

# What's Decidable about Causally Consistent Shared Memory?

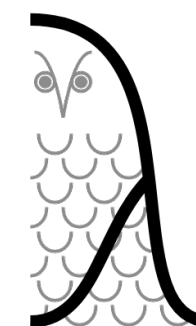
**Ori Lahav**



**Udi Boker**



OWLS  
January 20, 2021

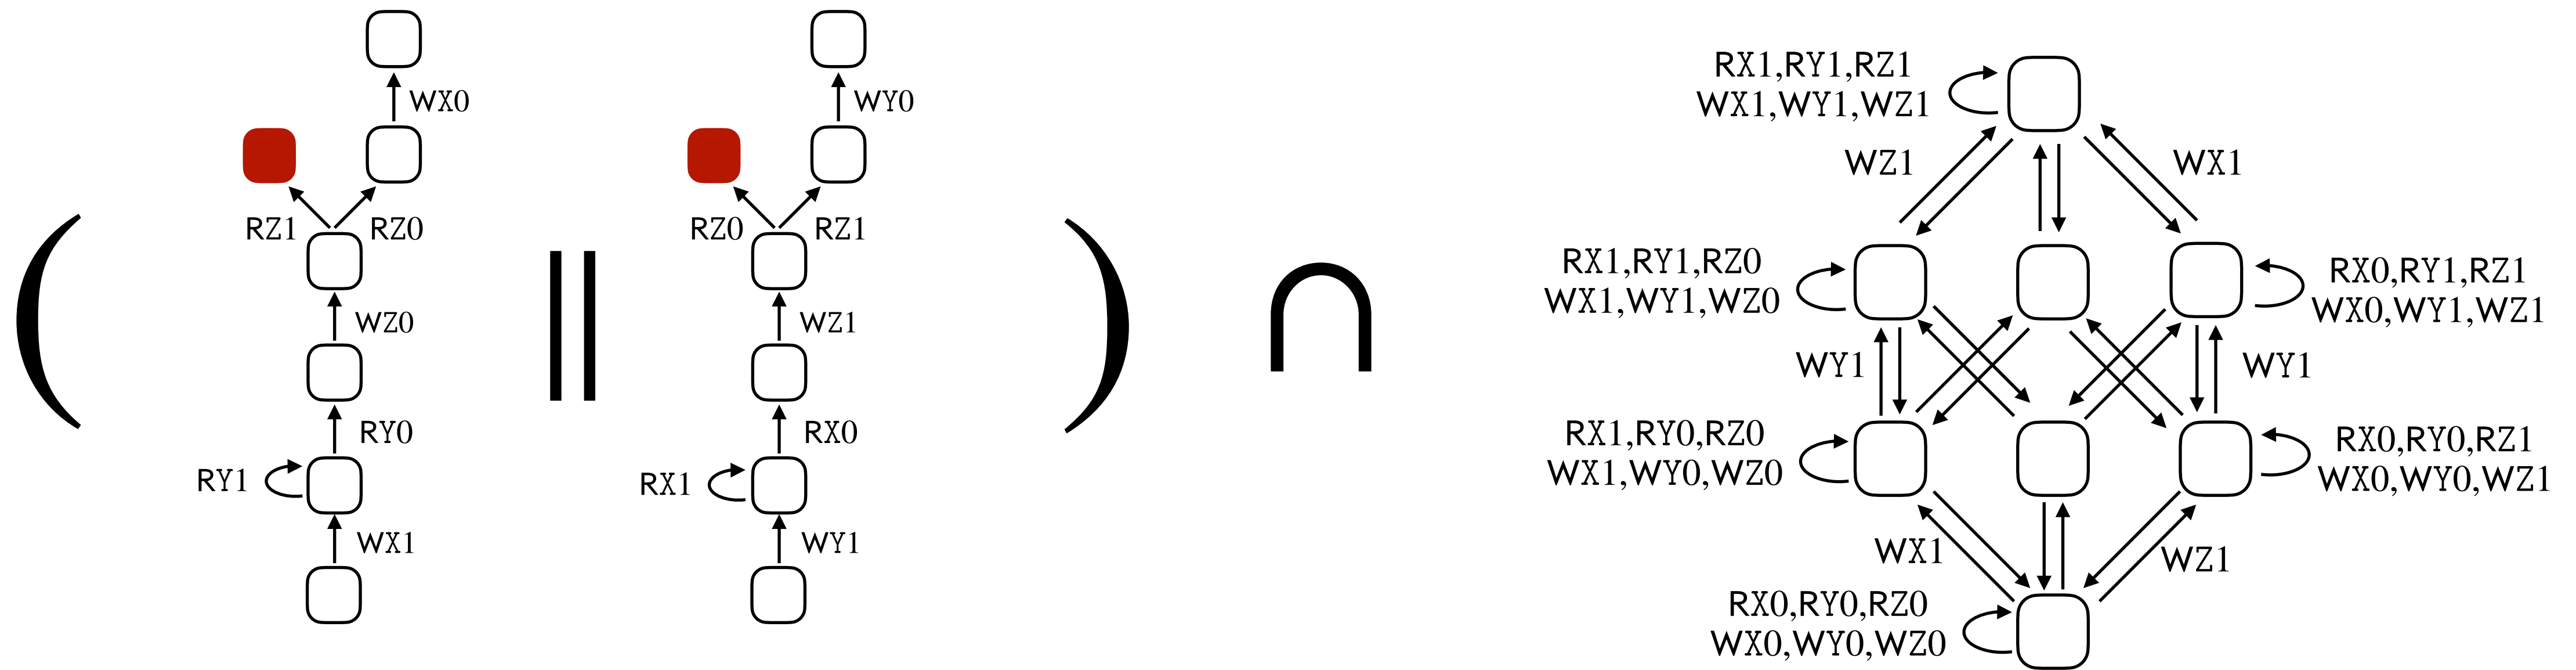


# Safety Verification

```
X := 1;  
repeat  
  a := Y;  
until (a = 0);  
Z := 0;  
assert (Z = 0);  
X := 0;
```

```
Y := 1;  
repeat  
  b := X;  
until (b = 0);  
Z := 1;  
assert (Z = 1);  
Y := 0;
```

# Simple Solution

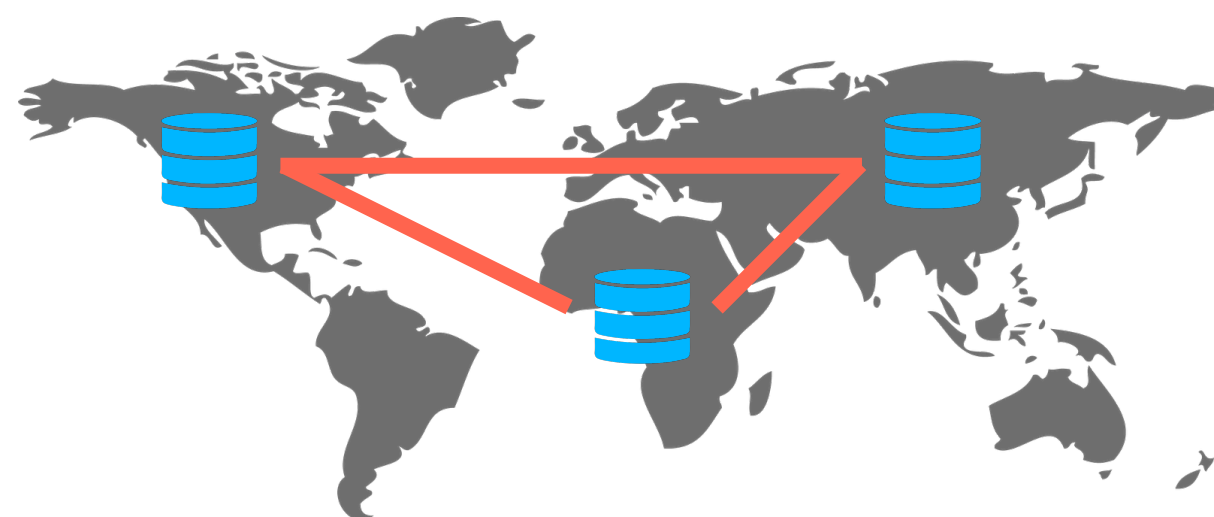
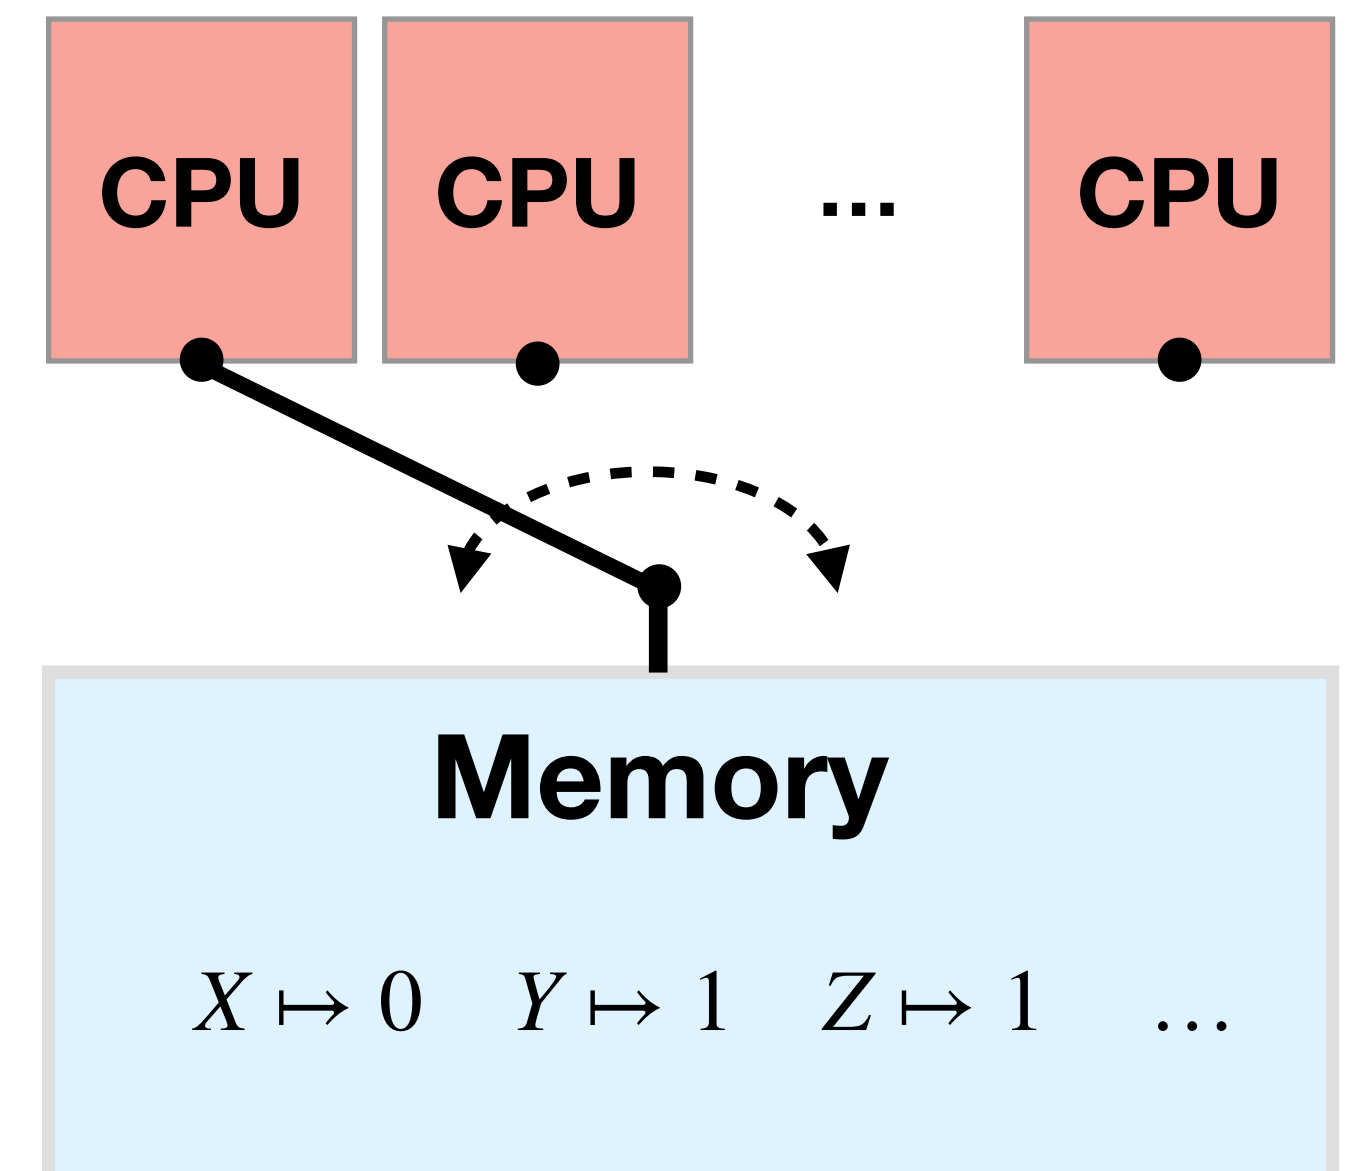


