


# THE SEMANTICS OF TRANSACTIONS AND WEAK MEMORY IN X86, POWER, ARM, AND C++



**Nathan Chong**  
Arm Ltd.



**Tyler Sorensen**  
Princeton University



**John Wickerson**  
Imperial College London

**USENIX Annual Technical Conference, 11 July 2019**

# WEAK MEMORY

<code>MOV [x] 1</code>	<code>MOV [y] 1</code>
<code>MOV r0 [y]</code>	<code>MOV r1 [x]</code>

`r0=1`

`r0=0`

`r0=1`

`r0=0`

`r1=1`






`r1=1`

`r1=0`

`r1=0`

SC

x86



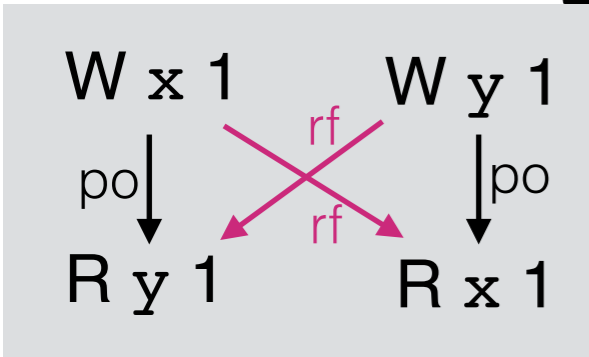
# WEAK MEMORY IS HARD!

- x86 proved tricky to formalise correctly [*Sarkar et al., POPL'09; Owens et al., TPHOLs'09*]
- Bug found in deployed "Power 5" processors [*Alglave et al., CAV'10*]
- C++ specification did not guarantee its own key property [*Batty et al., POPL'11*]
- Routine compiler optimisations are invalid under Java and C++ memory models [*Sevcik, PLDI'11; Vafeiadis et al. POPL'15*]
- Behaviour of NVIDIA graphics processors contradicted NVIDIA's programming guide [*Alglave et al., ASPLOS'15*]

# MODELLING WEAK MEMORY


```

MOV [x] 1 | MOV [y] 1
MOV r0 [y] | MOV r1 [x]
    
```




r0=1 r1=1

SC: ✓  
x86: ✓




r0=0 r1=1

SC: ✓  
x86: ✓



r0=1 r1=0

SC: ✓  
x86: ✓



r0=0 r1=0

SC: ✗  
x86: ✓

# TRANSACTIONAL MEMORY

- X86: `XBEGIN`  
...  
`XEND`
- Power: `tbegin`  
...  
`tend`
- ARM: `tstart`  
...  
`tcommit`
- C++: `atomic {`  
...  
`}`

## Transactional Memory: Architectural Support for Lock-Free Data Structures

Maurice Herlihy  
Digital Equipment Corporation  
Cambridge Research Laboratory  
Cambridge MA 02139  
herlihy@crl.dec.com

J. Eliot B. Moss  
Dept. of Computer Science  
University of Massachusetts  
Amherst, MA 01003  
moss@cs.umass.edu

### Abstract

A shared data structure is *lock-free* if its operations do not require mutual exclusion. If one process is interrupted in the middle of an operation, other processes will not be prevented from operating on that object. In highly concurrent systems, lock-free data structures avoid common problems associated with conventional locking techniques, including priority inversion, convoying, and difficulty of avoiding deadlock. This paper introduces *transactional memory*, a new multiprocessor architecture intended to make lock-free synchronization as efficient (and easy to use) as conventional techniques based on mutual exclusion. Transactional memory allows programmers to describe operations that apply

structures avoid common problems associated with conventional locking techniques in highly concurrent systems.

- *Priority inversion* occurs when a process is preempted while holding a lock, preventing higher-priority processes from executing.
- *Convoying* occurs when a process is blocked from being scheduled, perhaps by exhaustion of the system, by a page fault, or by some other reason. When such an interruption occurs, the process is unable to run until the system is able to run again.
- *Deadlock* can occur if processes hold locks on the same set of objects in different orders. Deadlock avoidance can be awkward if the objects are multiple data objects, particularly in a multiprocessor environment.

