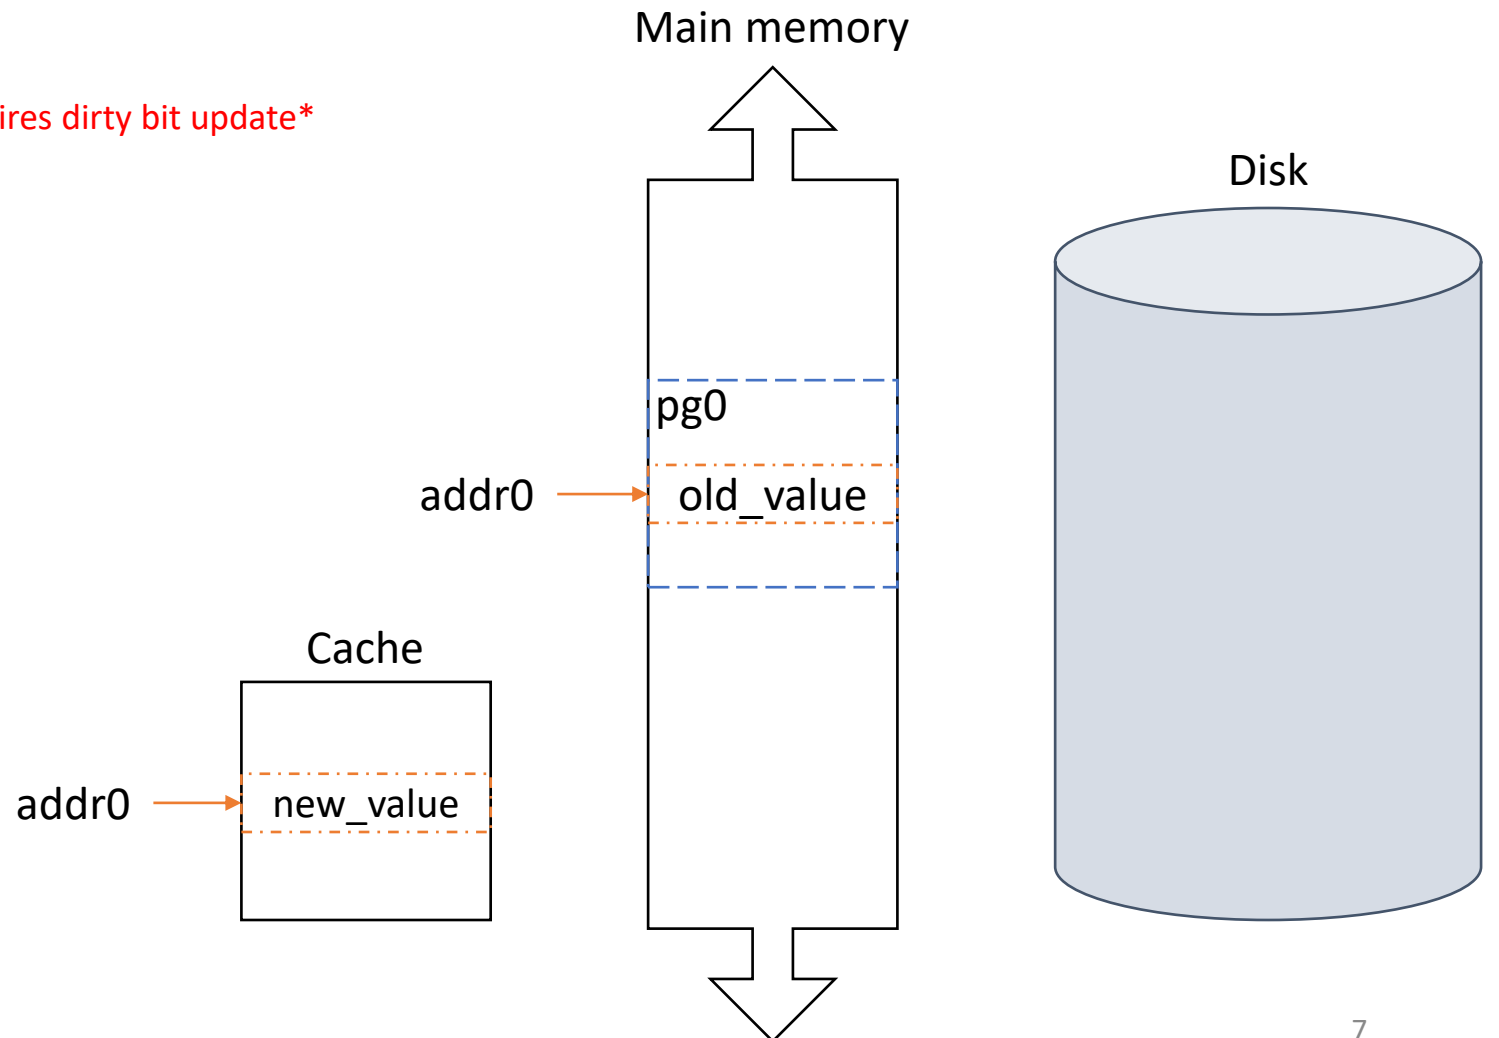
Page Table					
A	D	R	W	Physical page	
addr0	0	0	1	1	disk
addr1	0	0	1	0	pg0



# What if dirty bit is not updated before page swapped to disk?

```
*addr0 = new_value *requires dirty bit update*  
*  
*  
*  
data = *addr1
```

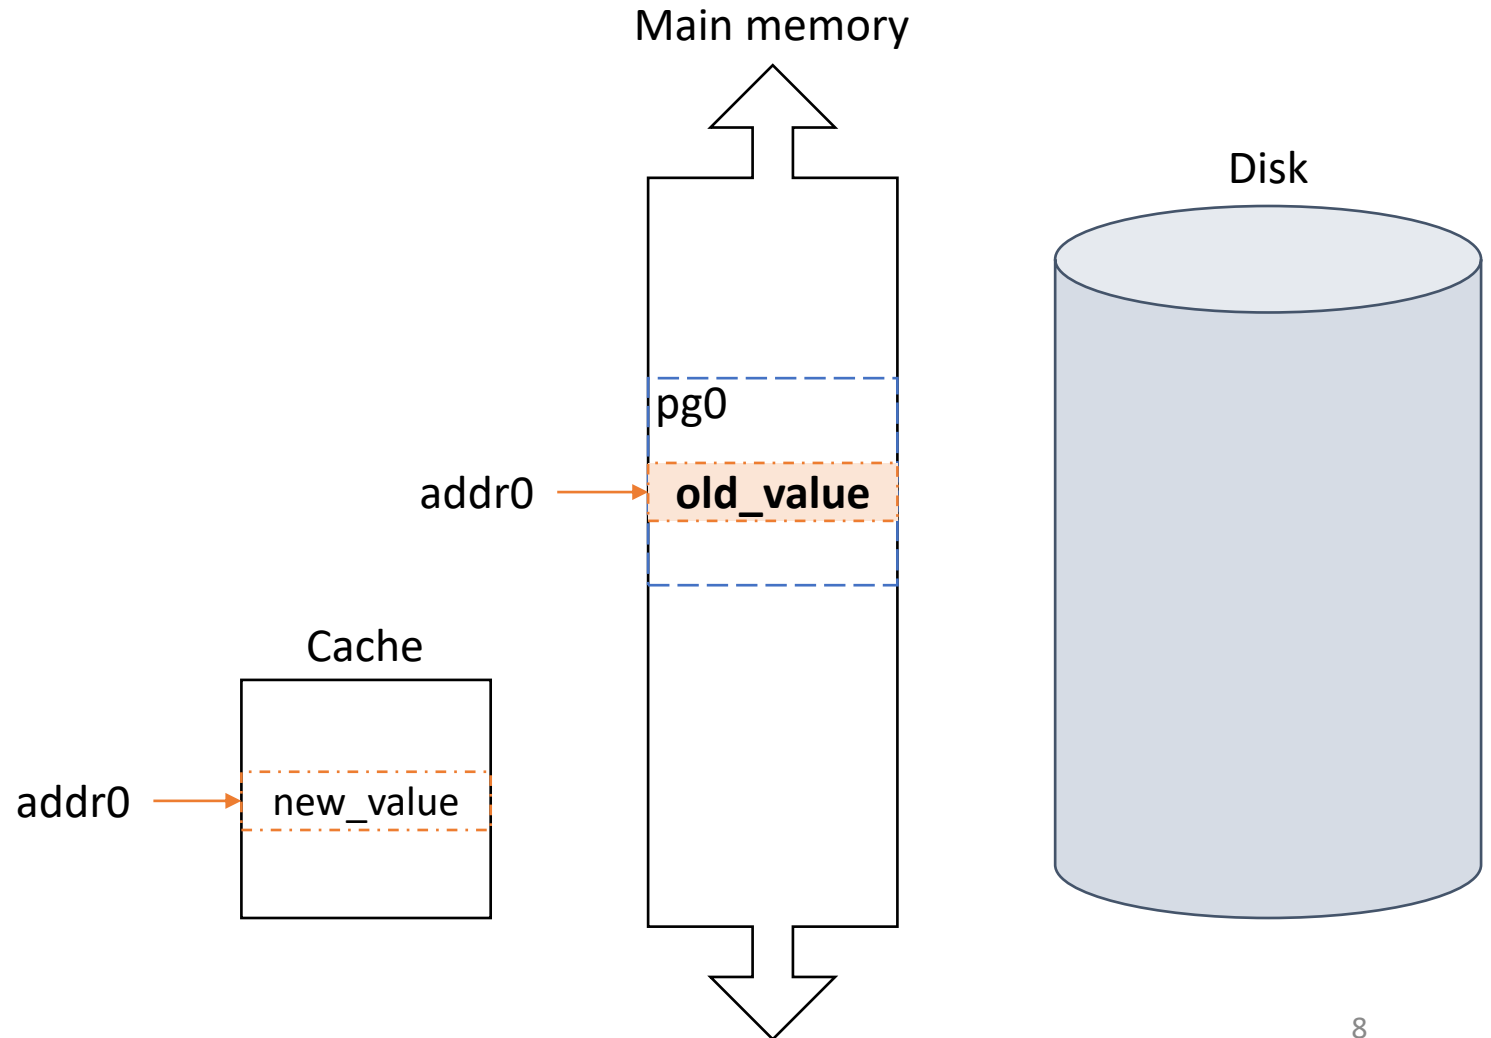
Page Table					
	A	D	R	W	Physical page
addr0	0	0	1	1	pg0
addr1	0	0	1	0	disk



# What if dirty bit is not updated before page swapped to disk?

```
*addr0 = new_value
*
*
*
data = *addr1
```

Page Table				
A	D	R	W	Physical page
addr0	0	0	1	pg0
addr1	0	0	1	disk






# What if dirty bit is not updated before page swapped to disk?

```
*addr0 = new_value
*
*
*
data = *addr1
```

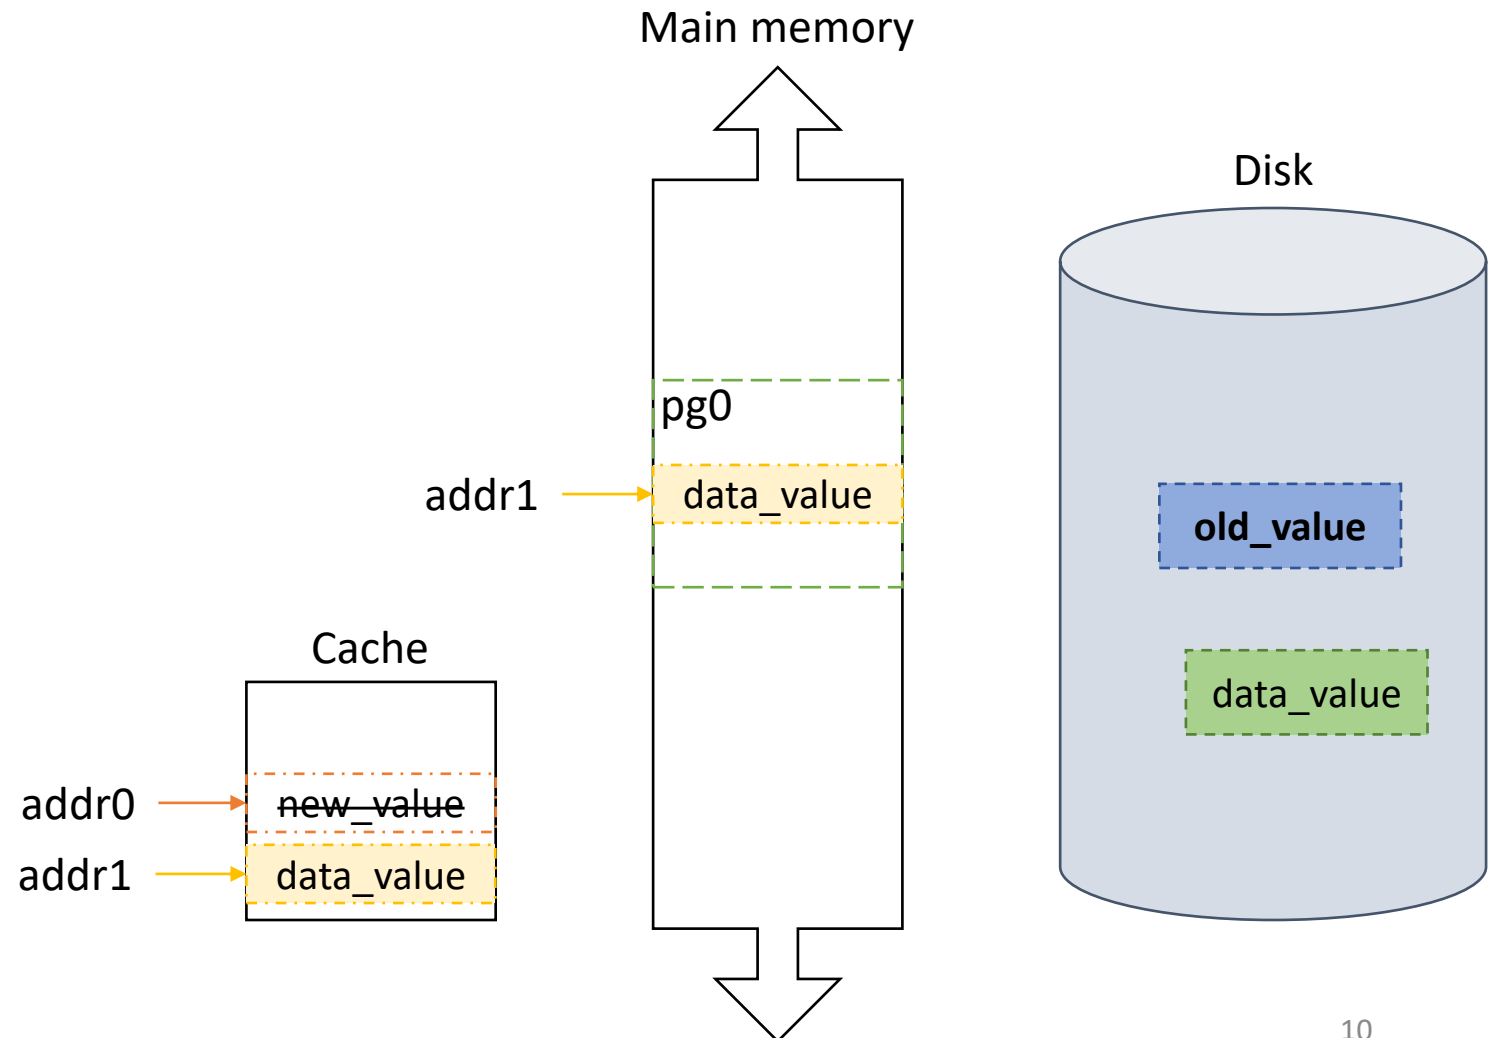
Page Table				
A	D	R	W	Physical page
addr0	0	0	1	disk
addr1	0	0	1	disk