For programs with a **bounded data domain**, this problem is clearly **decidable**:

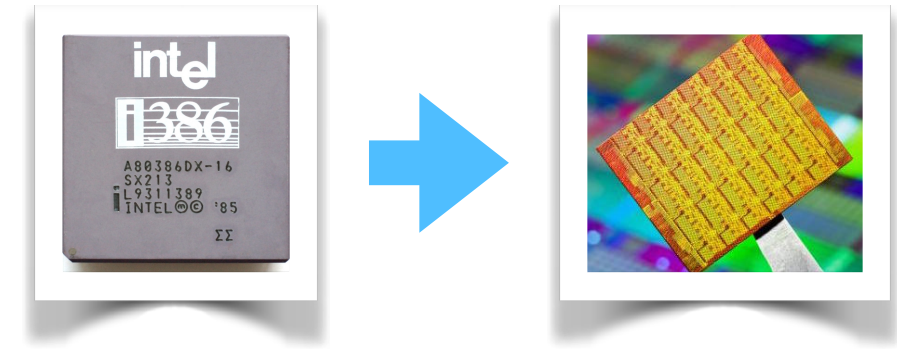
- Reduction to reachability in finite-state systems
- PSPACE-complete

# Sequential Consistency (SC)

- We assumed the classical shared-memory model.
- But it is **unrealistic**:  
for better **performance/scalability/availability/fault-tolerance** shared-memory implementations provide weaker semantics.
- Similar situation in *distributed data-stores*



# Weaker Memory Models



- x86-TSO →
- POWER
- ARM
- RISC-V
- C/C++11 →
- many many more...

## Decidable

[Atig, Bouajjani, Burckhardt, Musuvathi. POPL'2010]

[Abdulla, Atig, Bouajjani, Ngo. CUNCUR'2016]

## Undecidable

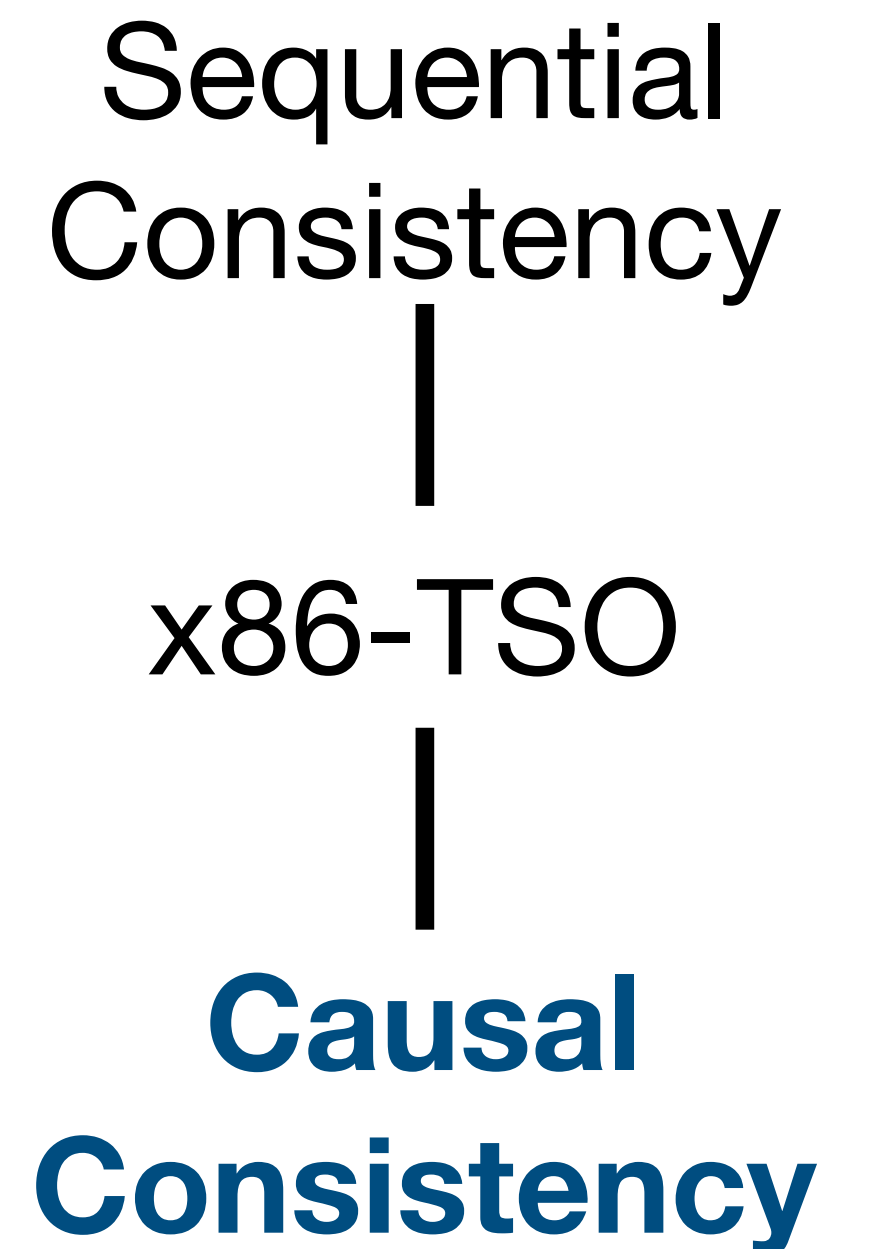
[Abdulla, Arora, Atig, Krishna. PLDI'2019]

- Even for the **Release/Acquire** fragment:  
`memory_order_release & memory_order_acquire`
- Reduction from *Post correspondence problem*

bba	ab	bba	a
bb	aa	oo	baa

# Causal Consistency

- A classical model originated from *replicated data stores*:
  - nodes may disagree on the order of some operations
  - consensus on the order of “*causally related*” operations
- Relatively simple and intuitive but more scalable than SC



# Safety Verification

```
X := 1;  
repeat  
  a := Y; // 0  
until (a = 0);  
Z := 0;  
assert (Z = 0);  
X := 0;
```

```
Y := 1;  
repeat  
  b := X; // 0  
until (b = 0);  
Z := 1;  
assert (Z = 1);  
Y := 0;
```