# WEAK MEMORY + TM = ?

**XBEGIN**

**MOV** [x] 1

**MOV** r0 [y]

**XEND**

**XBEGIN**

**MOV** [y] 1

**MOV** r1 [x]

**XEND**

r0=1

r0=0

r0=1

r0=0

r1=1

r1=1

r1=0

r1=0

SC

x86

# WEAK MEMORY + TM = ?

**XBEGIN**

**MOV [x] 1**

**MOV r0 [y]**

**XEND**

**XBEGIN**

**MOV [y] 1**

**MOV r1 [x]**

**XEND**

r0=1

r0=0

r0=1

r0=0

r1=1

r1=1

r1=0

r1=0

transactional SC

SC

x86



# BUILDING OUR MODELS

**x86:**

$\text{acyclic}(po_{\text{loc}} \cup \text{com})$  (COHERENCE)  
 $\text{empty}(rmw \cap (fr_e; co_e))$  (RMWISOL)  
 $\text{acyclic}(hb)$  (ORDER)  
 where  $ppo = ((W \times W) \cup (R \times W) \cup (R \times R)) \cap po$   
 $tfence = po \cap ((\neg stxn; stxn) \cup (stxn; \neg stxn))$   
 $L = \text{domain}(rmw) \cup \text{range}(rmw)$   
 $\text{implied} = [L]; po \cup po; [L] \cup tfence$   
 $hb = mfence \cup ppo \cup \text{implied} \cup rf_e \cup fr \cup co$   
 $\text{acyclic}(\text{stronglift}(\text{com}, stxn))$  (STRONGISOL)  
 $\text{acyclic}(\text{stronglift}(hb, stxn))$  (TXNORDER)

**Power:**

$\text{acyclic}(po_{\text{loc}} \cup \text{com})$  (COHERENCE)  
 $\text{empty}(rmw \cap (fr_e; co_e))$  (RMWISOL)  
 $\text{acyclic}(hb)$  (ORDER)  
 where  $ppo = (\text{preserved program order, elided})$   
 $tfence = po \cap ((\neg stxn; stxn) \cup (stxn; \neg stxn))$   
 $fence = \text{sync} \cup tfence \cup (\text{lwsync} \setminus (W \times R))$   
 $ihb = ppo \cup fence$   
 $thb = (rf_e \cup ((fr_e \cup co_e)^*; ihb)^*; (fr_e \cup co_e)^*; rf_e^?)$   
 $hb = (rf_e^?; ihb; rf_e^?) \cup \text{weaklift}(thb, stxn)$   
 $\text{acyclic}(co \cup \text{prop})$  (PROPAGATION)  
 where  $efence = rf_e^?; fence; rf_e^?$   
 $\text{prop}_1 = [W]; efence; hb^*; [W]$   
 $\text{prop}_2 = \text{com}_e^*; efence^*; hb^*; (\text{sync} \cup tfence); hb^*$   
 $tprop_1 = rf_e; stxn; [W]$   
 $tprop_2 = stxn; rf_e$   
 $\text{prop} = \text{prop}_1 \cup \text{prop}_2 \cup tprop_1 \cup tprop_2$   
 $\text{irreflexive}(fr_e; \text{prop}; hb^*)$  (OBSERVATION)  
 $\text{acyclic}(\text{stronglift}(\text{com}, stxn))$  (STRONGISOL)  
 $\text{acyclic}(\text{stronglift}(hb, stxn))$  (TXNORDER)  
 $\text{empty}(rmw \cap tfence^*)$  (TXNCANCELSRMW)

**ARM:**

$\text{acyclic}(po_{\text{loc}} \cup \text{com})$  (COHERENCE)  
 $\text{acyclic}(ob)$  (ORDER)  
 where  $dob = (\text{order imposed by dependencies, elided})$   
 $aob = (\text{order imposed by atomic RMWs, elided})$   
 $bob = (\text{order imposed by barriers, elided})$   
 $tfence = po \cap ((\neg stxn; stxn) \cup (stxn; \neg stxn))$   
 $ob = \text{com}_e \cup dob \cup aob \cup bob \cup tfence$   
 $\text{empty}(rmw \cap (fr_e; co_e))$  (RMWISOL)  
 $\text{acyclic}(\text{stronglift}(\text{com}, stxn))$  (STRONGISOL)  
 $\text{acyclic}(\text{stronglift}(ob, stxn))$  (TXNORDER)  
 $\text{empty}(rmw \cap tfence^*)$  (TXNCANCELSRMW)