# What if dirty bit is not updated before page swapped to disk?

```
*addr0 = new_value
*
*
*
data = *addr1
*
*
*
data2 = *addr0
```

Page Table					
	A	D	R	W	Physical page
addr0	0	0	1	1	disk
addr1	0	0	1	0	pg0



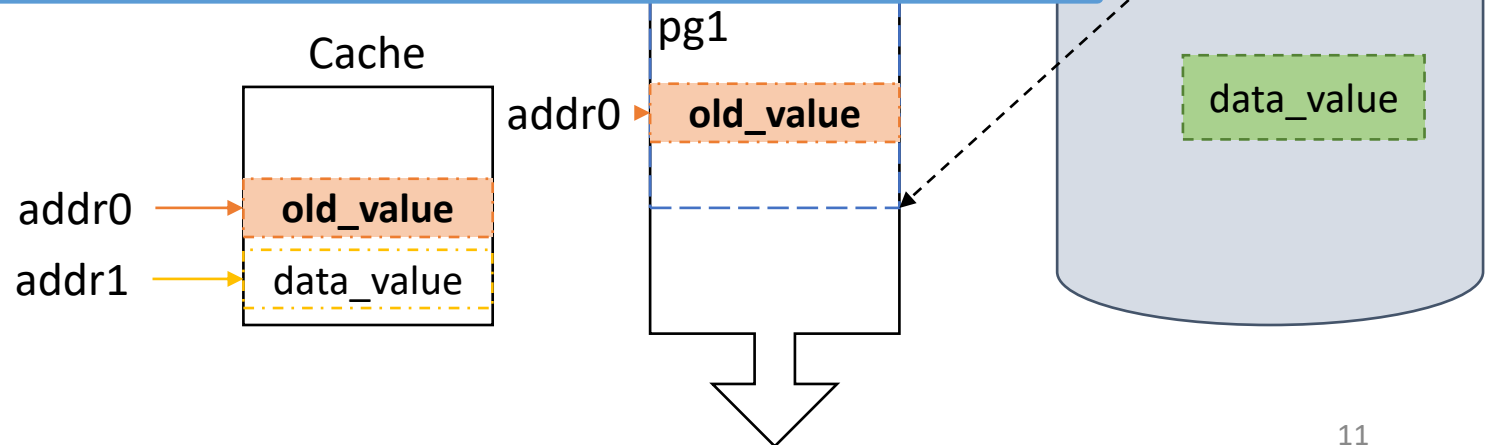
# What if dirty bit is not updated before page swapped to disk?

```
*addr0 = new_value
```

data  
data

Like user-facing memory operations, the dirty bit updates are another **memory reference** that requires **correct ordering** for correct program executions.


Page Table					
	A	D	R	W	Physical page
addr0	0	0	1	1	pg1
addr1	0	0	1	0	pg0



# Memory Consistency Models (MCMs) specify observable behaviors for concurrent programs

- MCMs specify rules for legal values that can be returned when software loads from memory on a shared memory system

**Problem:** don't reason about program executions impacted by shared virtual memory (VM) state with existing MCM constructs



Litmus test: small diagnostic

**This work: Memory Transistency Models (MTMs) – the superset of MCMs that additionally capture VM-aware ordering specifications**

memory fences where needed)

# Prior work

Leslie Lamport defined first MCM: **sequential consistency** [Lamport 1979]

## Formal MCM specification examples:

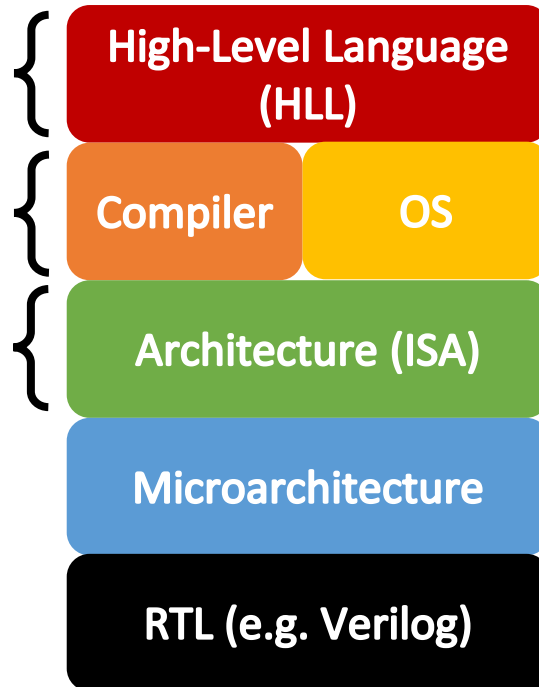
Java [Manson et al., POPL '05]  
C11 [Boehm & Adve, PLDI '08]

Linux Kernel [Alglave et al., ASPLOS '18]

x86-TSO [Owens et al., TPHOLs '09]  
POWER [Mador-Haim et al., CAV '12]

## ISA-level MCM tools aiding verification:

diy tool [Alglave et al., CAV '10]  
litmus tool [Alglave et al., TACAS '11]  
herd tool [Alglave et al., TOPLAS '14]  
Automated MCM litmus test synthesis  
[Lustig et al., ASPLOS 2017]



## Check tool suite:

TriCheck [Trippel et al., ASPLOS '17]

COATCheck [Lustig et al., ASPLOS '16]

PipeCheck [Lustig et al., MICRO '14]

CCICheck [Manerkar et al., MICRO '15]


PipeProof [Manerkar et al., MICRO '18]

RTLCheck [Manerkar et al., MICRO '17]

# Prior work

## Specification and verification of MTMs

Romanescu et al. introduce the concept of separate MCMs for VAs and PAs [ASPLOS '10]


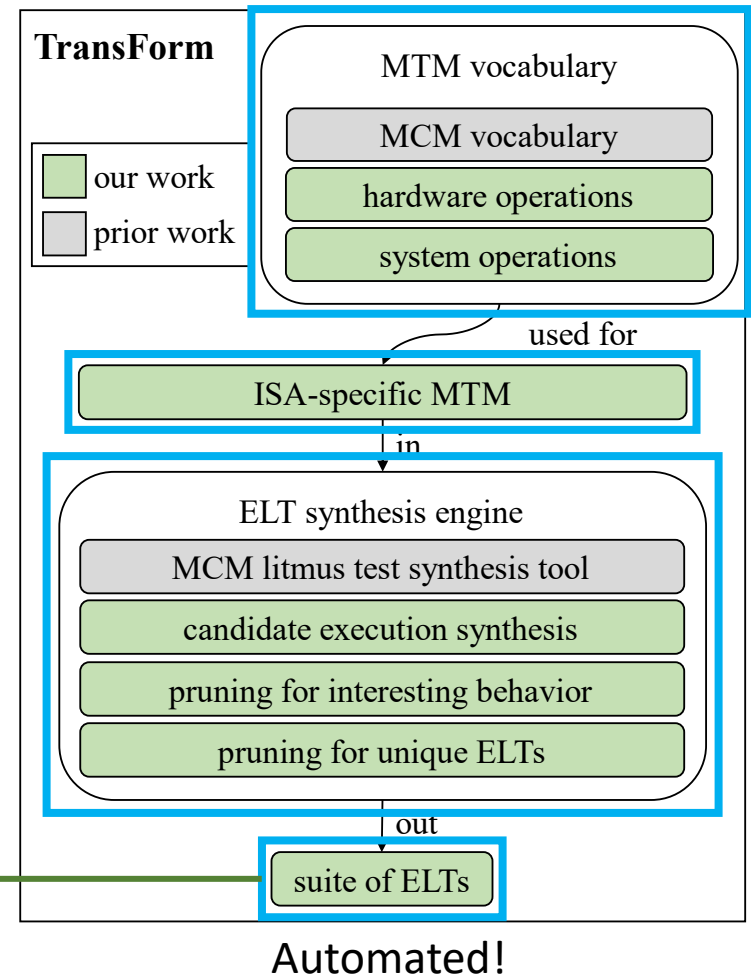


**Check tool suite:**

} COATCheck [Lustig et al., ASPLOS '16]

# TransForm introduces constructs for ISA-level MTM specification and ELT synthesis

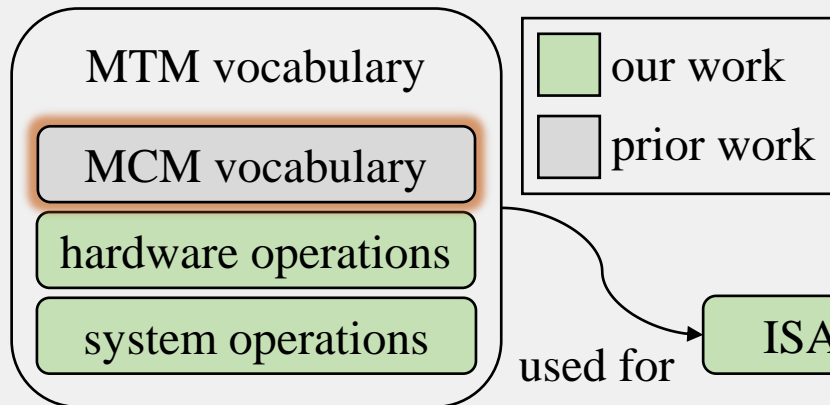
- Formal MTM vocabulary captures system- and hardware-level VM events and interactions with user-facing program instructions
- Enables ISA-level MTM specification
- Enables automated *enhanced litmus test (ELT)* synthesis




Allows for verification against *formally specified* MTM

# Outline

## TransForm




- Background on ISA-level MCM vocabulary
  - Background on virtual memory systems
  - Novel ISA-level MTM vocabulary
  - Automating synthesis of ELTs
  - Case Study: an estimated MTM for x86
  - Future Work & Conclusions
- Prior Work
- My Work





# Approach to defining vocabulary for formally reasoning about MTMs

- MCMs can be defined **axiomatically**
  - Axiomatic MCM specifications use sets of **relations** that can describe user-facing program executions
  - MCM relations describe user-facing event executions of programs with one-to-one V-to-P address mappings
- MTMs are *superset* of MCMs
  - Axiomatic MTM specifications can use MCM relations but require additional relations to similarly describe transistency events and V-to-P address mappings that can have synonyms and be modified



# ISA-level MCM relations can describe programs and their *candidate executions*

**Candidate executions** – set of possible executions of a program and their outcomes

## Program

### Instructions

Core 0	Core 1
a. st data = 1;	c. ld r1 = flag;
b. st flag = 1;	d. ld r2 = data;


$\{W_0, W_1, R_2, R_3\}$

$x, W_1 \rightarrow y, R_2 \rightarrow y, R_3 \rightarrow x$

*program order (po)*

**po** =  $\{W_0 \rightarrow W_1, R_2 \rightarrow R_3\}$

## Graph



Numerical subscripts serve as instruction ID.

## Candidate execution

Core 0	Core 1
$W_0 x = 1$	$R_2 y = 1$
$W_1 y = 1$	$R_3 x = 0$

### Communication (com) relations

*reads from (rf)*

**rf** =  $\{W_1 \rightarrow R_2\}$

*coherence order (co)*

**co** =  $\{\}$

*from reads (fr)*

**fr** =  $\{R_3 \rightarrow W_0\}$

Accessed data (outcome) symbolically represented by com relations

[Shasha & Snir, 1988]  
[Alglave et al., 2014]

# ISA-level MCM relations can describe programs and their *candidate executions*

**Candidate executions** – set of possible executions of a program and their outcomes

## Program

Core 0	Core 1
$W_0 x$	$R_2 y$
$W_1 y$	$R_3 x$

## Instructions

**Event** =  $\{W_0, W_1, R_2, R_3\}$

**MemoryEvent** =  $\{W_0, W_1, R_2, R_3\}$


**Location** =  $\{x, y\}$

**address**  $\{W_0 \rightarrow x, W_1 \rightarrow y, R_2 \rightarrow y, R_3 \rightarrow x\}$

*program order (po)*

**po** =  $\{W_0 \rightarrow W_1, R_2 \rightarrow R_3\}$

## Graph



Numerical subscripts serve as instruction ID.