**Both can read 0 in the same execution!**

# Is it a Problem?

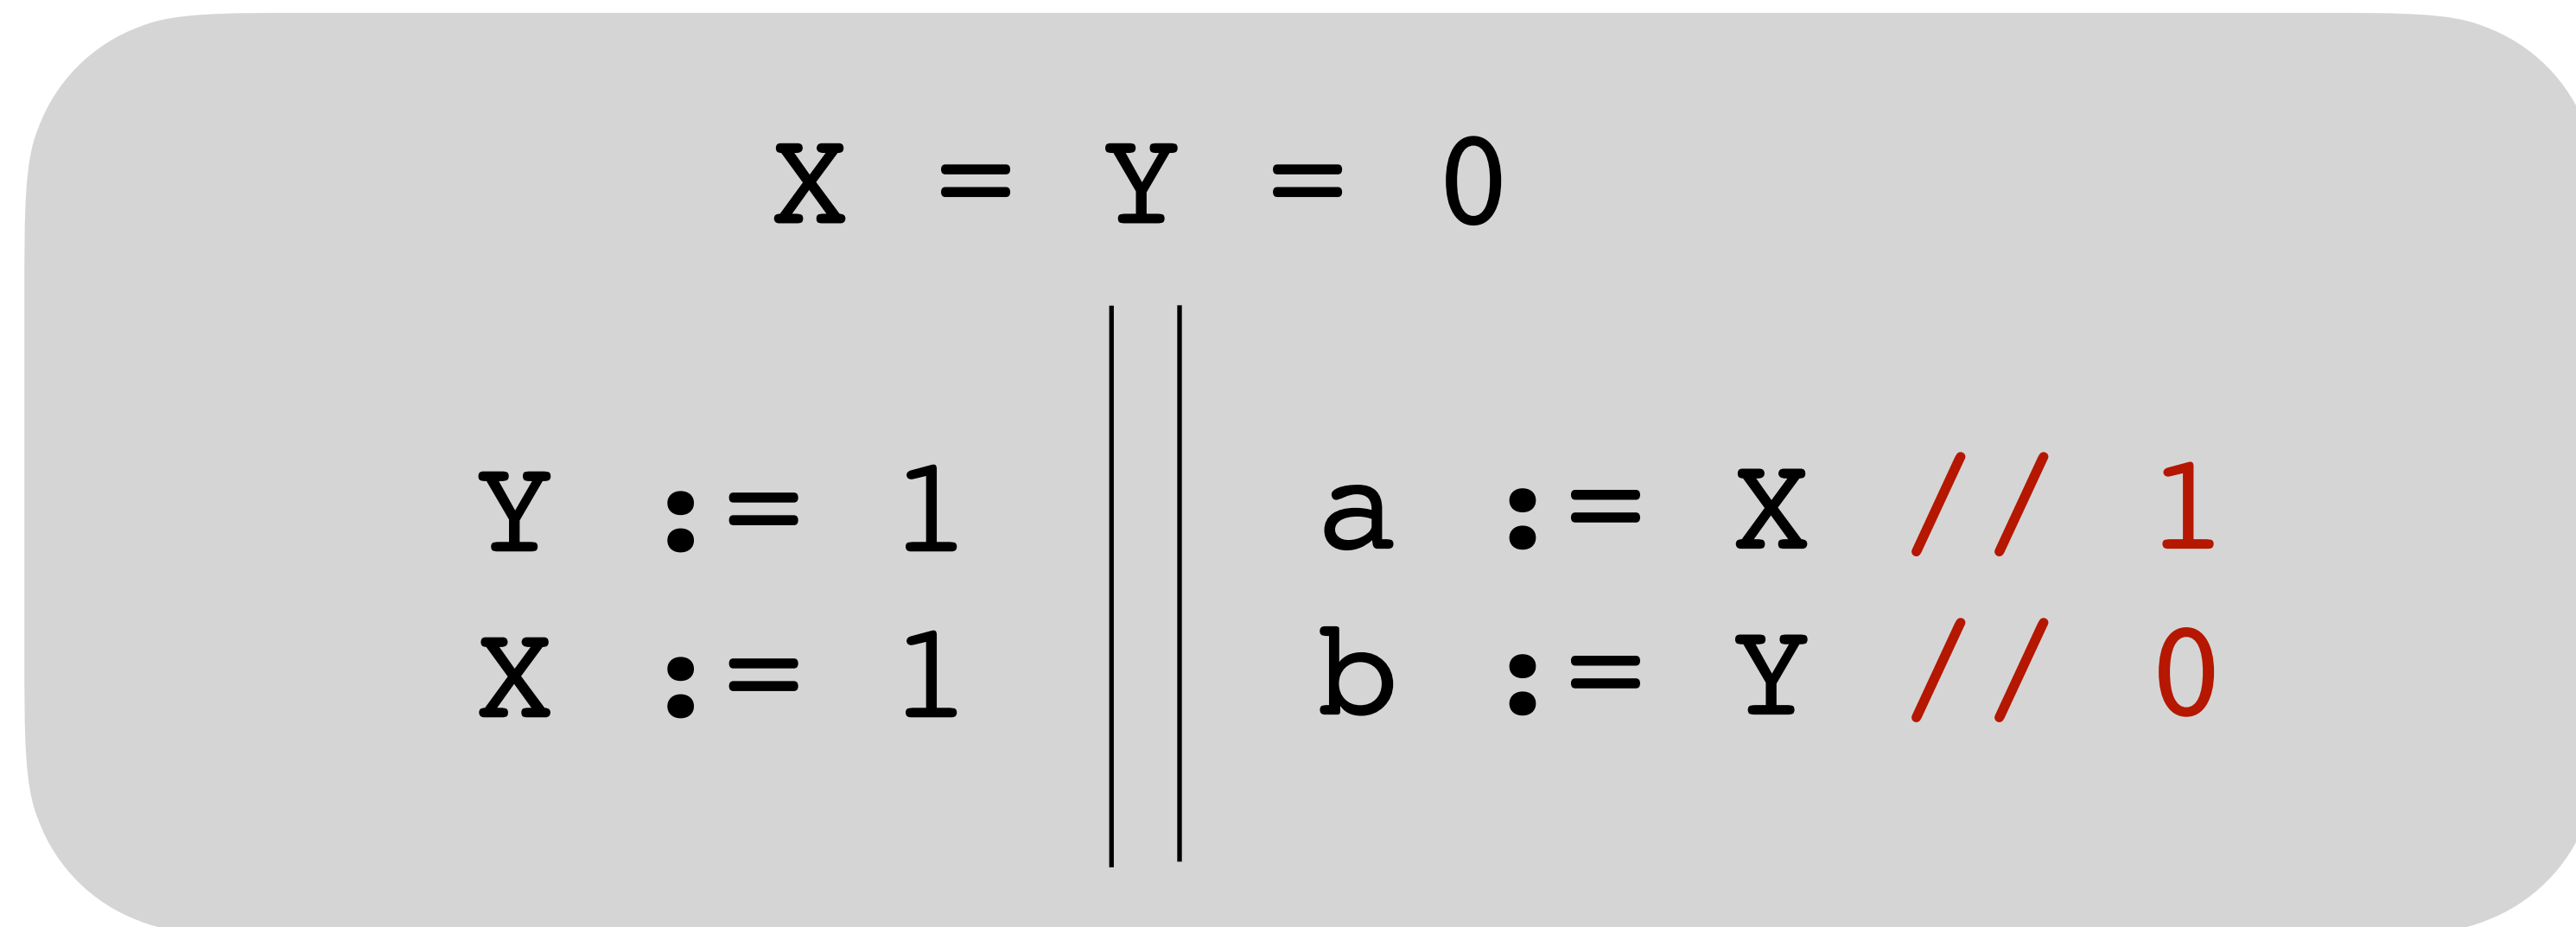
[https://en.wikipedia.org/wiki/Dekker's\\_algorithm](https://en.wikipedia.org/wiki/Dekker's_algorithm)

- How come *airplanes don't crash*?
- There are ways to **demand sequential consistency** when we need it.
- We often **don't need** sequential consistency in its full power.