**C++:**

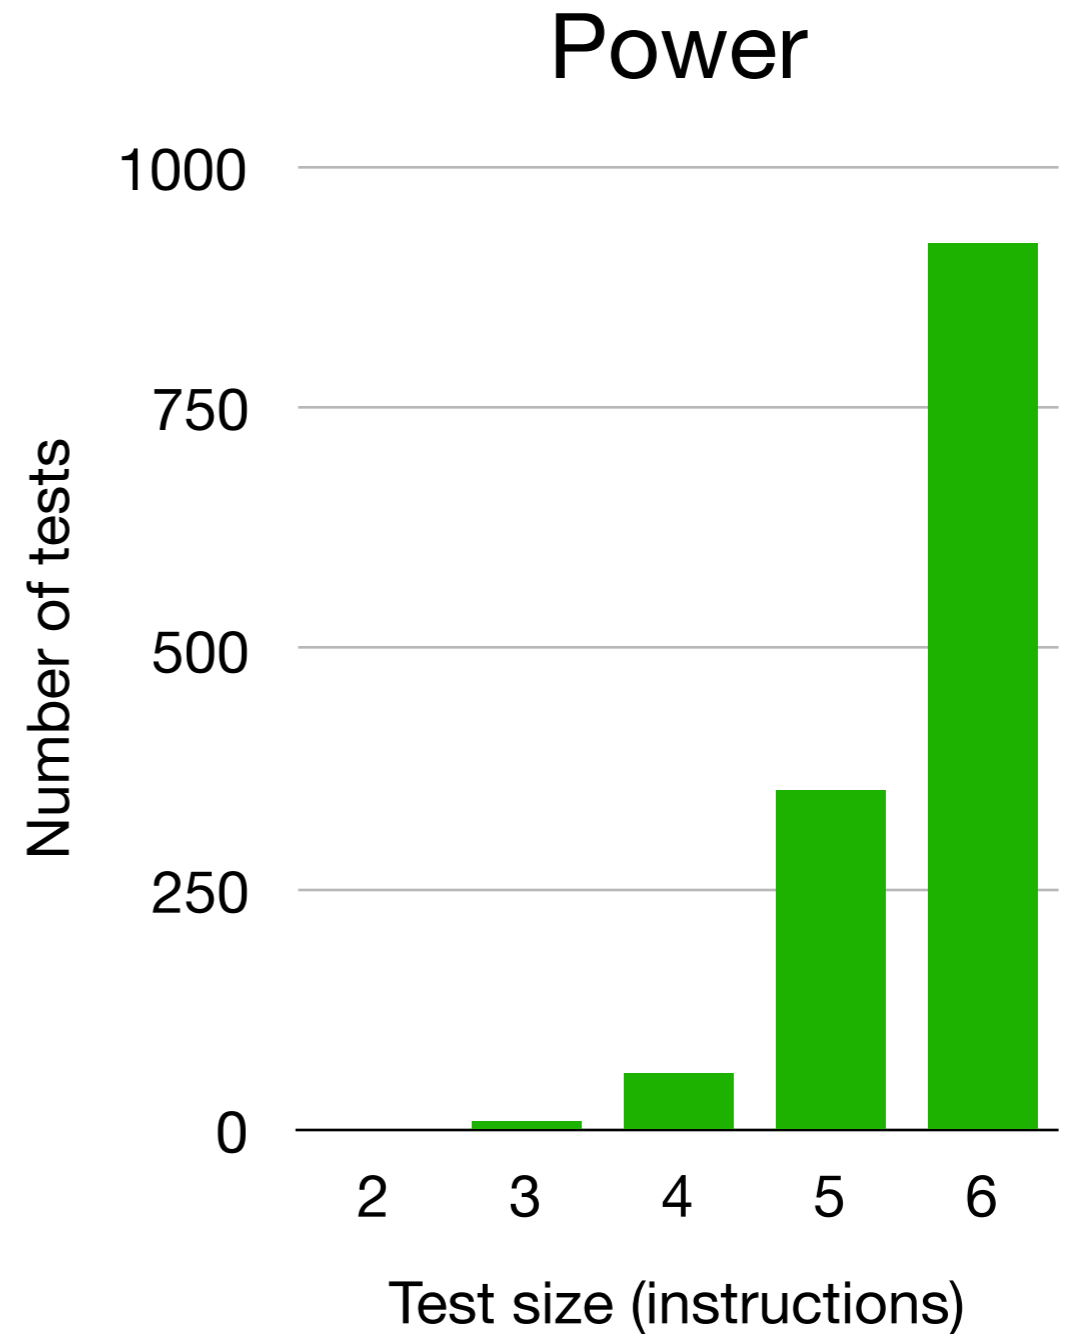

$\text{irreflexive}(hb; \text{com}^*)$  (HBCom)  
 where  $sw = (\text{synchronises-with, elided})$   
 $ecom = \text{com} \cup (co; rf)$   
 $tsw = \text{weaklift}(ecom, stxn)$   
 $hb = (sw \cup tsw \cup po)^+$   
 $\text{empty}(rmw \cap (fr_e; co_e))$  (RMWISOL)  
 $\text{acyclic}(po \cup rf)$  (NoTHINAIR)  
 $\text{acyclic}(psc)$  (SEQCST)  
 where  $psc = (\text{constraints on SC events, elided})$   
 $\text{empty}(\text{cnf} \setminus \text{Ato}^2 \setminus (hb \cup hb^{-1}))$  (NoRACE)  
 where  $\text{cnf} = ((W \times W) \cup (R \times W) \cup (W \times R)) \cap \text{sloc} \setminus id$

# VALIDATING OUR MODELS

1. Consult architecture manuals.
2. Interview engineers.
3. Check models have reasonable mathematical properties (e.g. adding/extending/coalescing transactions is safe).
4. Check that models validate existing compiler mappings.
5. Generate conformance tests and run them on hardware.