## Candidate execution

Core 0	Core 1
$W_0 x = 1$	$R_2 y = 1$
$W_1 y = 1$	$R_3 x = 0$

## Communication (com) relations

*reads from (rf)*

**rf** =  $\{W_1 \rightarrow R_2\}$

*coherence order (co)*

**co** =  $\{\}$

*from reads (fr)*

**fr** =  $\{R_3 \rightarrow W_0\}$

Accessed data (outcome) symbolically represented by com relations

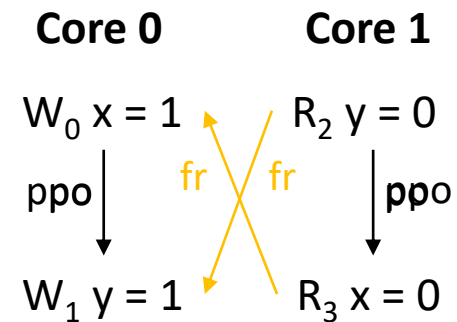
[Shasha & Snir, 1988]  
[Alglave et al., 2014]

# MCM specifications place constraints on permitted execution behaviors

Axiomatic MCM specifications use MCM relations to describe axioms that constrain candidate execution behaviors

Intel x86 processors use the **total store order (TSO)** memory model (**x86-TSO**) [Owens et al., 2009]: strict sequential memory access orderings but relaxed Store->Load orderings to allow for store buffering

**Causality** – axiom for x86-TSO:  
**acyclic(rfe + co + fr + ppo + fence)**




# MCM specifications place constraints on permitted execution behaviors

Axiomatic MCM specifications use MCM relations to describe axioms that constrain candidate execution behaviors

Intel x86 processors use the **total store order (TSO)** memory model (**x86-TSO**) [Owens et al., 2009]: strict sequential memory access orderings but relaxed Store->Load orderings to allow for store buffering

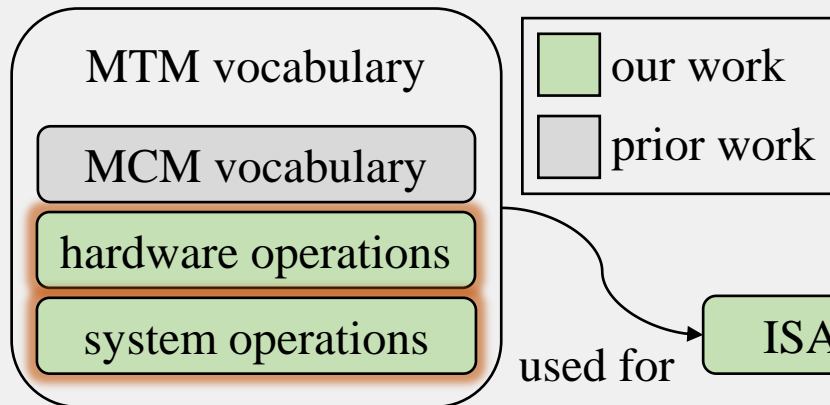
**Causality** – axiom for x86-TSO:  
**acyclic(fr + co + fr + ppo + fence)**

mp (“message passing”) litmus test



# Outline

## TransForm



- Background on ISA-level MCM vocabulary
  - Background on virtual memory systems
  - Novel ISA-level MTM vocabulary
  - Automating synthesis of ELTs
  - Case Study: a estimated MTM for x86
  - Future Work & Conclusions
- Prior Work
- My Work



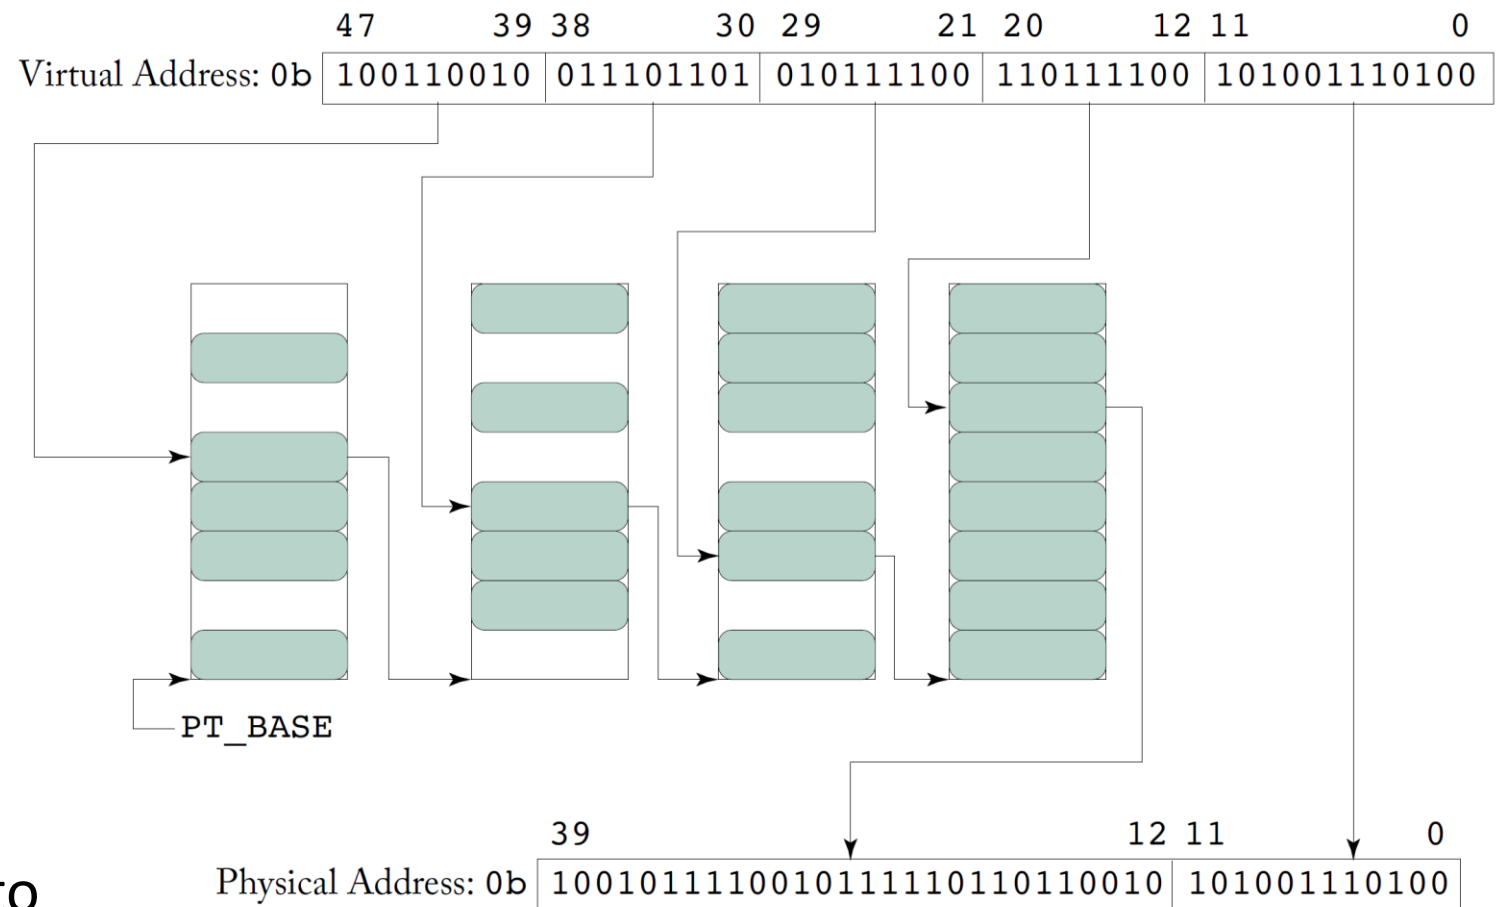
# V-to-P address mappings need to be stored and accessed during memory events

V-to-P address mappings are stored in **page tables**.

**Page table entries (PTEs)** hold address mapping and status bits (permissions, access, dirty).

Page tables are usually structured hierarchically.

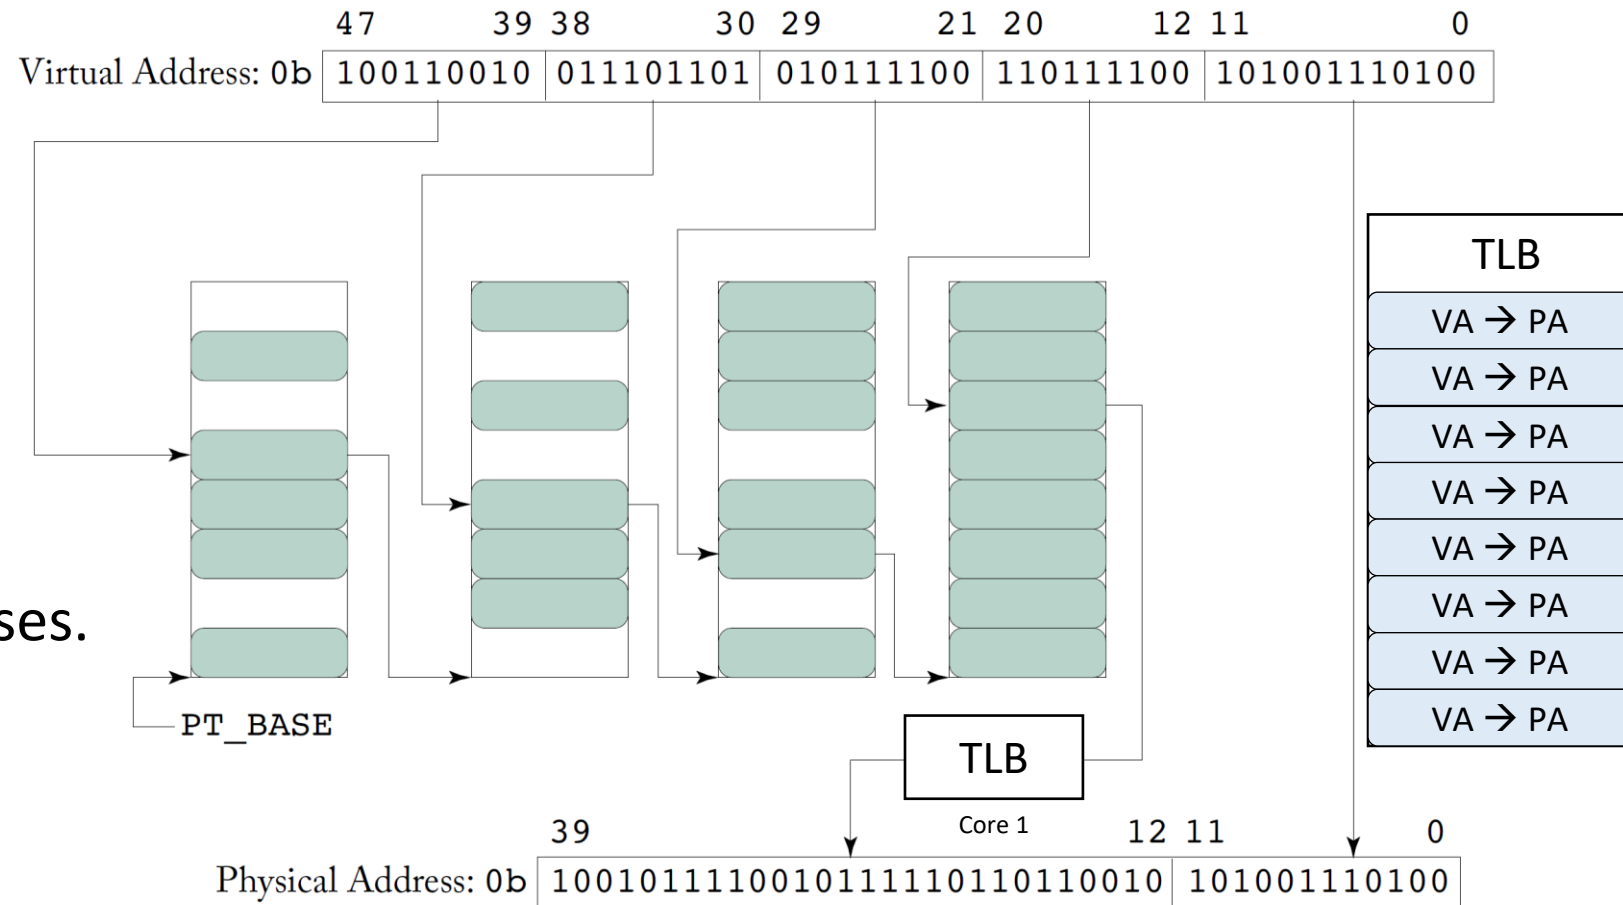
When address translation is needed, a **page table walk** traverses the page table levels to find the desired address mapping.



# V-to-P address mappings are cached in the *translation lookaside buffer (TLB)*

Hierarchical page tables require **additional** memory accesses during address translation – big performance hit.

Page mappings cached in TLB to reduce latency of memory accesses.