*We have to define and understand  
the semantics of shared-memory concurrency.*



# Flag-Based Synchronization

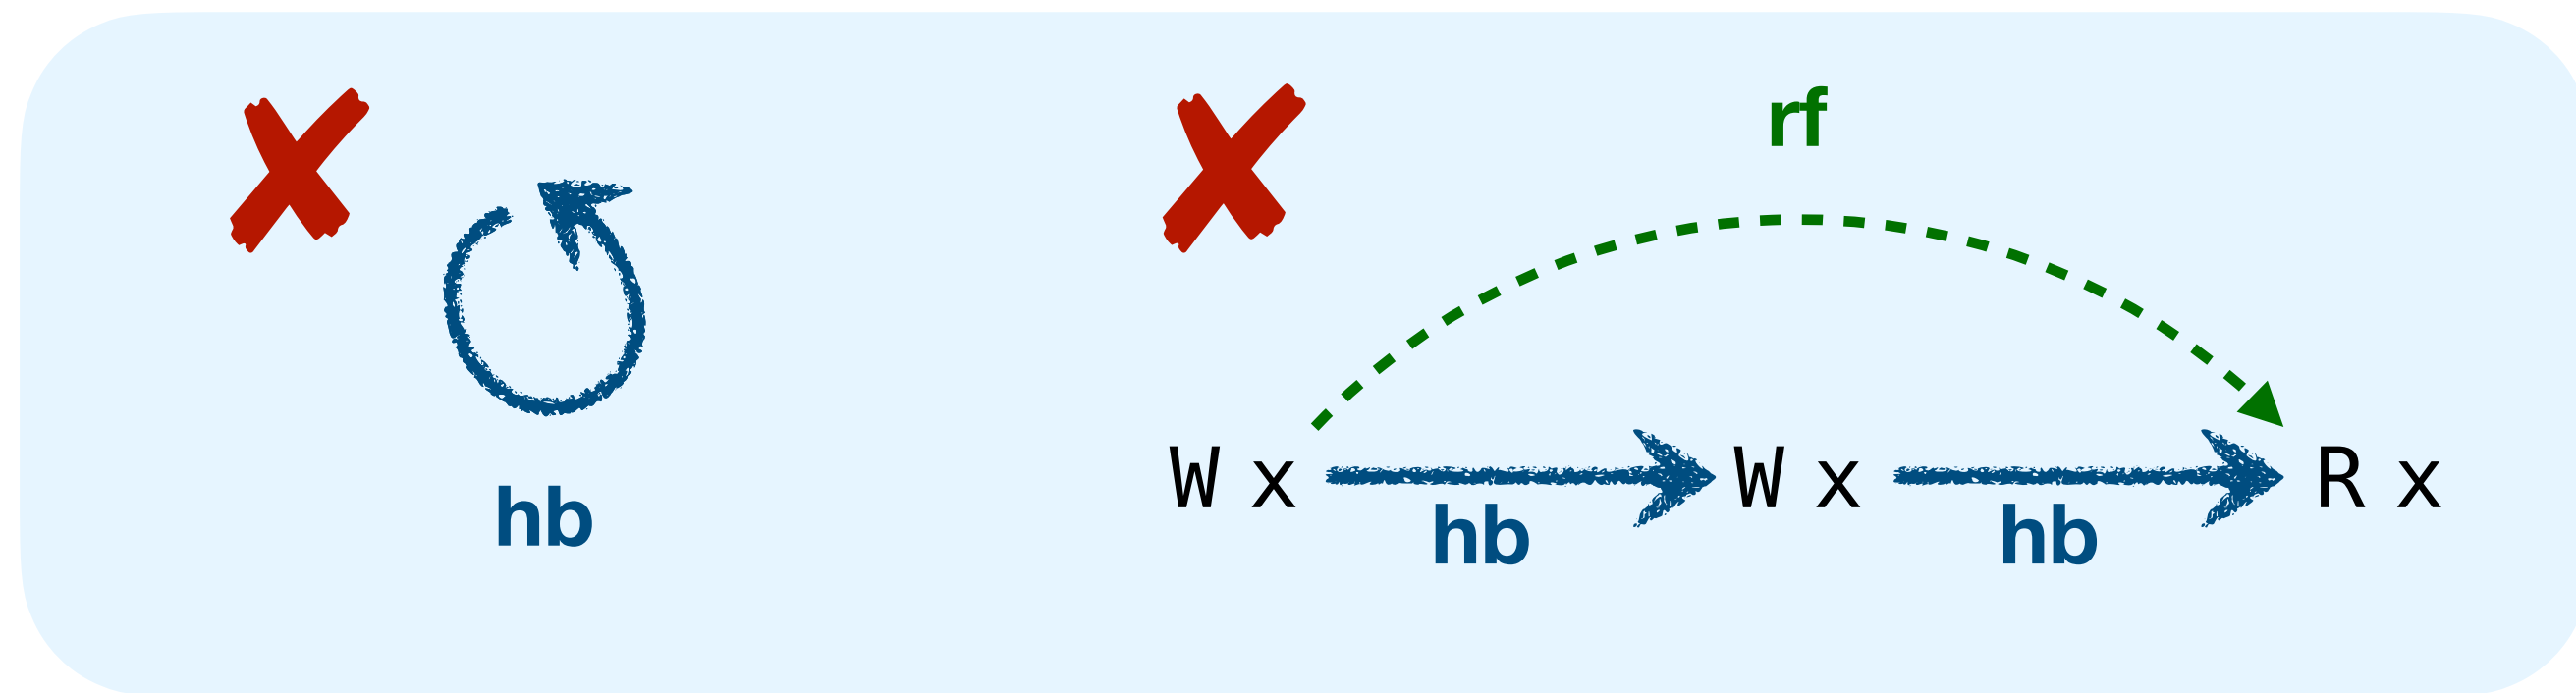


**X** This behavior is forbidden under causal consistency

# Formal Semantics

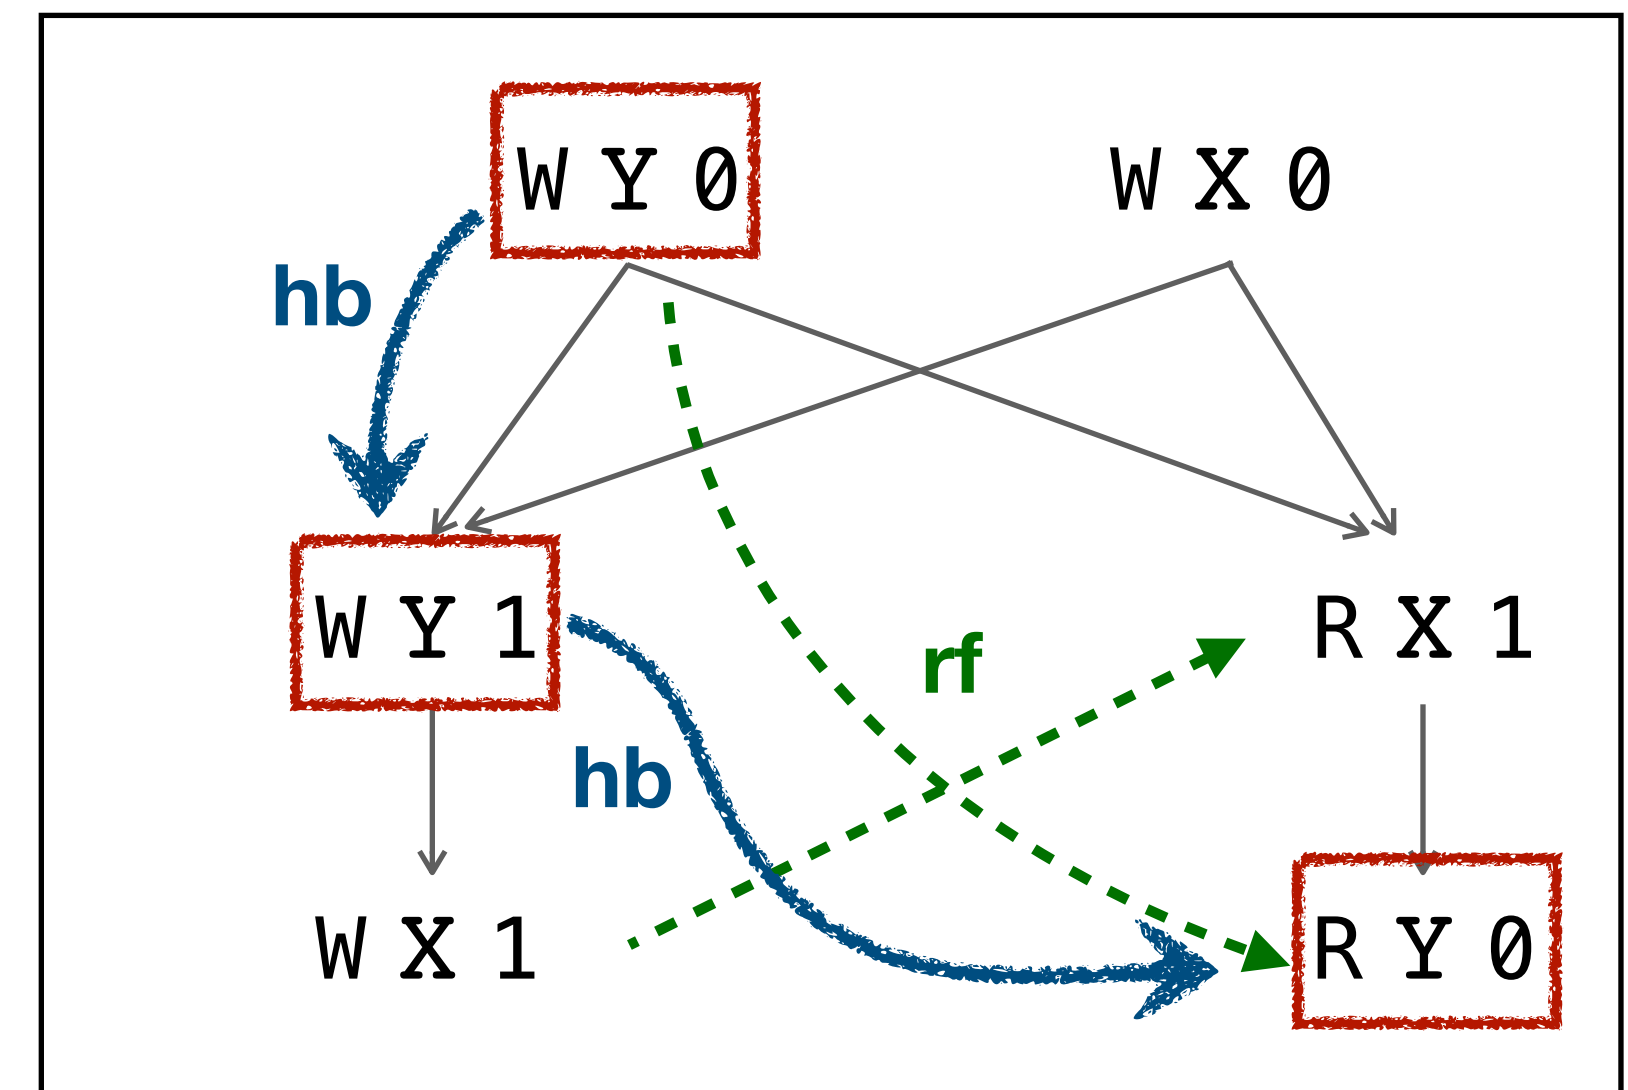
Defined **declaratively** using *execution graphs*

**happens-before** = ( program-order  $\cup$  **reads-from** )<sup>+</sup>



```

X = Y = 0
Y := 1 || a := X // 1
X := 1 || b := Y // 0
    
```