# VALIDATING OUR MODELS

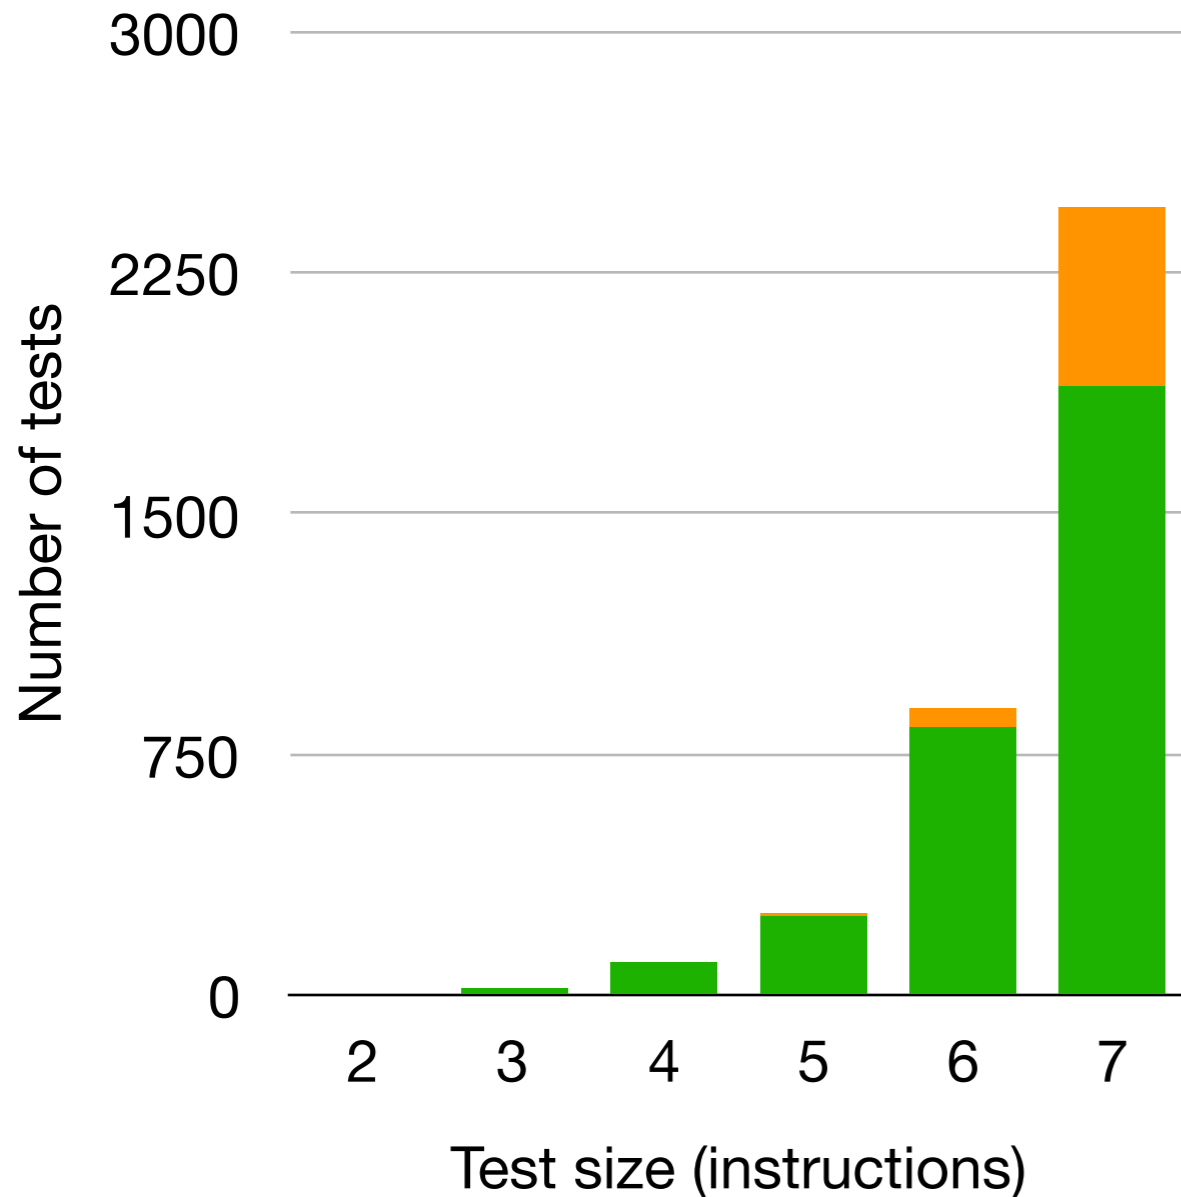
Behaviours that must be forbidden




# VALIDATING OUR MODELS

Behaviours that should be allowed

x86



Power



# USING OUR MODELS

# LOCK ELISION

```
lock()  
X = X + 2  
unlock()
```

```
lock()  
X = 1  
unlock()
```

# LOCK ELISION

```
lock()  
ldr W5,[X]  
add W5,W5,#2  
str W5,[X]  
unlock()
```

```
lock()  
mov W7,#1  
str W7,[X]  
unlock()
```

# LOCK ELISION

Loop:

```
ldaxr W2,[M]  
cbnz W2,Loop  
mov W3,#1  
stxr W4,W3,[M]  
cbnz W4,Loop  
ldr W5,[X]  
add W5,W5,#2  
str W5,[X]  
stlr WZR,[M]
```

```
lock()  
mov W7,#1  
str W7,[X]  
unlock()
```



# LOCK ELISION

Loop:

```
ldaxr W2,[M]
cbnz W2,Loop
mov W3,#1
stxr W4,W3,[M]
cbnz W4,Loop
ldr W5,[X]
add W5,W5,#2
str W5,[X]
stlr WZR,[M]
```