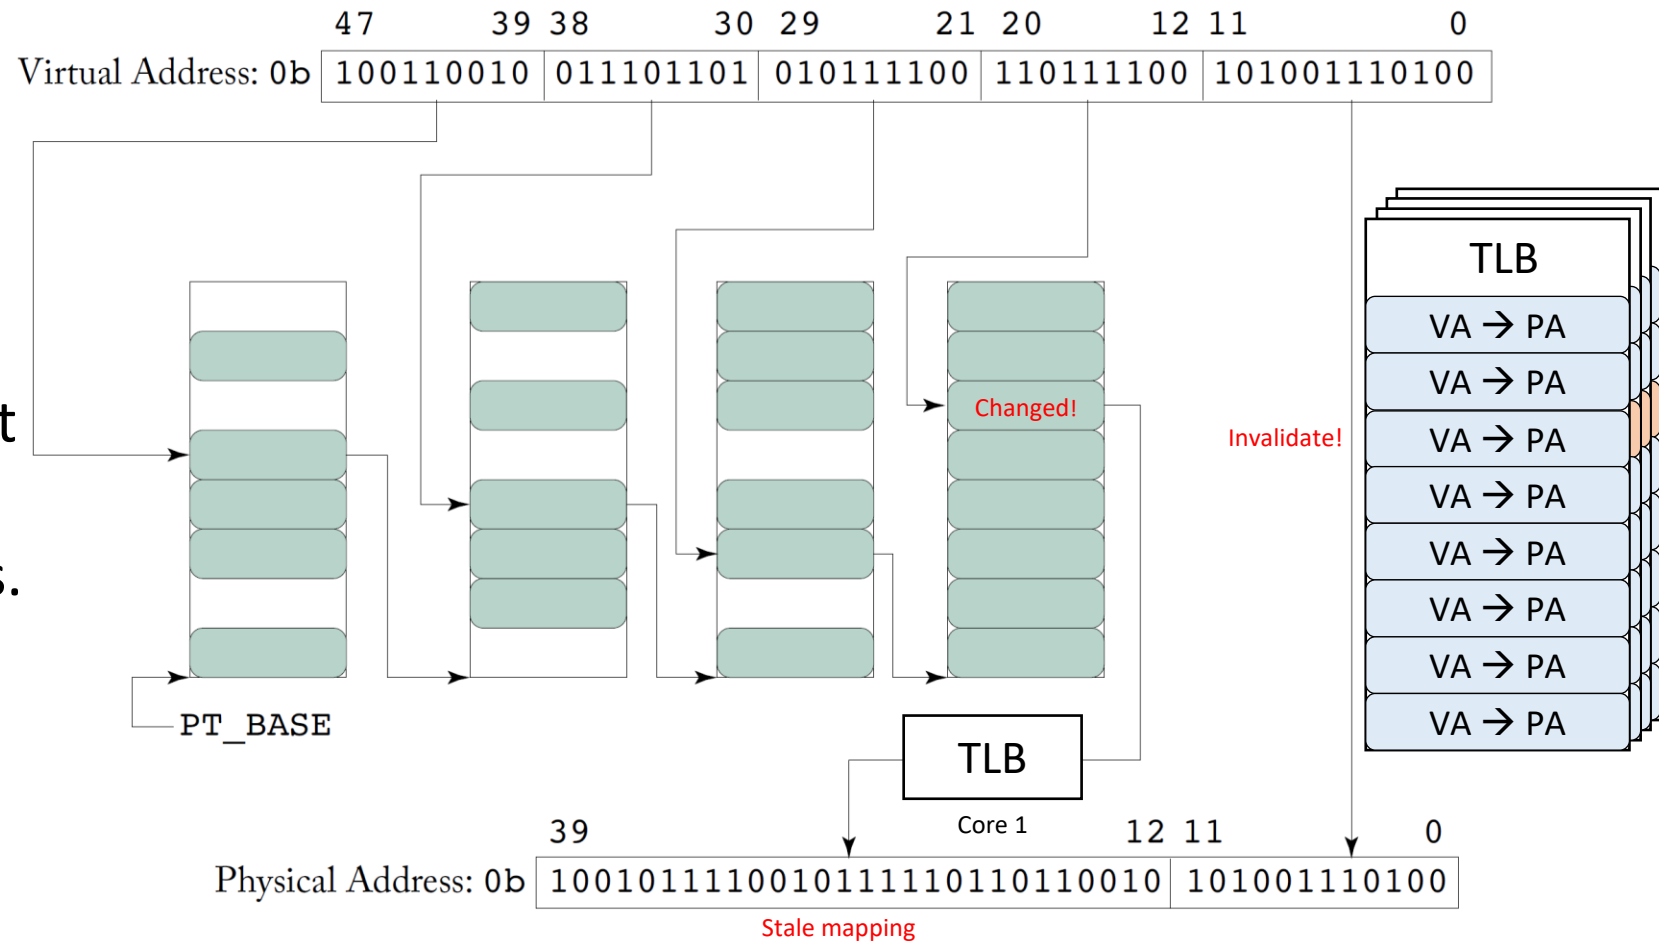




# V-to-P address mappings can be changed by OS

Operating system (OS) may change address mappings in page tables.

Corresponding TLB entries must be invalidated on *each core* to prevent stale mapping accesses.

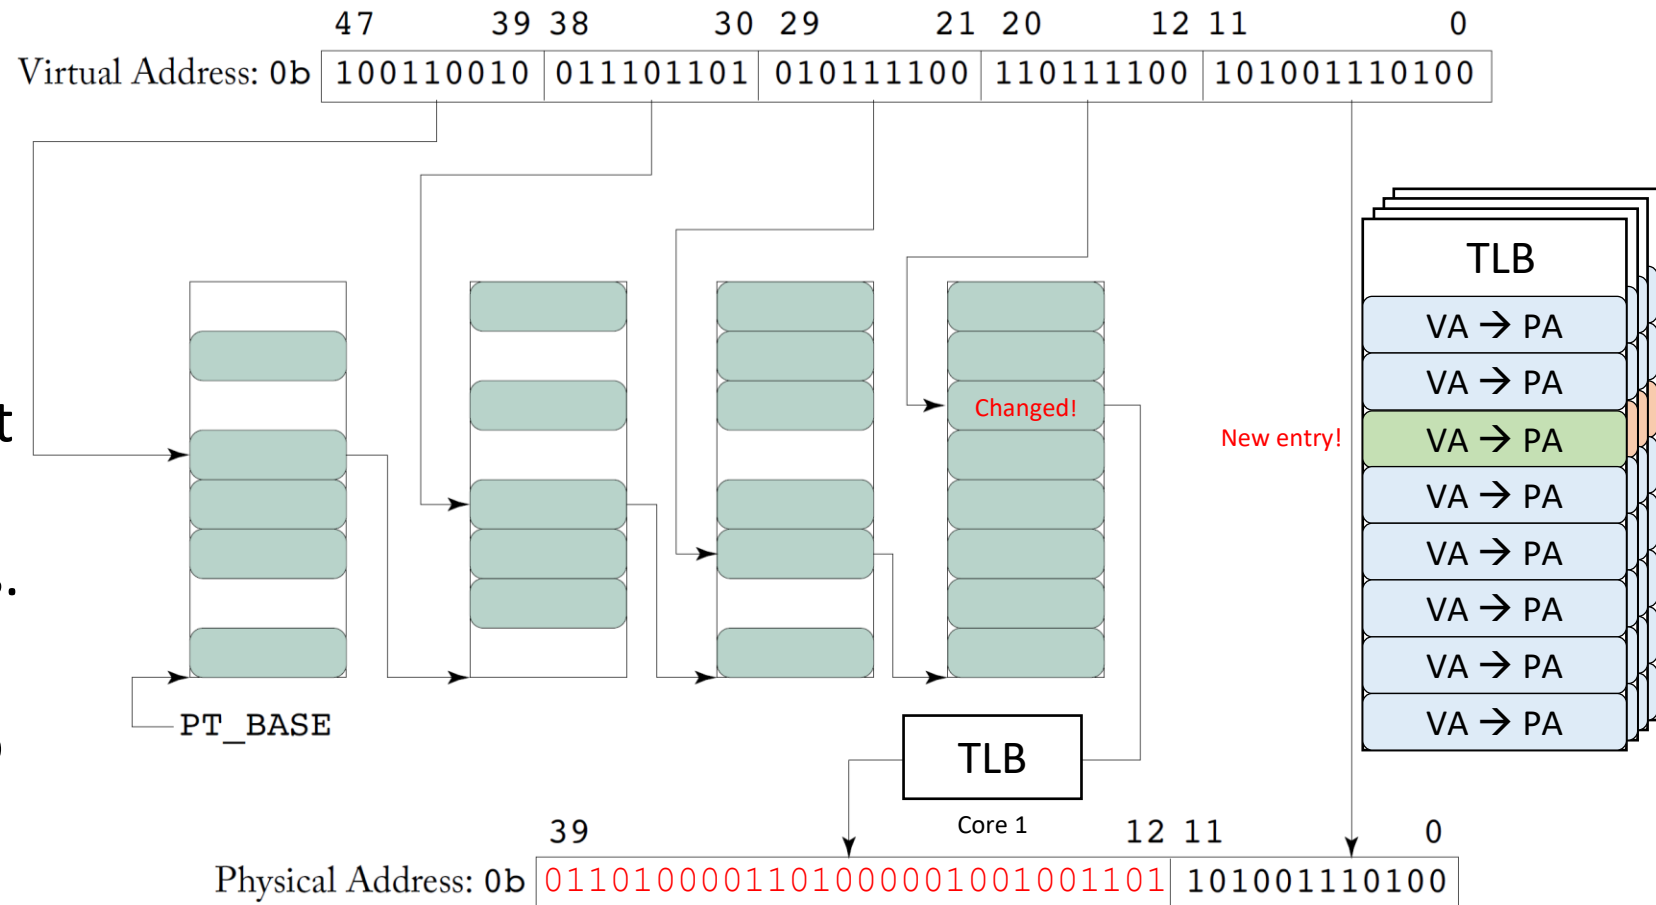


# V-to-P address mappings can be changed by OS

Operating system (OS) may change address mappings in page tables.

Corresponding TLB entries must be invalidated on *each core* to prevent stale mapping accesses.

New page table walk needed to load new mapping into TLB.

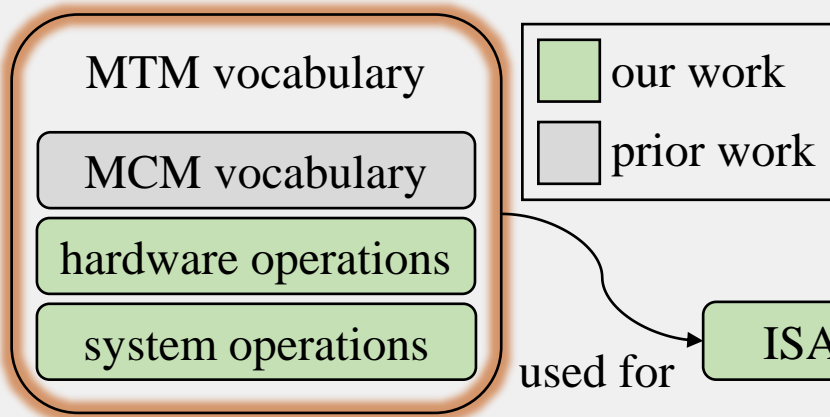


# Virtual memory events TransForm needs to support

Hardware-level events	System-level events
Page table walk Loads TLB entries on memory access	Address mapping changes V-to-P address mapping must be modifiable like data
PTE status bit updates TransForm supports dirty bit updates on memory stores	TLB entry invalidations May be invoked on multiple cores by address mapping changes

# Outline

## TransForm



- Background on ISA-level MCM vocabulary
  - Background on virtual memory systems
  - Novel ISA-level MTM vocabulary
  - Automating synthesis of ELTs
  - Case Study: a estimated MTM for x86
  - Future Work & Conclusions
- Prior Work
- My Work




# MTM Vocabulary: Hardware-level events

TransForm supports **page table walks (PTW)** and **dirty bit updates**

## Ghost instructions

**PTW:** loads translation  
lookaside buffer (TLB) entry

**dirty bit update:** modifies dirty  
bit in PTE



# MTM Vocabulary: Hardware-level events

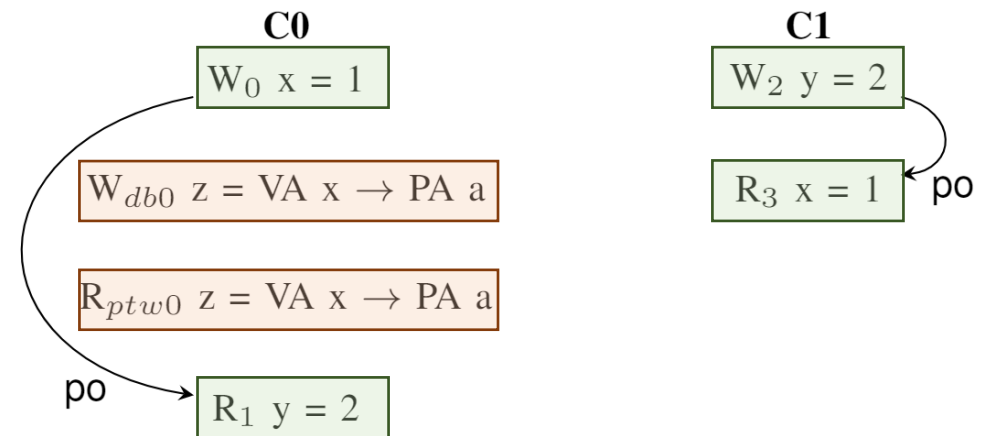
TransForm supports **page table walks (PTW)** and **dirty bit updates**