**inconsistent** execution graph  
**disallowed** program outcome

```

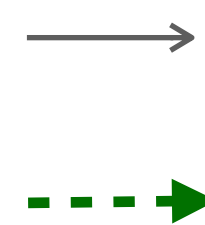
X := 1;
repeat
  a := Y;
until (a = 0);
Z := 0;
assert (Z = 0);

```

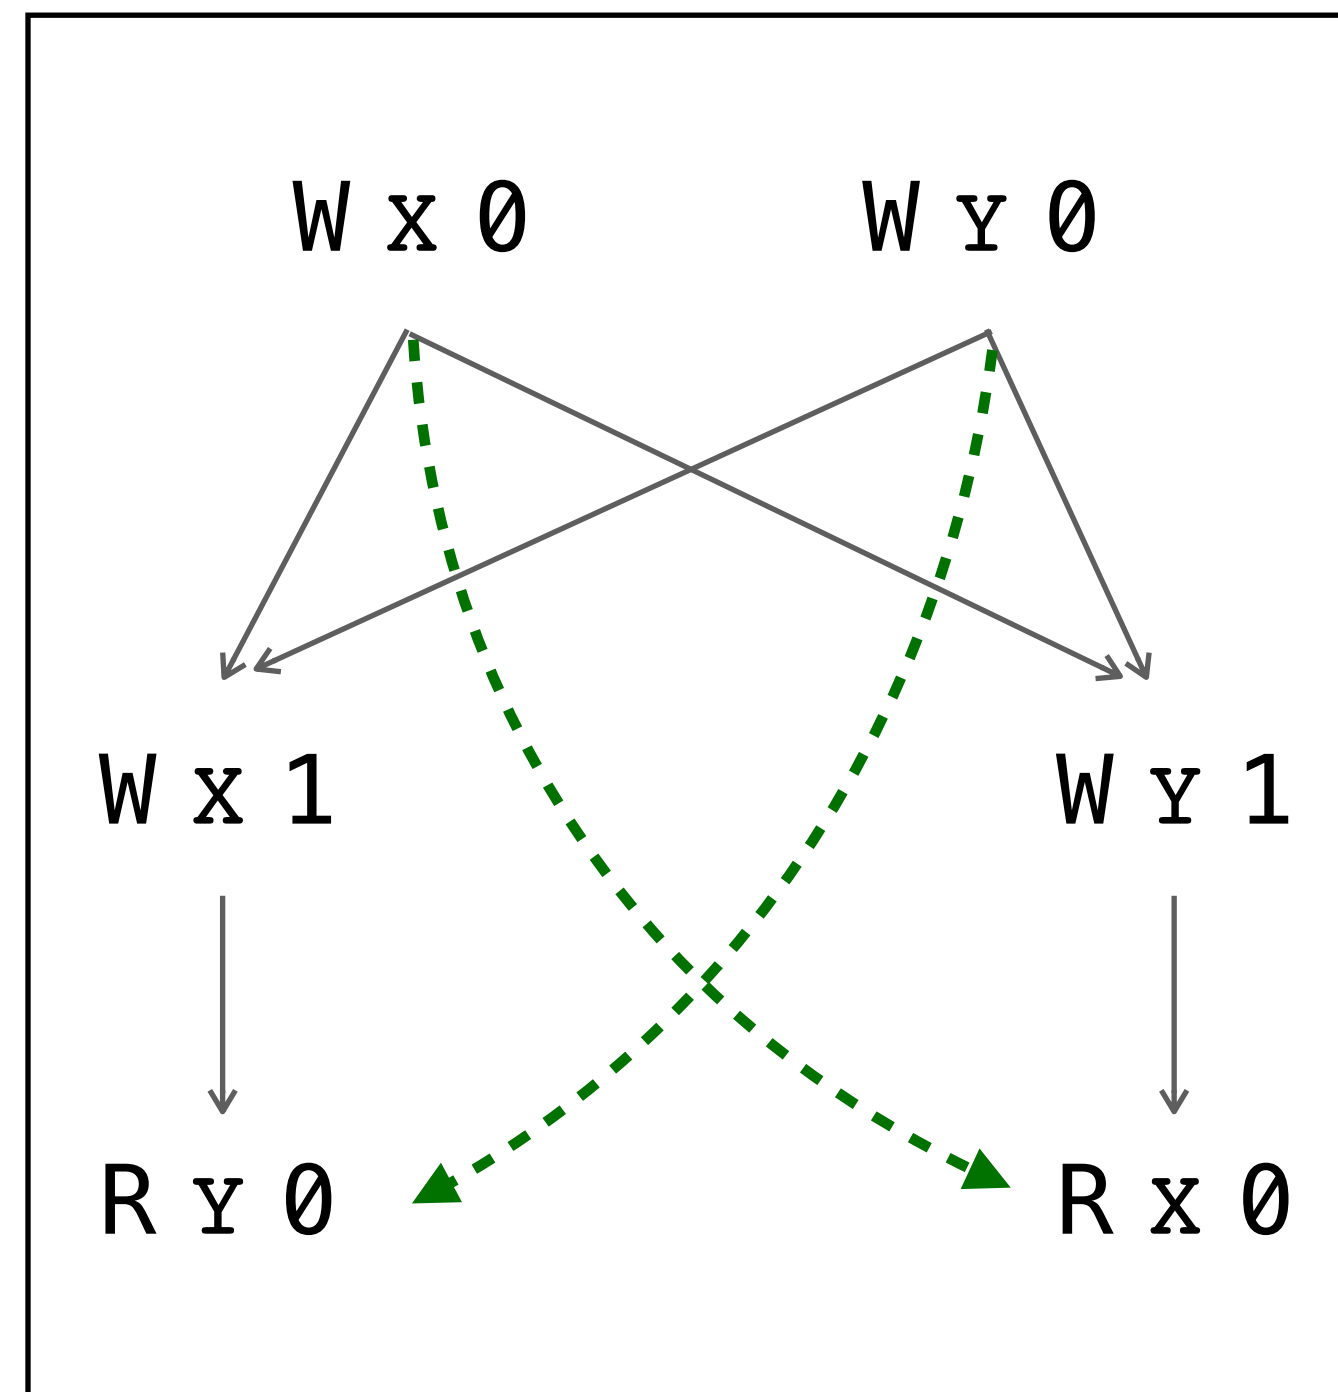
```

Y := 1;
repeat
  b := X;
until (b = 0);
Z := 1;
assert (Z = 1);

```



program-order  
**reads-from**



The execution graph is **consistent**.