**tstart**

```
ldr W6,[M]
```

```
cbz W6,L1
```

**tcancel**

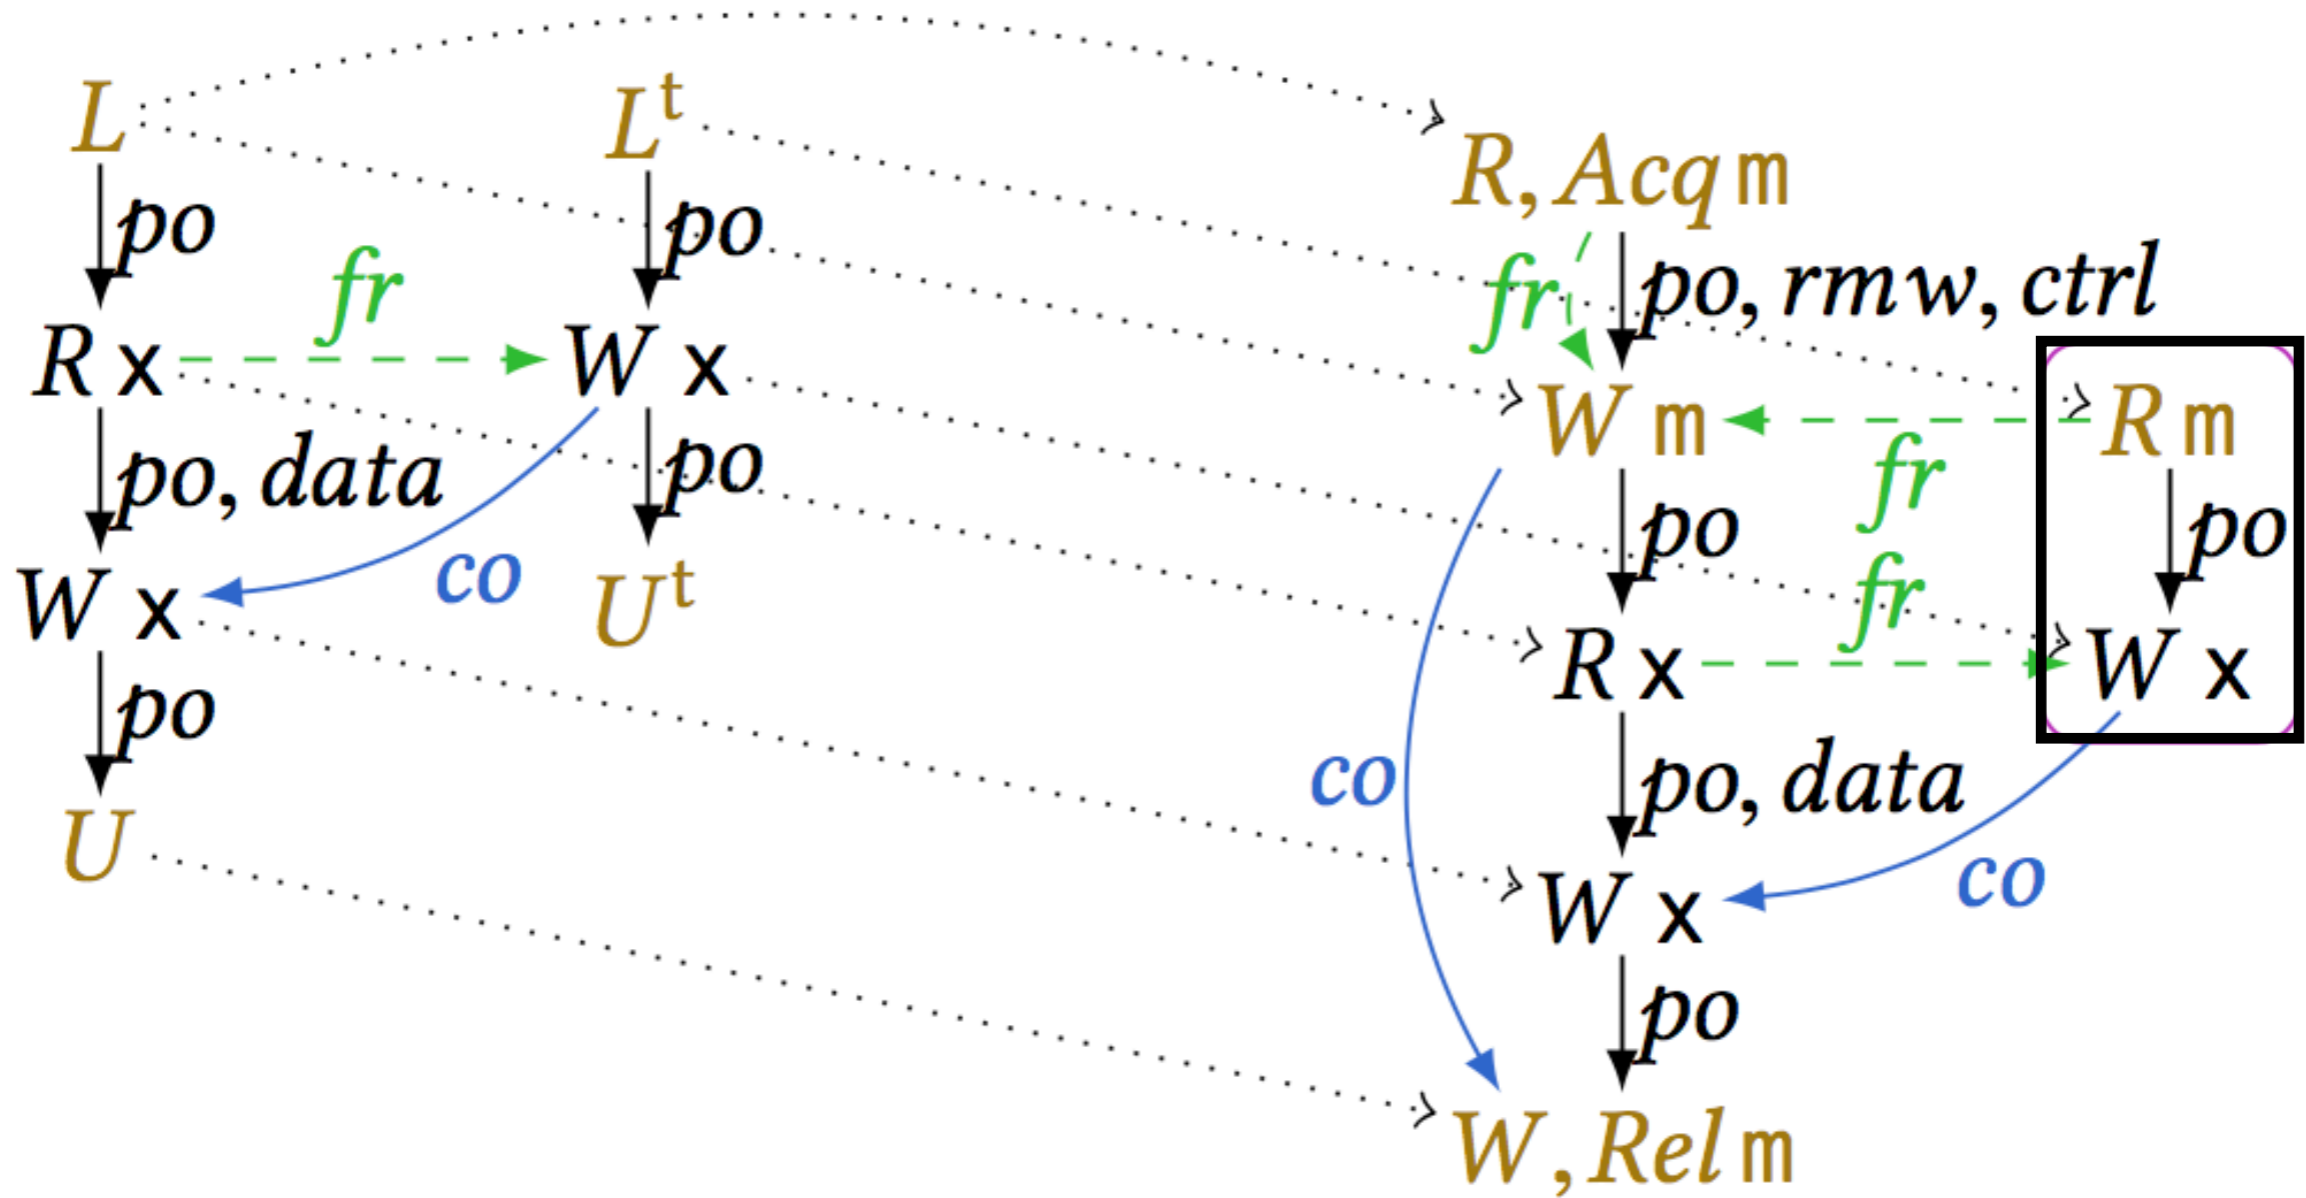
L1:

```
mov W7,#1
```

```
str W7,[X]
```

**tcommit**


# LOCK ELISION



# CONCLUSION

- Weak memory is pervasive, and transactional memory is entering the mainstream.
- We have designed and validated formal models of how these features interact in x86, Power, ARM, and C++.
- Weak memory + transactions + lock elision = tricky!

# THE SEMANTICS OF TRANSACTIONS AND WEAK MEMORY IN X86, POWER, ARM, AND C++



**Nathan Chong**  
Arm Ltd.



**Tyler Sorensen**  
Princeton University



**John Wickerson**  
Imperial College London

**USENIX Annual Technical Conference, 11 July 2019**