## Ghost instructions

**PTW:** loads translation  
lookaside buffer (TLB) entry

**dirty bit update:** modifies dirty  
bit in PTE

**ghost** – relates user-facing MemoryEvent to  
invoked ghost instructions (numerical subscripts)



# MTM Vocabulary: Hardware-level events

TransForm supports **page table walks (PTW)** and **dirty bit updates**

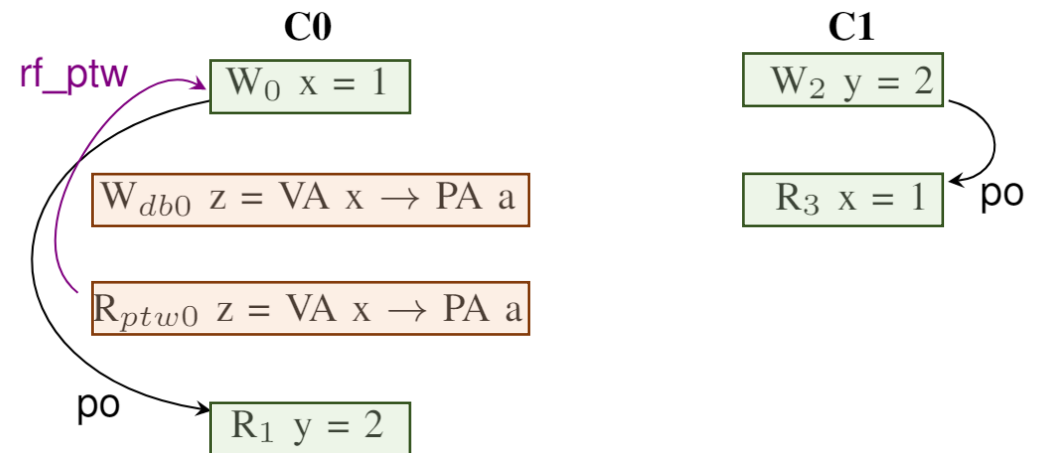
## Ghost instructions

**PTW**: loads translation  
lookaside buffer (TLB) entry

**dirty bit update**: modifies dirty  
bit in PTE

**ghost** – relates user-facing MemoryEvent to  
invoked ghost instructions (numerical subscripts)

**rf\_ptw** – relates PTW to user-facing  
MemoryEvents that access loaded TLB entry



# MTM Vocabulary: Hardware-level events

TransForm supports **page table walks (PTW)** and **dirty bit updates**

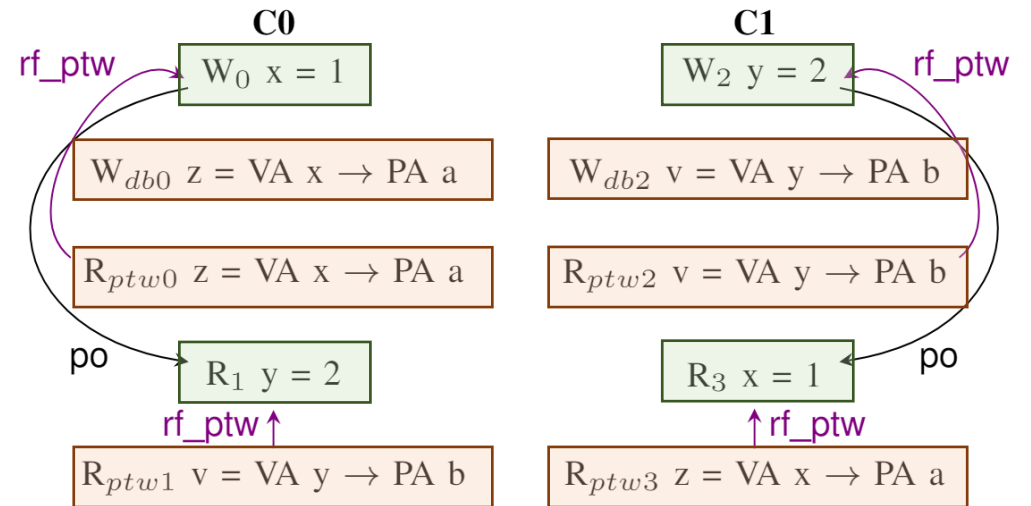
## Ghost instructions

**PTW**: loads translation  
lookaside buffer (TLB) entry

**dirty bit update**: modifies dirty  
bit in PTE

**ghost** – relates user-facing MemoryEvent to  
invoked ghost instructions (numerical subscripts)

**rf\_ptw** – relates PTW to user-facing  
MemoryEvents that access loaded TLB entry





# MTM Vocabulary: Hardware-level events

TransForm supports **page table walks (PTW)** and **dirty bit updates**


## Ghost instructions

**PTW**: loads translation  
lookaside buffer (TLB) entry

**dirty bit update**: modifies dirty  
bit in PTE

**ghost** – relates user-facing MemoryEvent to  
invoked ghost instructions (numerical subscripts)

**rf\_ptw** – relates PTW to user-facing  
MemoryEvents that access loaded TLB entry

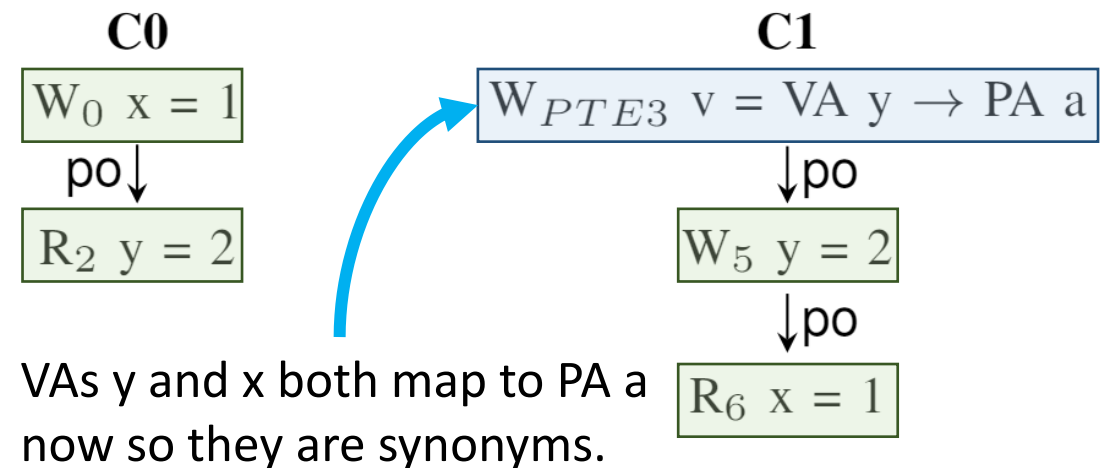


# MTM Vocabulary: System-level events

TransForm supports **address remappings via PTE Writes** and **TLB entry invalidations**

## Support instructions

**PTE Write:** changes address mapping stored in a PTE for some VA  $v$



# MTM Vocabulary: System-level events

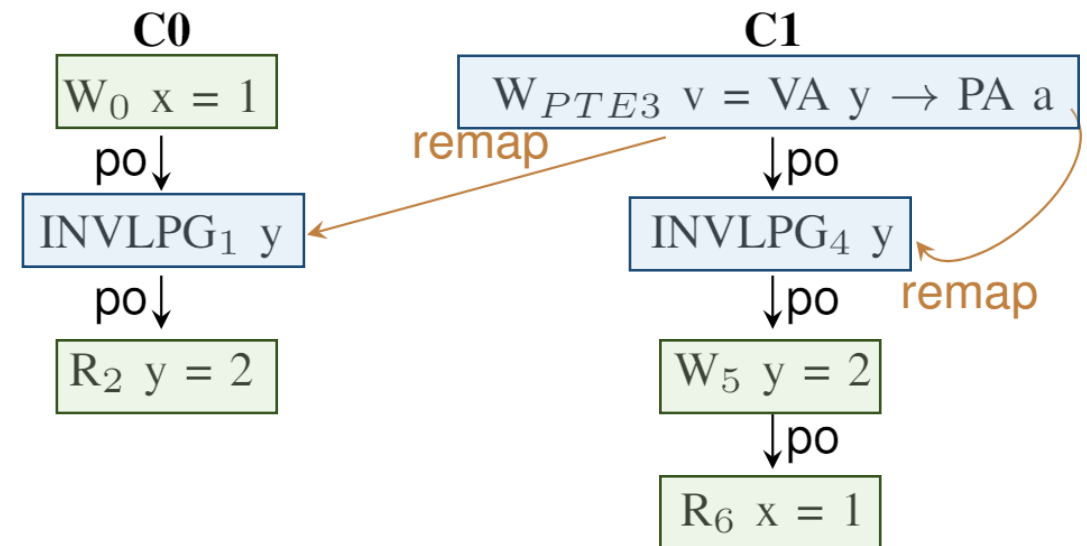
TransForm supports **address remappings via PTE Writes** and **TLB entry invalidations**

## Support instructions

**PTE Write:** changes address mapping stored in a PTE for some VA  $v$

**INVLPG:** invalidates TLB entry (named after x86 instruction)

**remap** – relates PTE Writes to invoked INVLPGs



# MTM Vocabulary: System-level events

TransForm supports **address remappings via PTE Writes** and **TLB entry invalidations**

## Support instructions

**PTE Write**: changes address mapping stored in a PTE for some VA  $v$

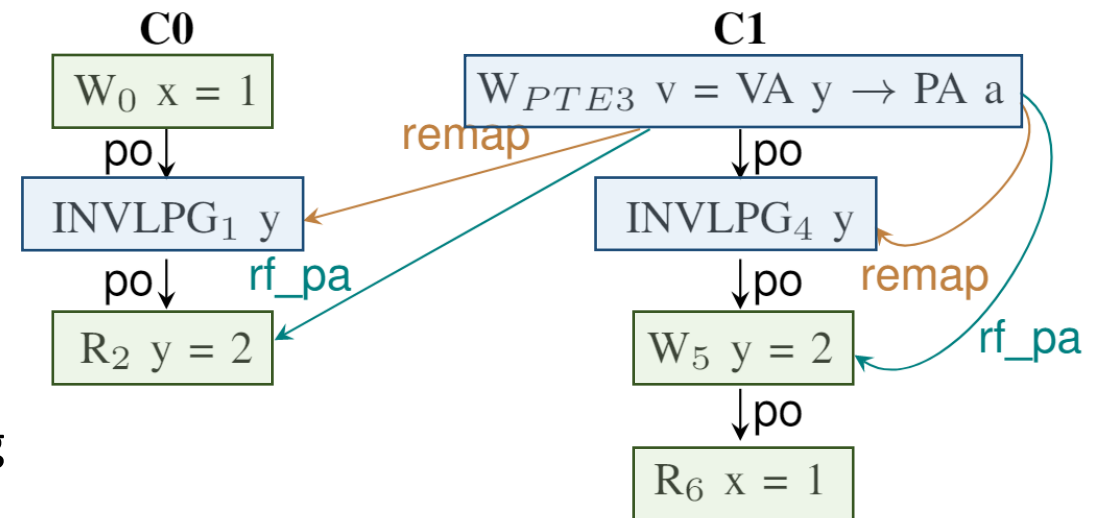
**INVLPG**: invalidates TLB entry (named after x86 instruction)

**remap** – relates PTE Writes to invoked INVLPGs

**rf\_pa** – relates PTE Write for VA  $v \rightarrow$  PA  $p$  to user-facing MemoryEvents accessing PA  $p$  via VA  $v$

**fr\_pa** – relates user-facing MemoryEvents accessing PA  $p$  via VA  $v$  to PTE Writes for VA  $v' \rightarrow$  PA  $p$

**co\_pa** and **fr\_va** follow similarly



# MTM Vocabulary: System-level events

TransForm supports **address remappings via PTE Writes** and **TLB entry invalidations**

## Support instructions

**PTE Write**: changes address mapping stored in a PTE for some VA  $v$

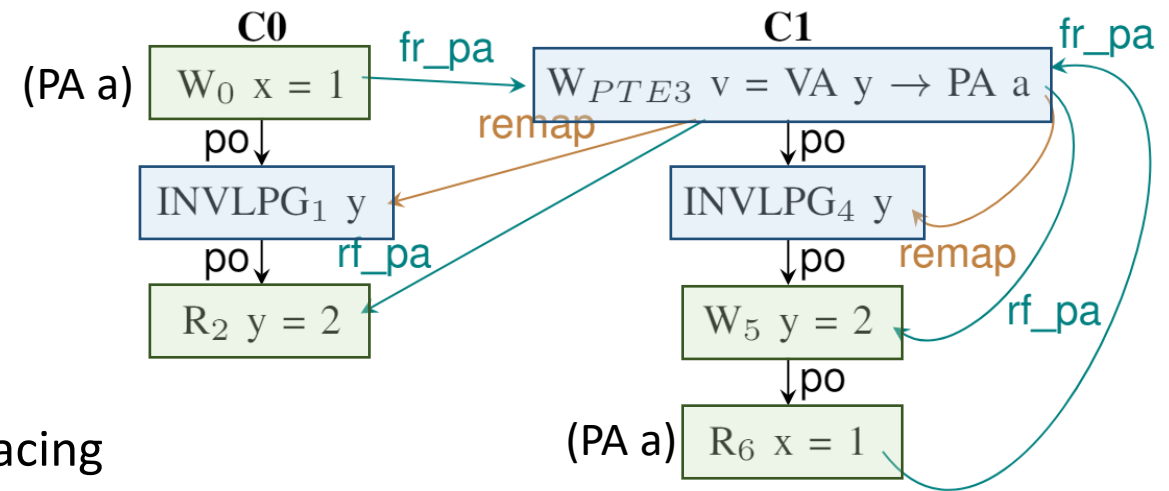
**INVLPG**: invalidates TLB entry (named after x86 instruction)

**remap** – relates PTE Writes to invoked INVLPGs

**rf\_pa** – relates PTE Write for VA  $v \rightarrow$  PA  $p$  to user-facing MemoryEvents accessing PA  $p$  via VA  $v$

**fr\_pa** – relates user-facing MemoryEvents accessing PA  $p$  via VA  $v$  to PTE Writes for VA  $v' \rightarrow$  PA  $p$

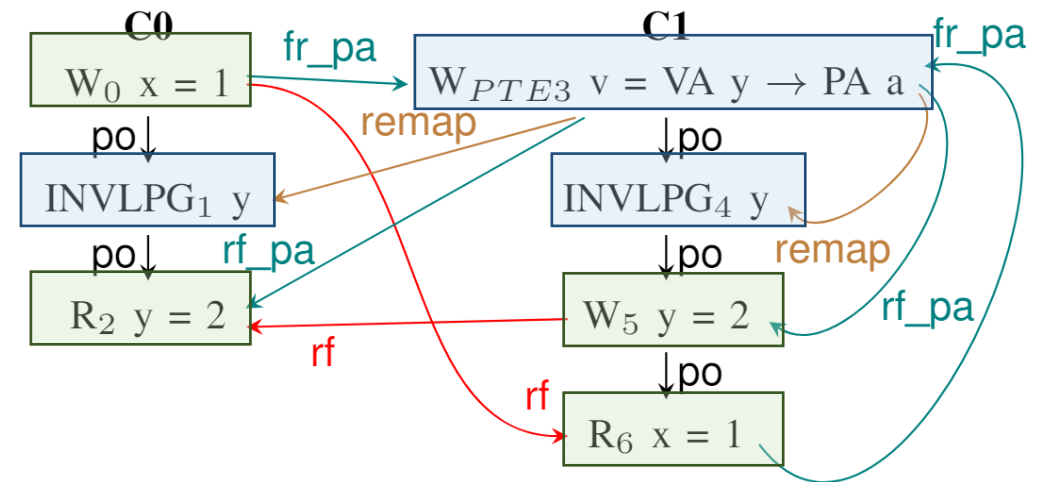
**co\_pa** and **fr\_va** follow similarly



# MTM Vocabulary: System-level events

TransForm supports **address remappings via PTE Writes** and **TLB entry invalidations**

These new com relations can be used to derive same PA accesses.