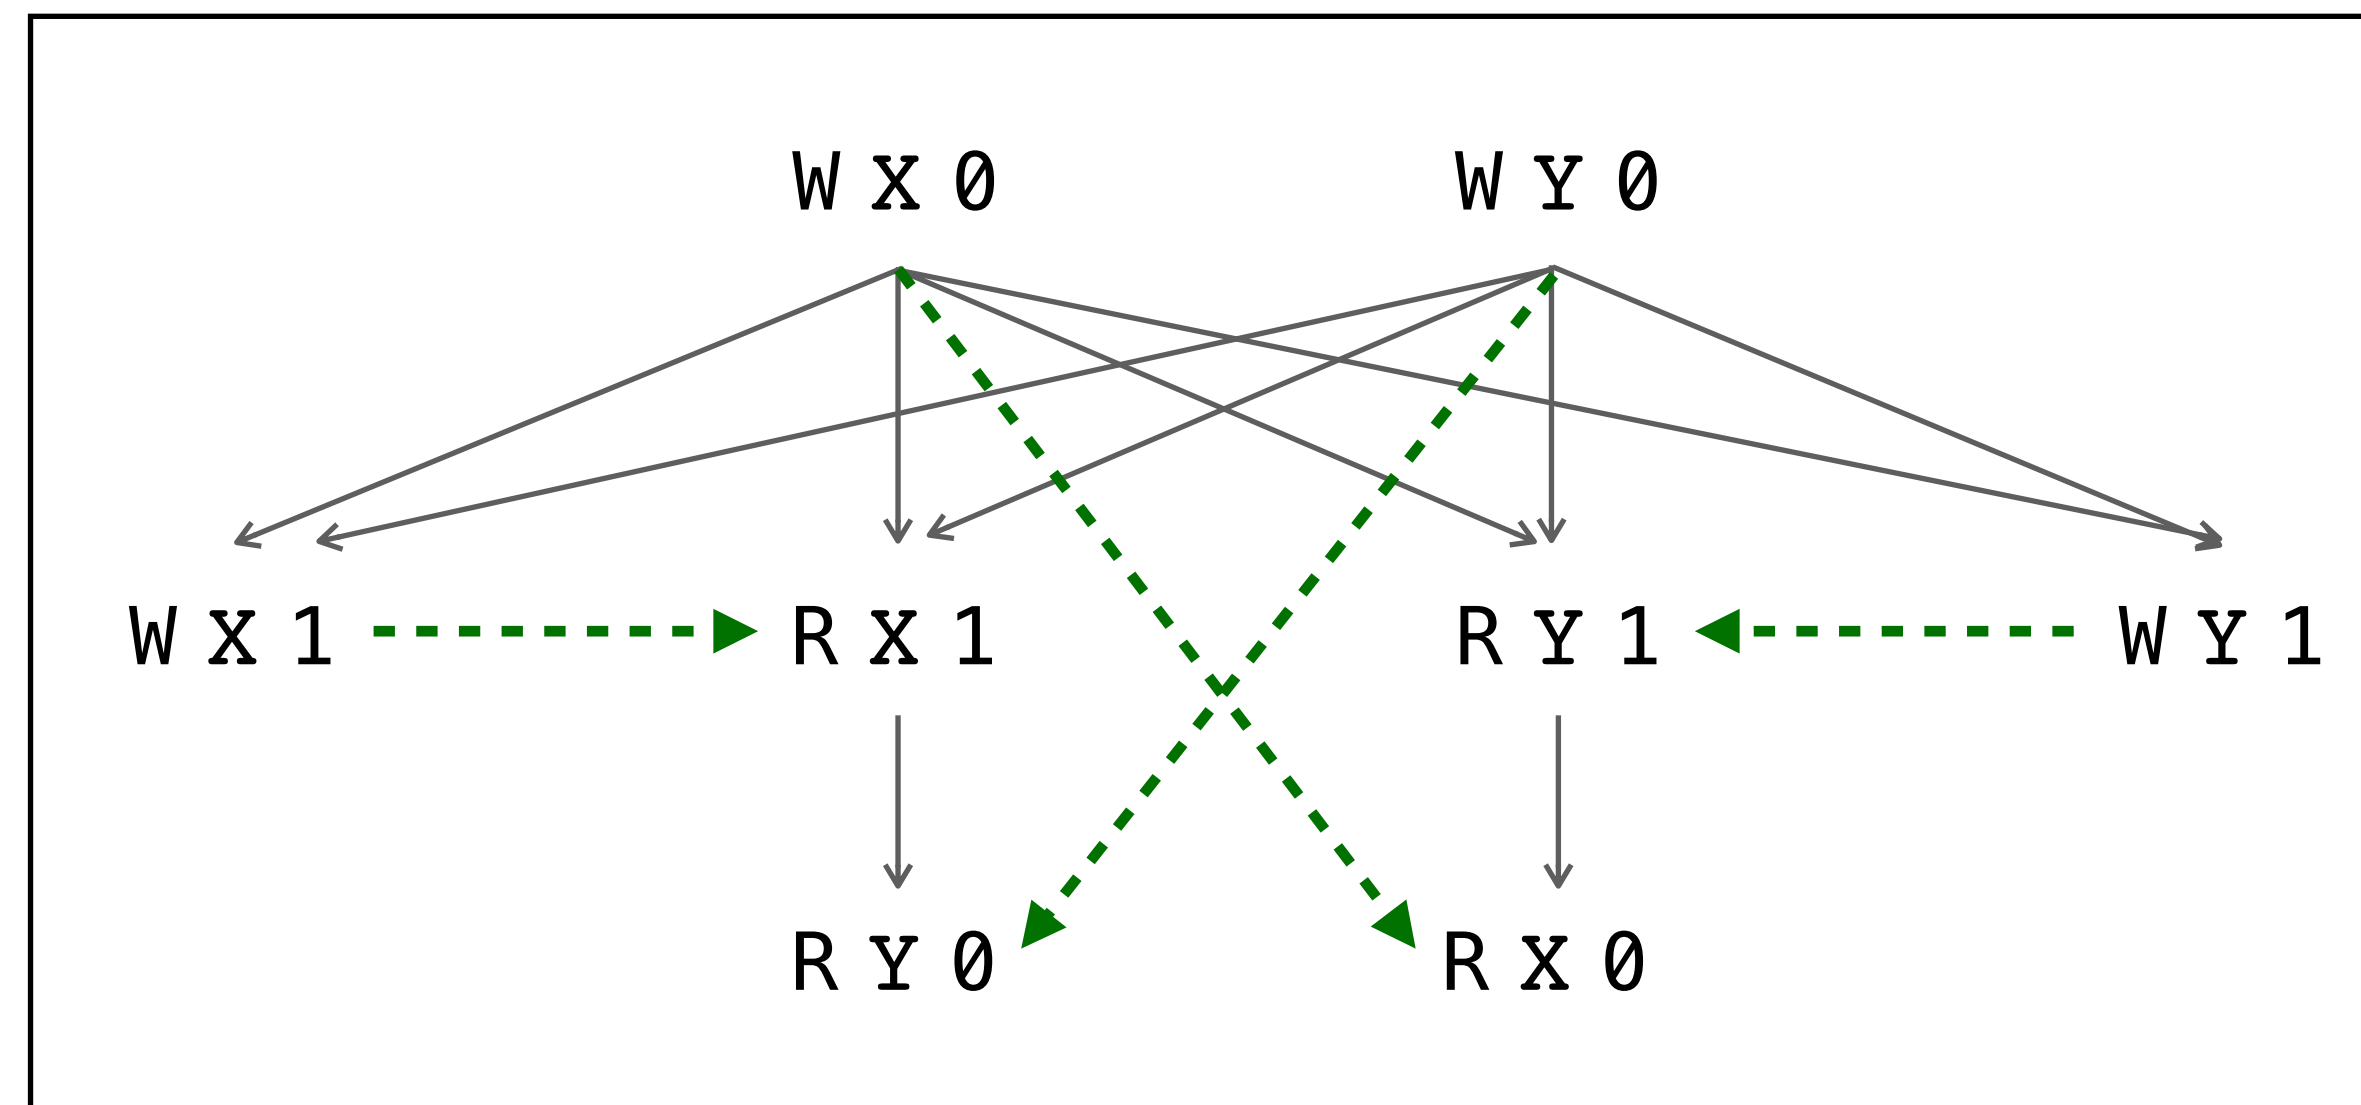
The annotated outcome is **allowed**.

# Non-Multi-Copy-Atomicity

Weaker than x86-TSO: different threads can observe writes in different orders

```
X = 1 ||| a = X // 1 ||| c = Y // 1 ||| Y = 1
      ||| b = Y // 0 ||| d = X // 0
```

→ program-order  
- - - reads-from



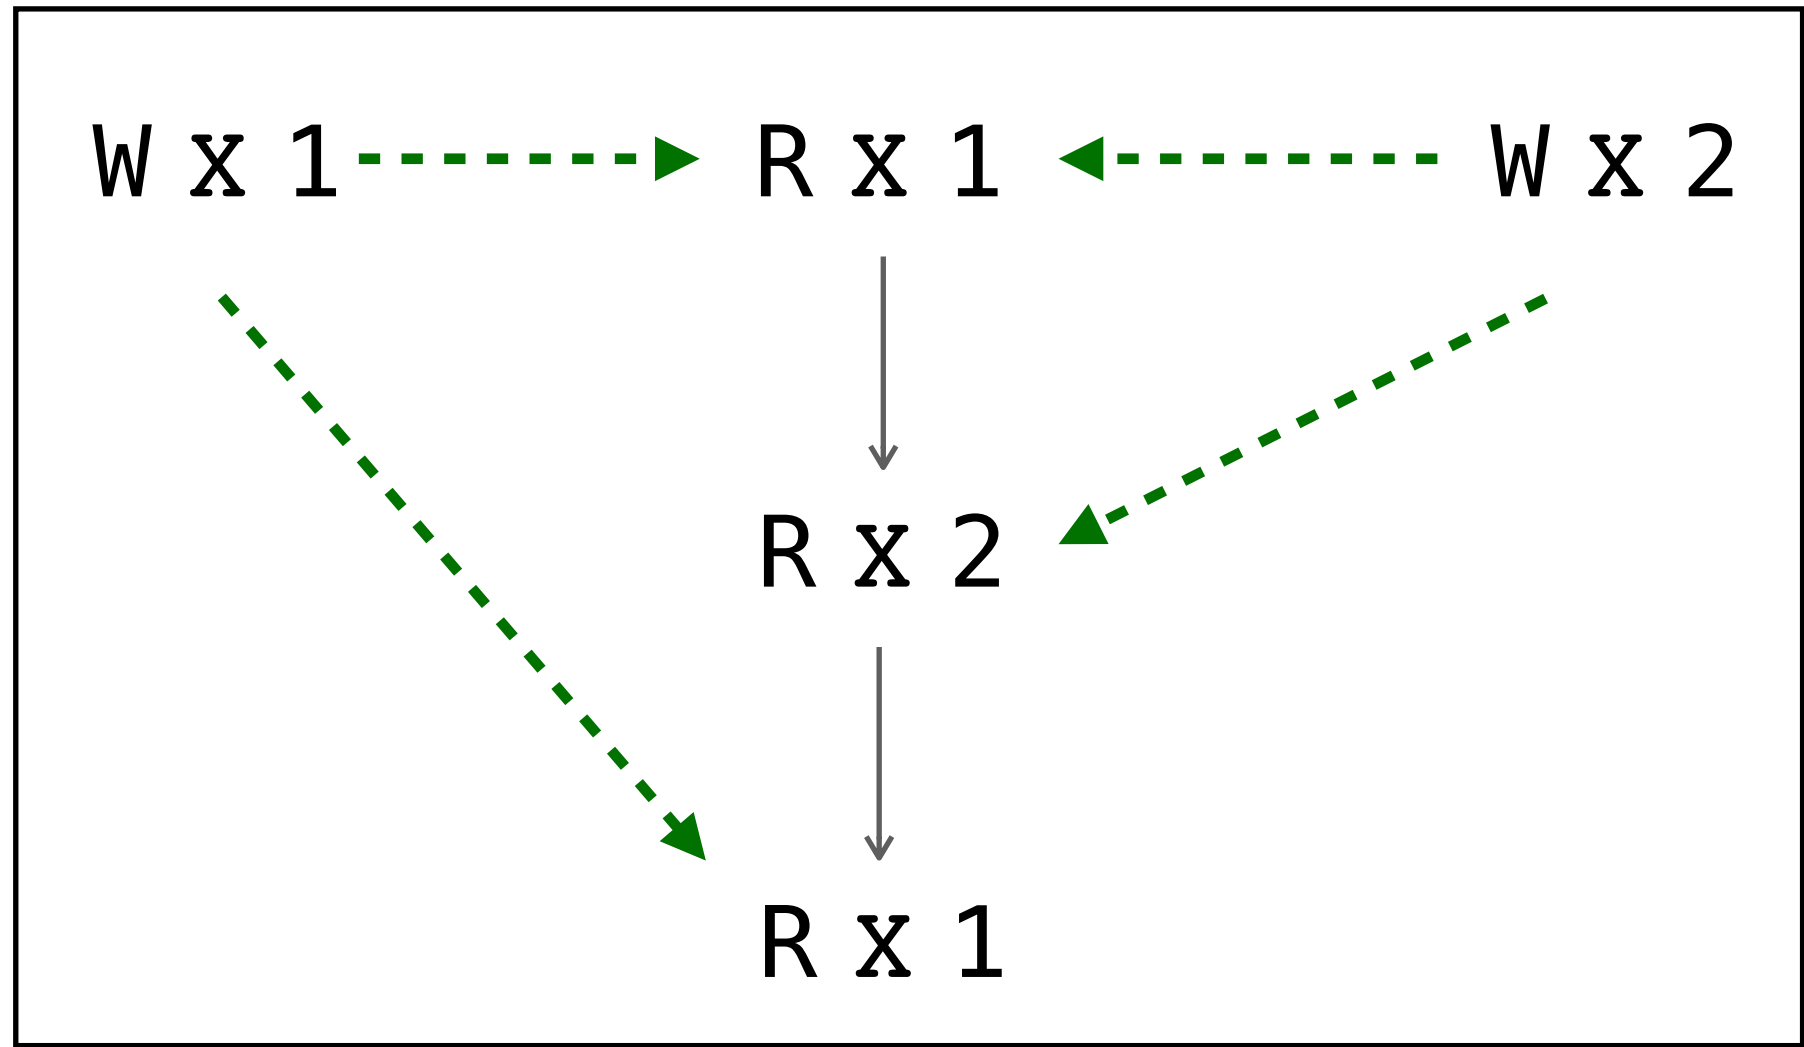
The execution graph is **consistent**.