**rf\_pa** – relates PTE Write for VA  $v \rightarrow PA \ p$  to user-facing MemoryEvents accessing PA  $p$  via VA  $v$

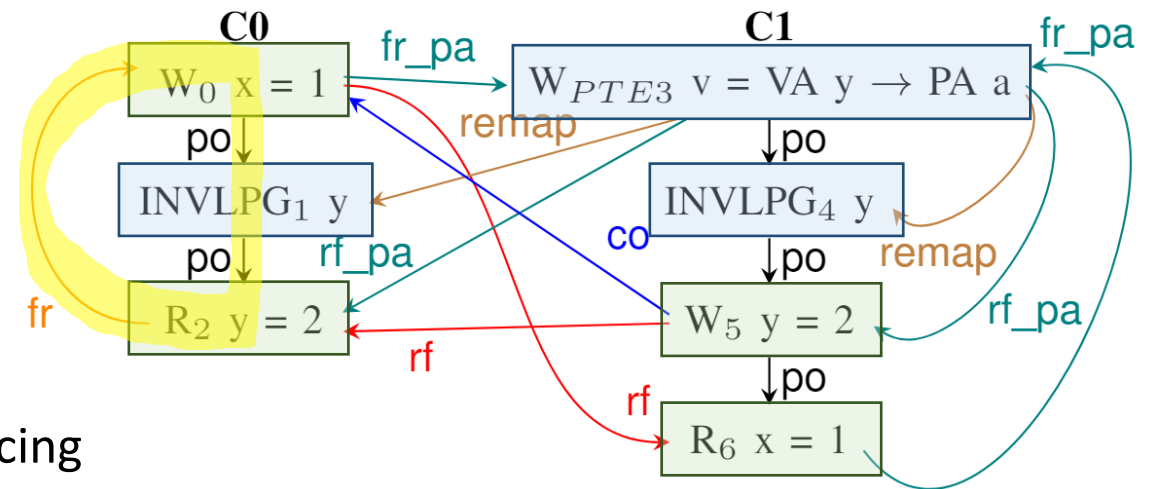
**fr\_pa** – relates user-facing MemoryEvents accessing PA  $p$  via VA  $v$  to PTE Writes for VA  $v' \rightarrow PA \ p$

**co\_pa** and **fr\_va** follow similarly

# MTM Vocabulary: System-level events

TransForm supports **address remappings via PTE Writes** and **TLB entry invalidations**

These new com relations can be used to derive same PA accesses.




**rf\_pa** – relates PTE Write for VA  $v \rightarrow PA p$  to user-facing MemoryEvents accessing PA  $p$  via VA  $v$

**fr\_pa** – relates user-facing MemoryEvents accessing PA  $p$  via VA  $v$  to PTE Writes for VA  $v' \rightarrow PA p$

**co\_pa** and **fr\_va** follow similarly

# MTM Vocabulary: Putting it all together

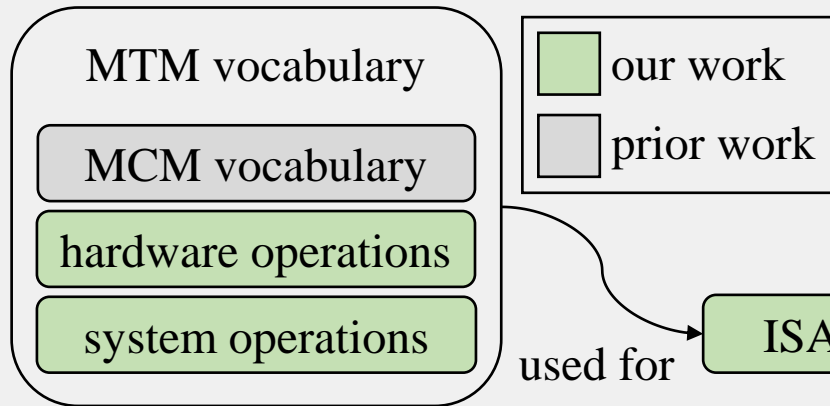
Program executions with transistency events and relations can get quite complex but they allow us to capture these additional interactions that can occur and impact the program's execution.






# Outline

## TransForm



- Background on ISA-level MCM vocabulary
  - Background on virtual memory systems
  - Novel ISA-level MTM vocabulary
  - Automating synthesis of ELTs
  - Case Study: an estimated MTM for x86
  - Future Work & Conclusions
- Prior Work
- My Work

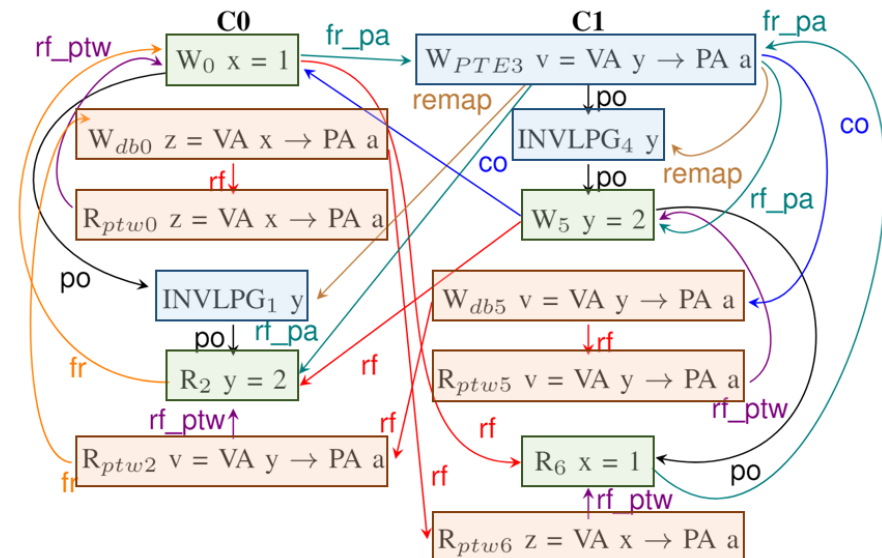
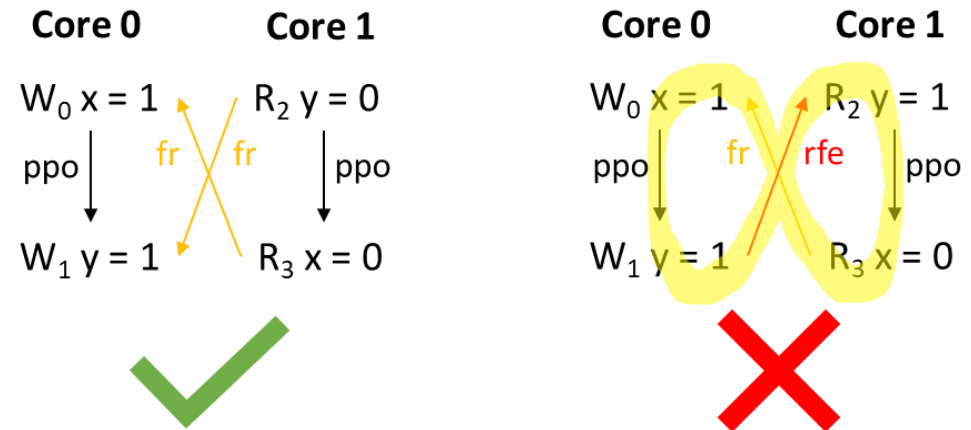


MCMs can be verified with litmus tests.

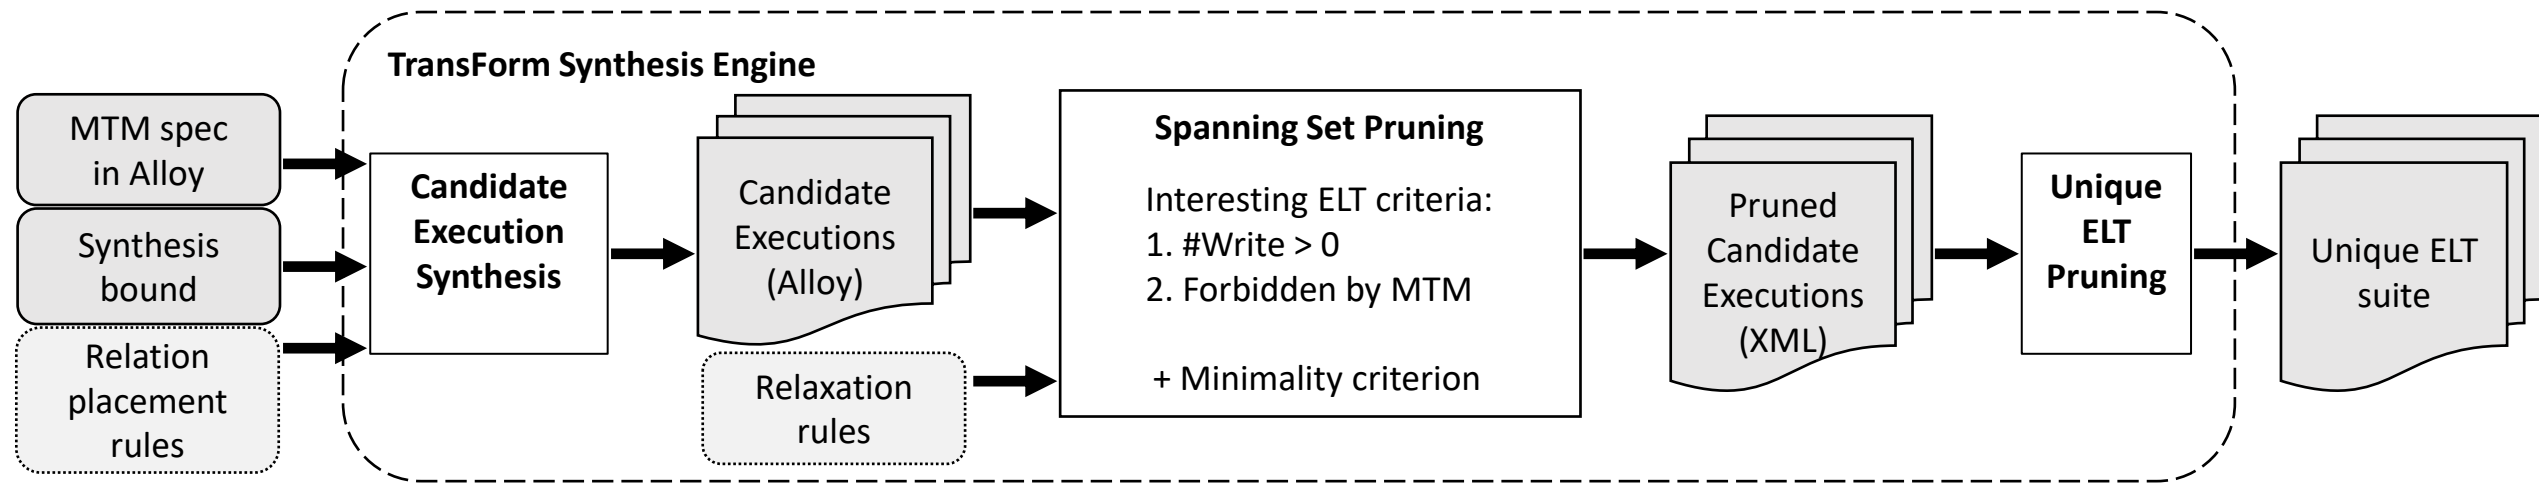
MTMs can be verified with *enhanced* litmus tests.

- **Litmus tests:** small diagnostic programs for validating MCM behaviors
  - Executions and their outcomes can be deemed **permitted** or **forbidden** by MCM specification

- **Enhanced litmus tests (ELTs):** litmus tests enhanced with system- and hardware-level events that facilitate address translation

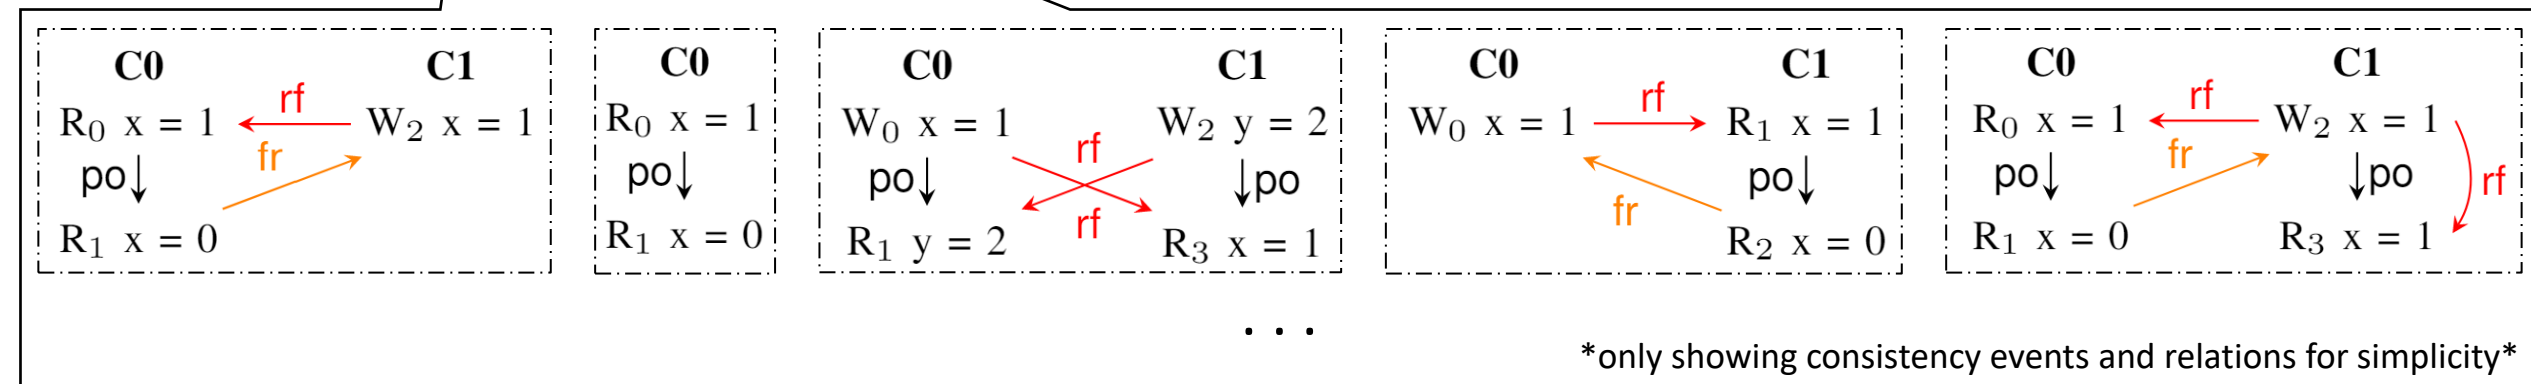
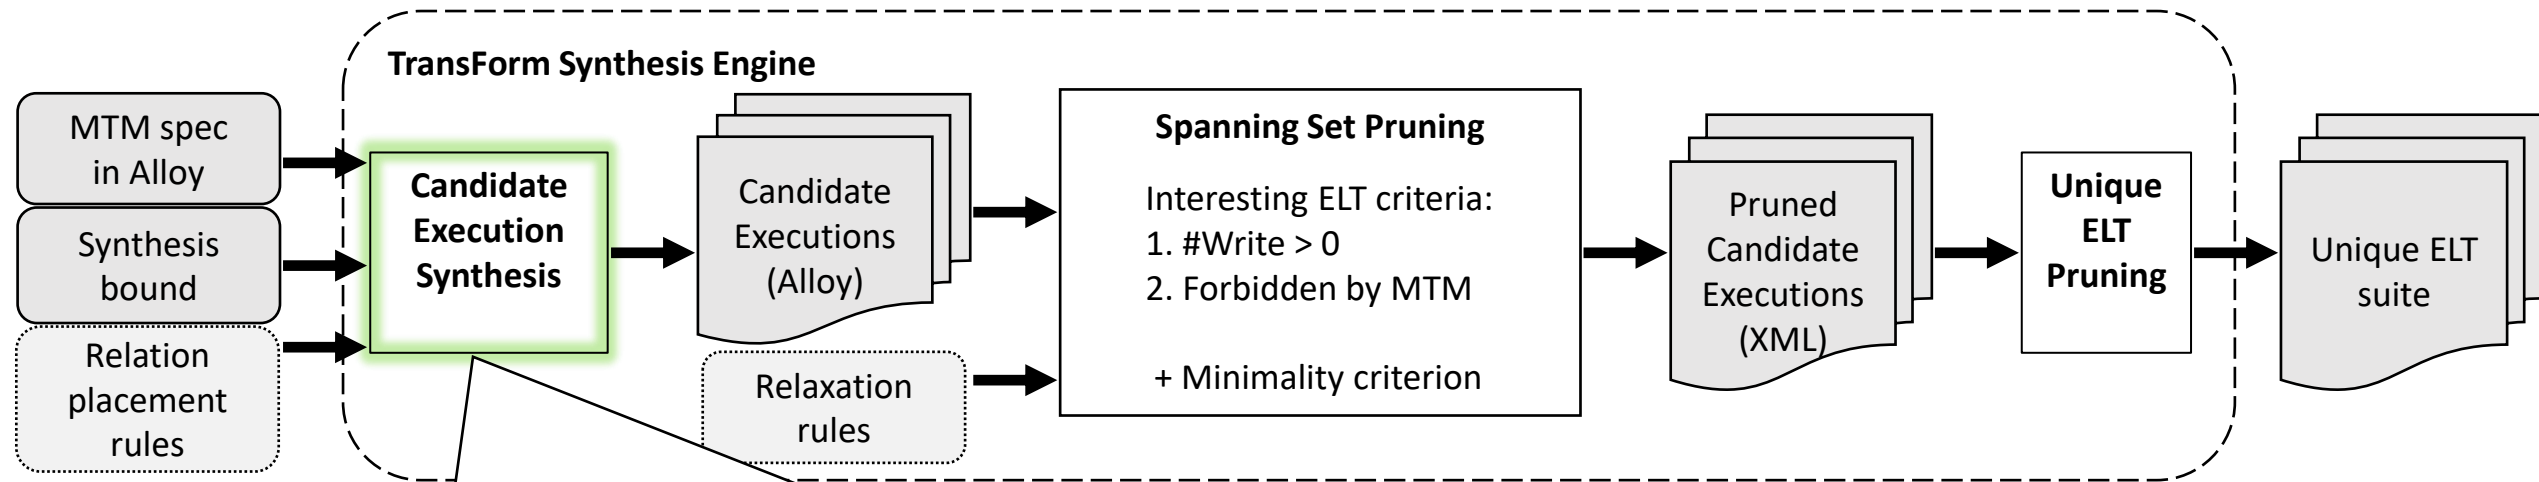


# From Specification to Test Synthesis

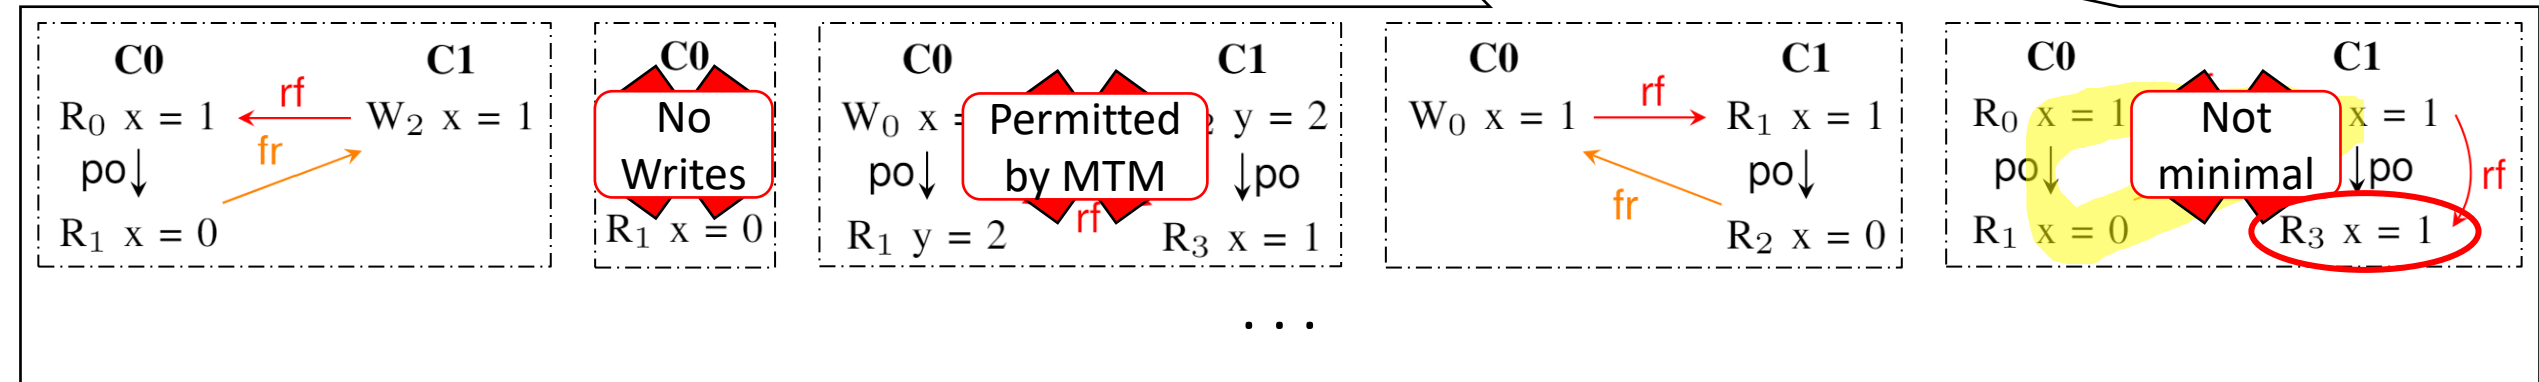
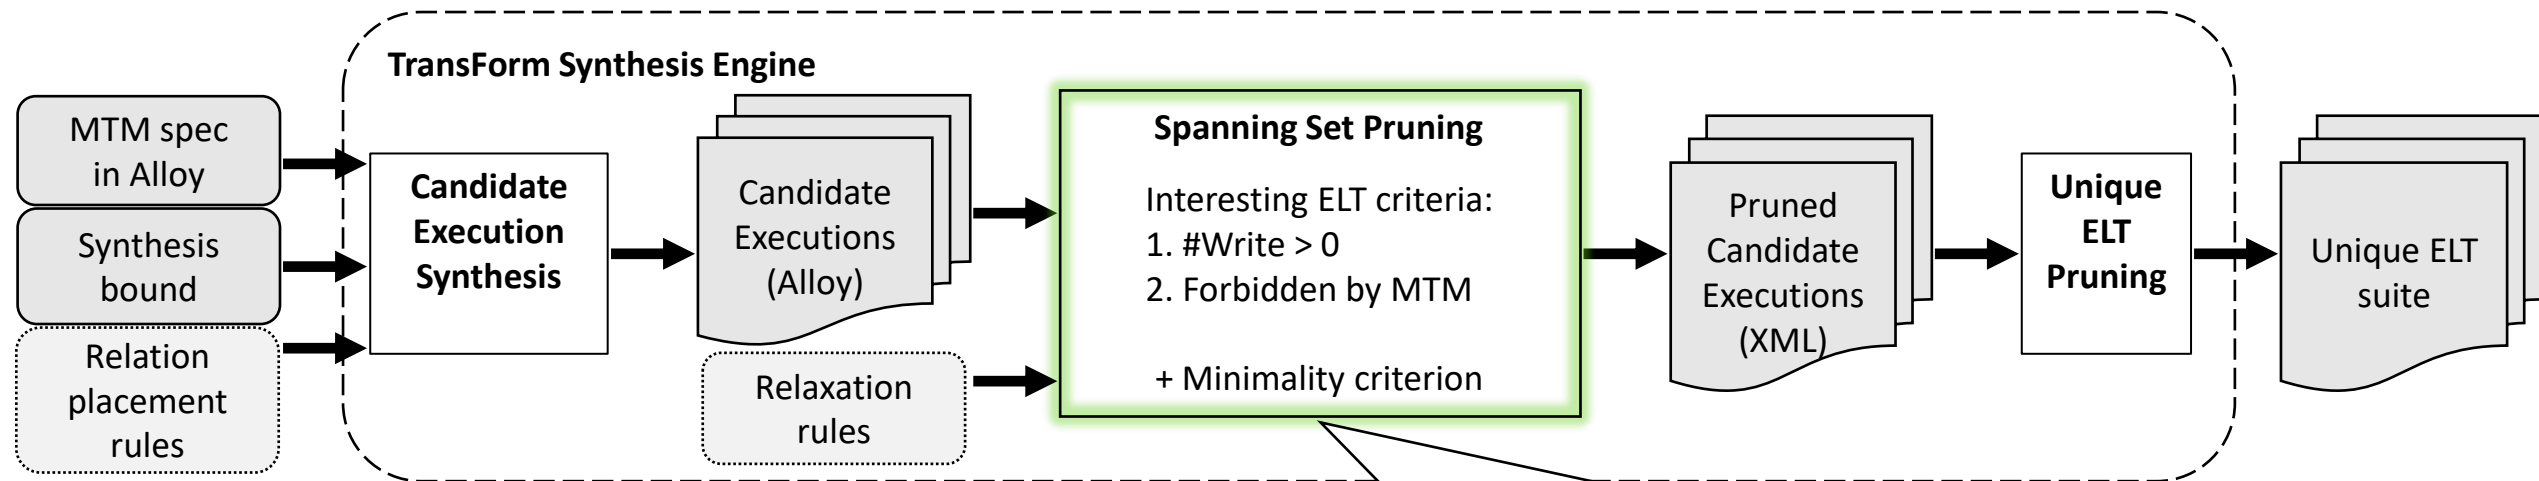


- ELTs can be described with MTM relations and support verification against an MTM spec
  - Goals:
    - Automated
    - Interesting and minimal (“Spanning set”)
    - Deduplicated
    - Comprehensive (to a bound)