The annotated outcome is **allowed**.

# What about *concurrent writes*?

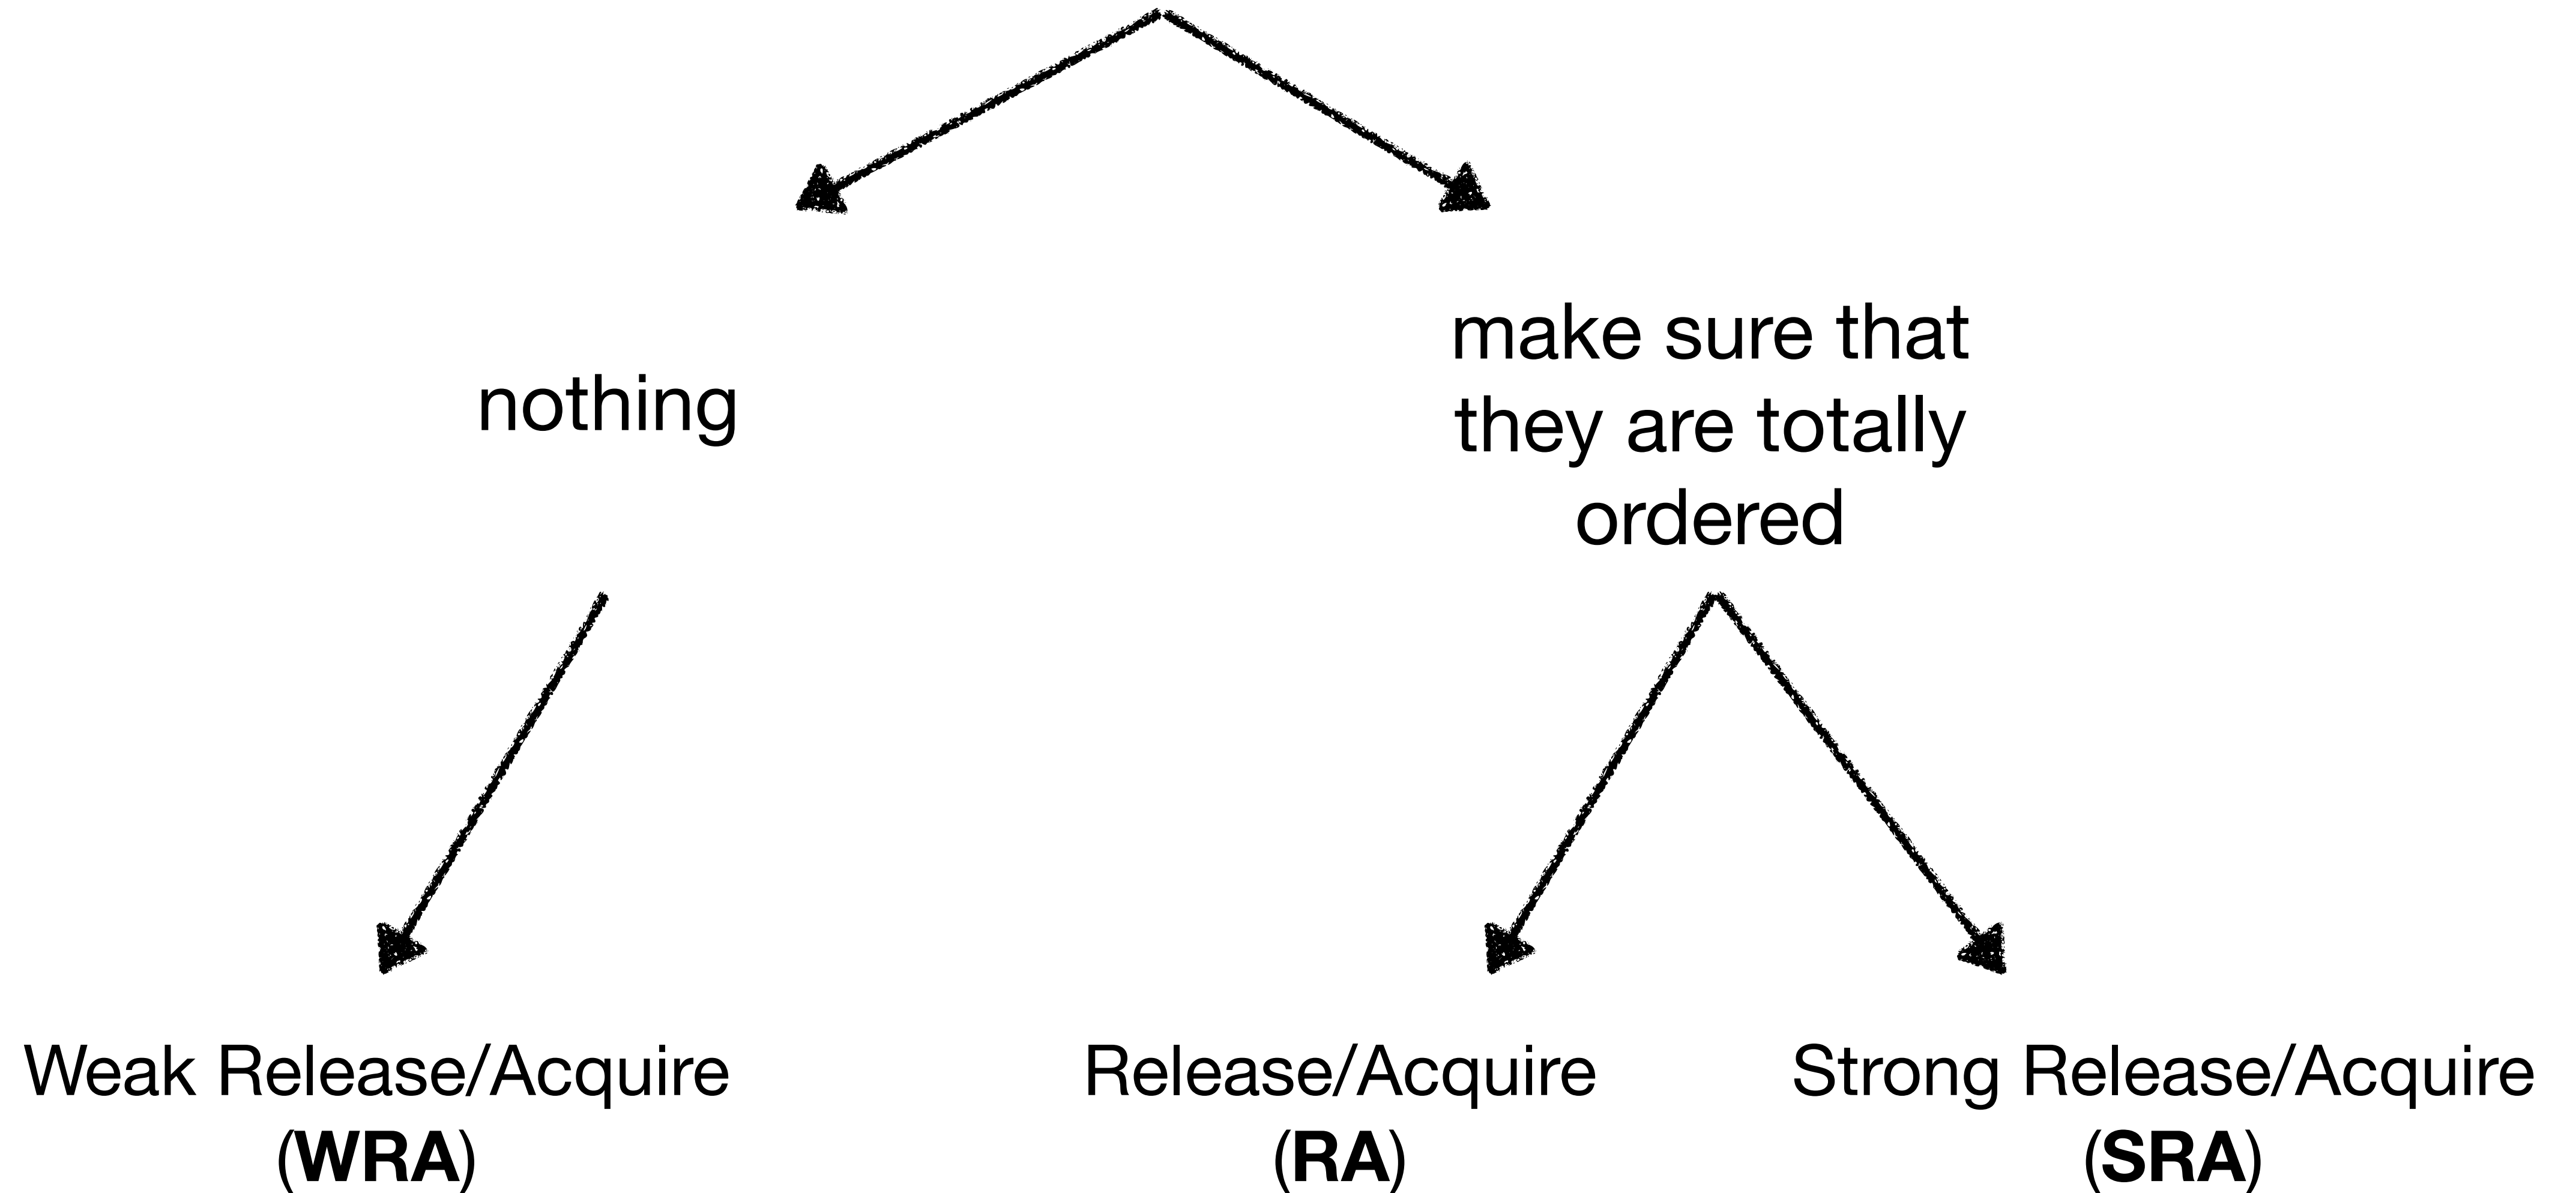
```
x = 1 ||| a = X // 1 ||| x = 2
        ||| b = X // 2 |||
        ||| c = X // 1 |||
```