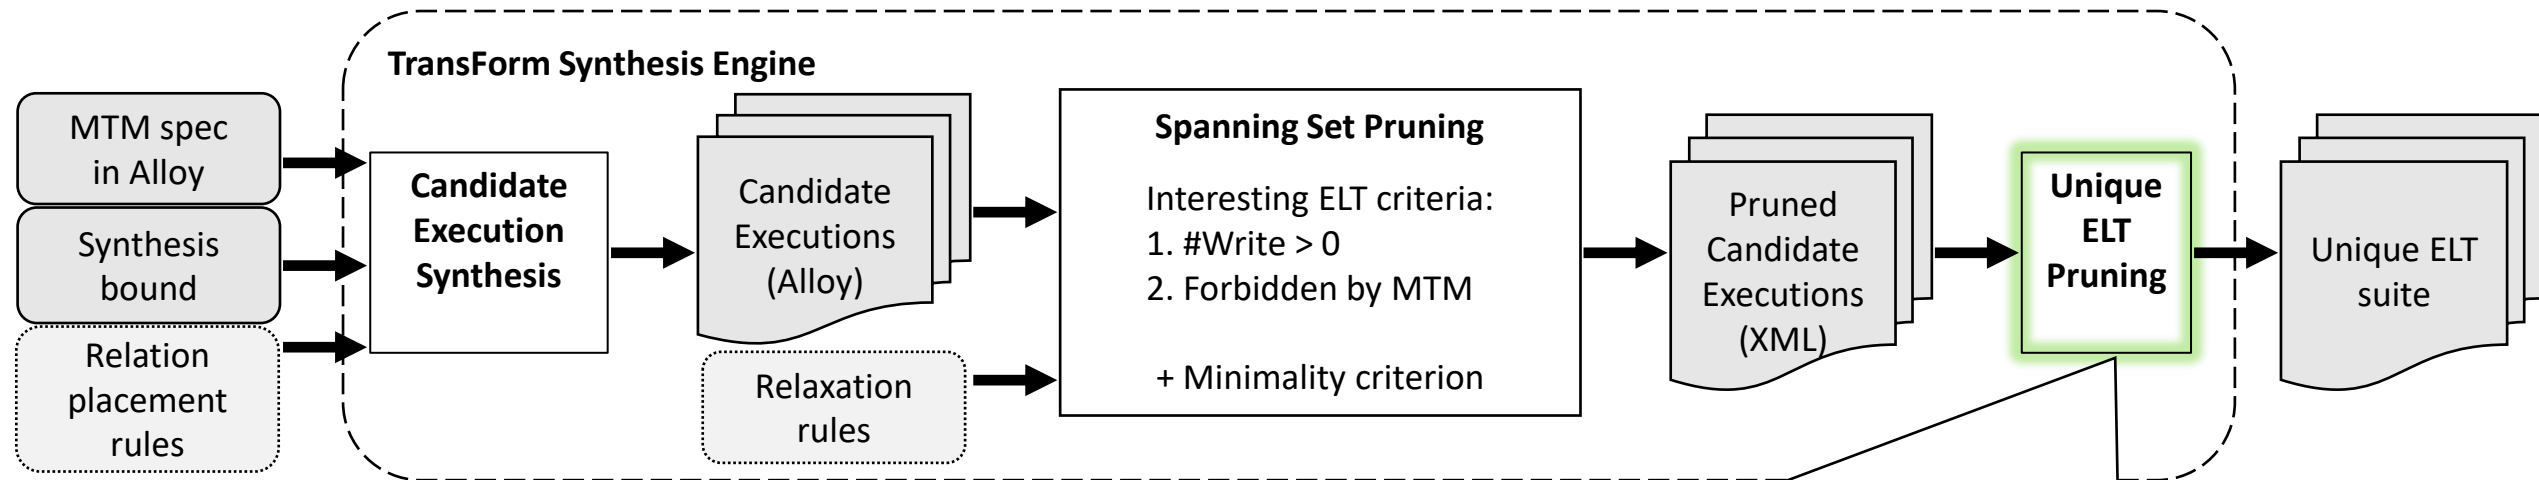
# TransForm's synthesis engine starts by synthesizing all possible candidate executions up to a bound



# Candidate executions are pruned for interesting ELT behaviors and checked for minimality

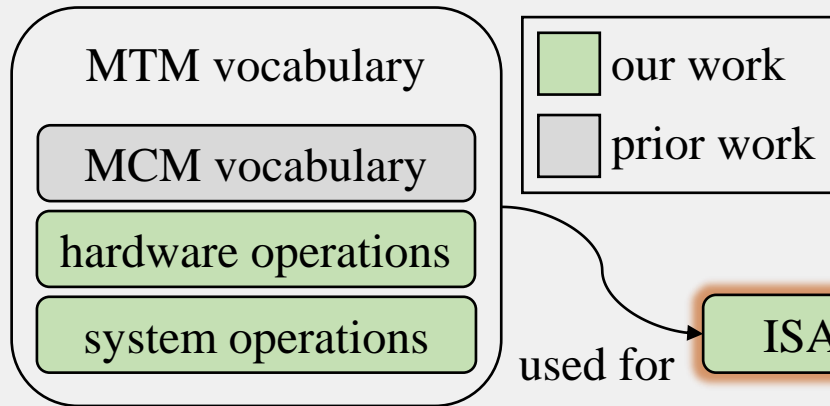


# Unique ELTs are found by deduplicating synthesized ELTs with a post-processing script




# Outline

## TransForm



- Background on ISA-level MCM vocabulary
  - Background on virtual memory systems
  - Novel ISA-level MTM vocabulary
  - Automating synthesis of ELTs
  - Case Study: an estimated MTM for x86
  - Future Work & Conclusions
- Prior Work
- My Work




# x86t\_elt predicates are composed of x86-TSO axioms and new transistency-specific axioms

- **x86t\_elt**: an approximate x86 transistency model based on prior work and publicly available documentation
- **x86-TSO**: `sc_per_loc`, `rmw_atomicity`, `causality`


**invlpg (required)**

`acylic[fr_va + remap + ^po]`



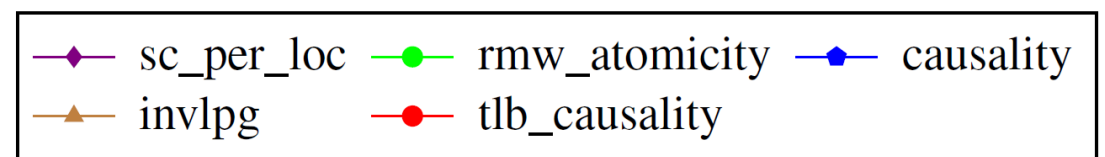
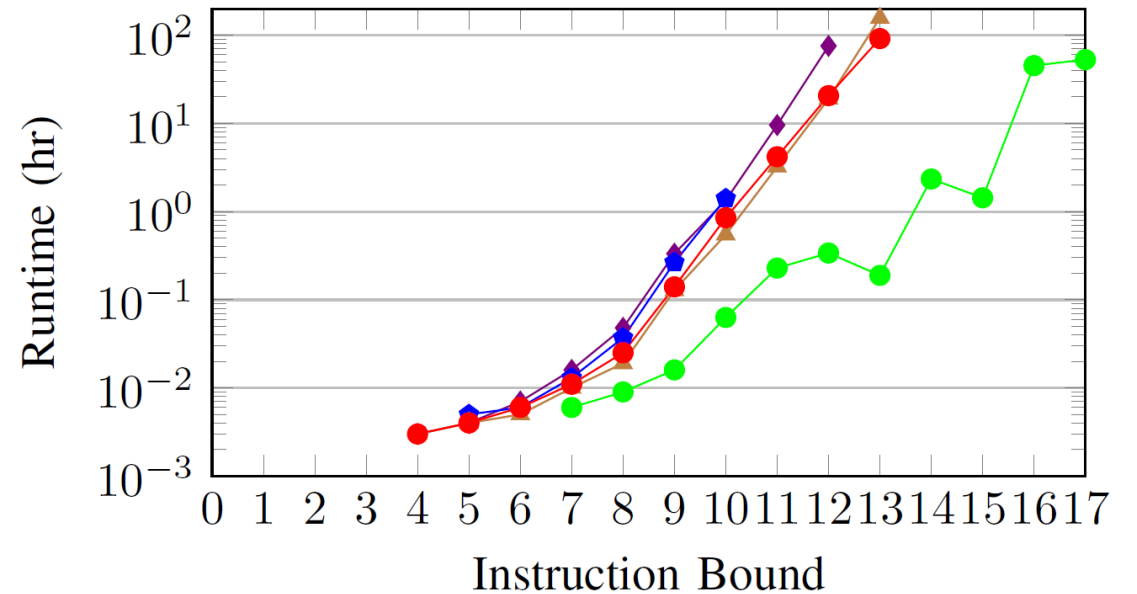
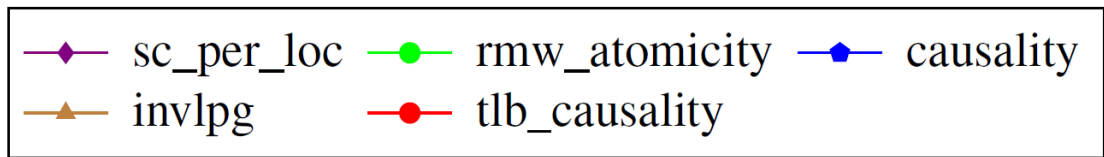

**tlb\_causality (auxiliary)**

`acylic[ptw_source + com]`





# A per-axiom suite was synthesized for each x86t\_elt axiom




140 total unique ELTs!

# The synthesized x86t\_elt suite consisted of all relevant ELTs from COATCheck and more

- All 22 relevant ELTs from COATCheck synthesized
  - 7 ELTs synthesized verbatim → map to 4 ELT programs in x86t\_elt suite
  - 15 ELTs can be reduced to a minimal ELT that is synthesized
- 4 ELTs from COATCheck, 136 new ELTs

# Outline

## TransForm



- Background on ISA-level MCM vocabulary
  - Background on virtual memory systems
  - Novel ISA-level MTM vocabulary
  - Automating synthesis of ELTs
  - Case Study: an estimated MTM for x86
  - Future Work & Conclusions
- Prior Work
- My Work



# Ongoing and Future Work

- Empirical MTM testing
- Validate x86t\_elt and verify x86 processors using synthesized ELTs
- Specify other MTMs (e.g., RISC-V)
- Model additional transistency interactions (e.g., permission bit updates)
- Formally reason about transistency and security

# Conclusions

- **TransForm**: framework for formal specification of MTMs and ELT synthesis
- Enables modern ISAs to have a formal specification that includes VM
- Offers systems programmers and hardware designers a stronger opportunity for verification of full systems
- Available at:  
<https://github.com/naorinh/TransForm>

# TransForm: Formally Specifying Transistency Models and Synthesizing Enhanced Litmus Tests

**Naorin Hossain**

Princeton University

*January 22, 2021*


<https://github.com/naorinh/TransForm>

# Backup Slides

---

Prior work developed a tool for verifying correctness of hardware MTM *implementations*

- **Enhanced litmus tests (ELTs)**: litmus tests enhanced with hardware- and system-level interactions that facilitate the VM abstraction

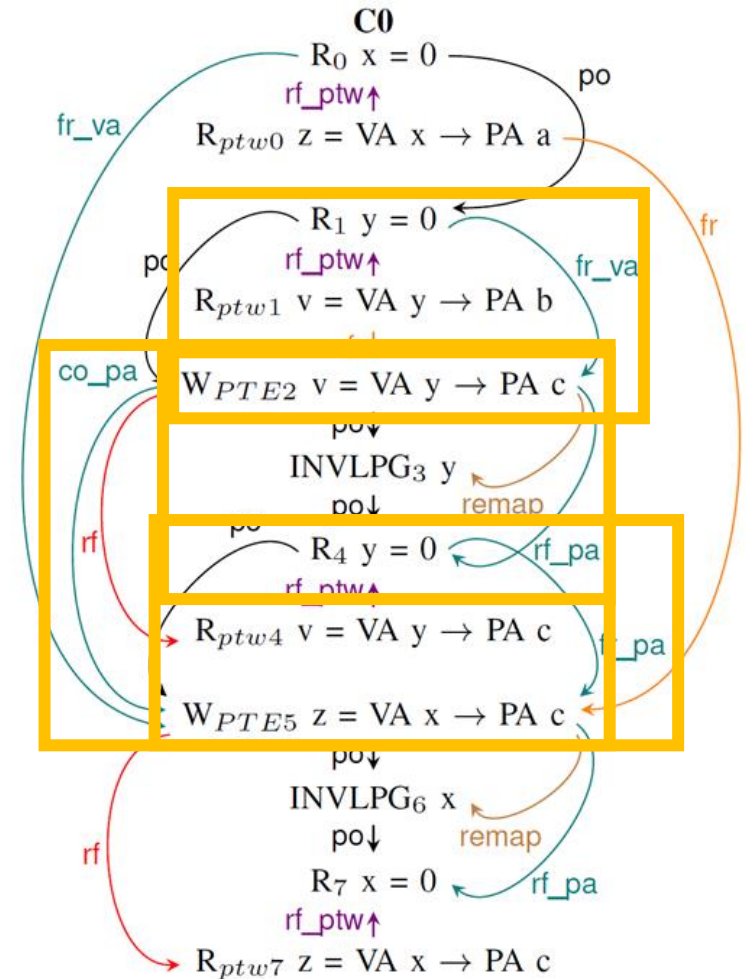


1. Formal MTM verification at *microarchitectural*-level
2. Lacked formal MTM specification
3. Hand-generated ELT suites



# Communication relations can be extended to apply to virtual-to-physical address mappings

- **rf\_pa**: maps a PTE write mapping VA  $v \rightarrow$  PA  $p$  to instructions that access VA  $v \rightarrow$  PA  $p$ .
- **fr\_pa**: maps instructions that access VA  $v \rightarrow$  PA  $p$  to PTE writes that map VA  $v' \rightarrow$  PA  $p$ .
- **co\_pa**: maps PTE writes that change mappings for different VAs to the same PA.
- **fr\_va**: maps instructions that access VA  $v \rightarrow$  PA  $p$  to PTE writes that map VA  $v \rightarrow$  PA  $p'$ .



# Tests with instructions that do not contribute to forbidden outcomes violate minimality criterion

**Minimality criterion:** any **relaxation** on the test results in satisfying the transistency predicates

**Relaxations** for transistency purposes are the removal of instructions and the following dependent instructions:

1. Invoked ghost instructions
2. Invoked INVLPGs
3. Dependent RMW operations

ELT execution should have a forbidden outcome that becomes legal under *every possible isolated* relaxation.

- **x86-TSO – causality axiom:**  
acyclic(rfe + co + fr + ppo + fence)
- Removal of  $W_4$  as relaxation does not result in satisfying causality