→ program-order  
- - - reads-from



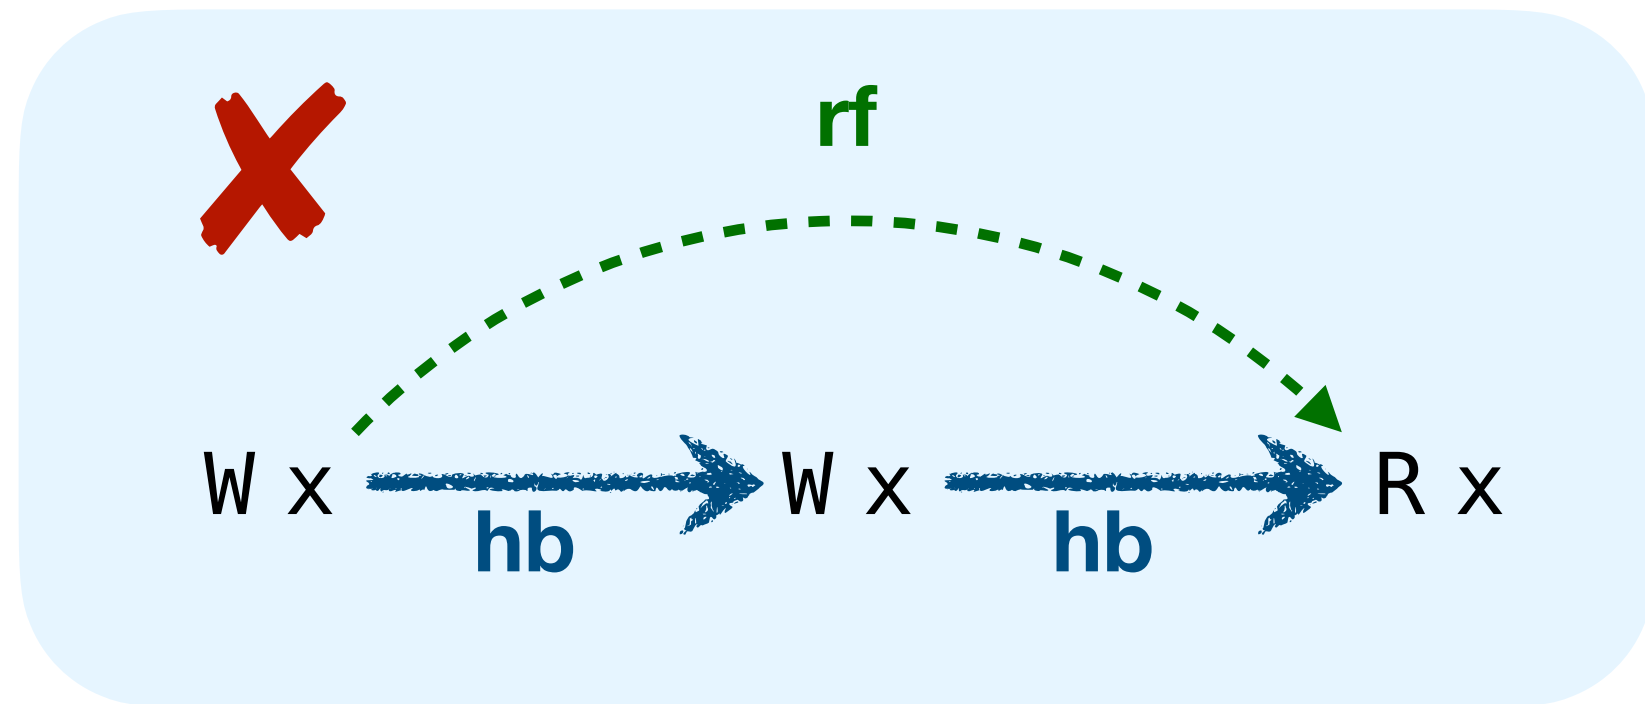
The execution graph is **consistent**.  
The annotated outcome is **allowed**.

# What about *concurrent writes*?



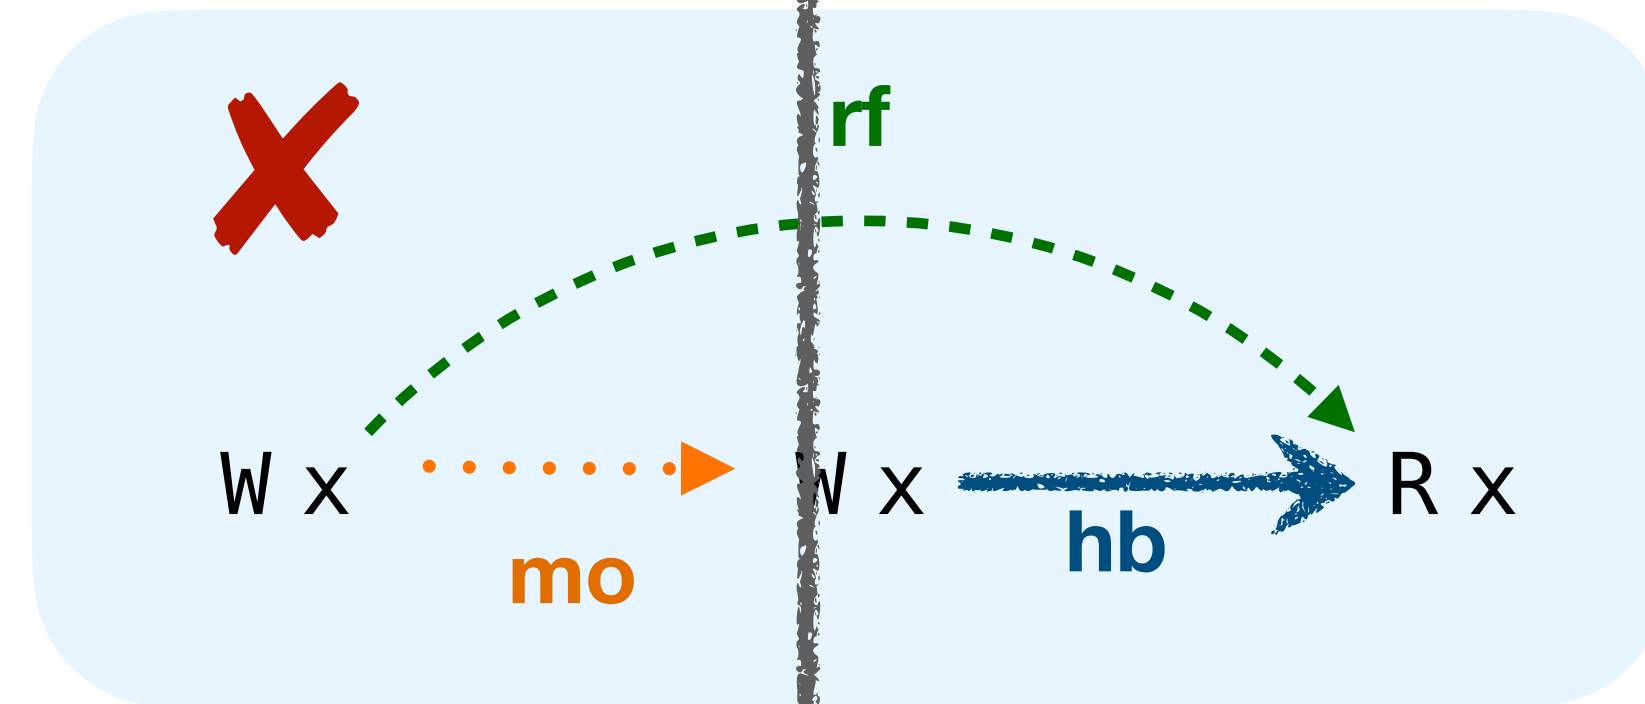
# Three Variants

Weak Release/Acquire  
(WRA)

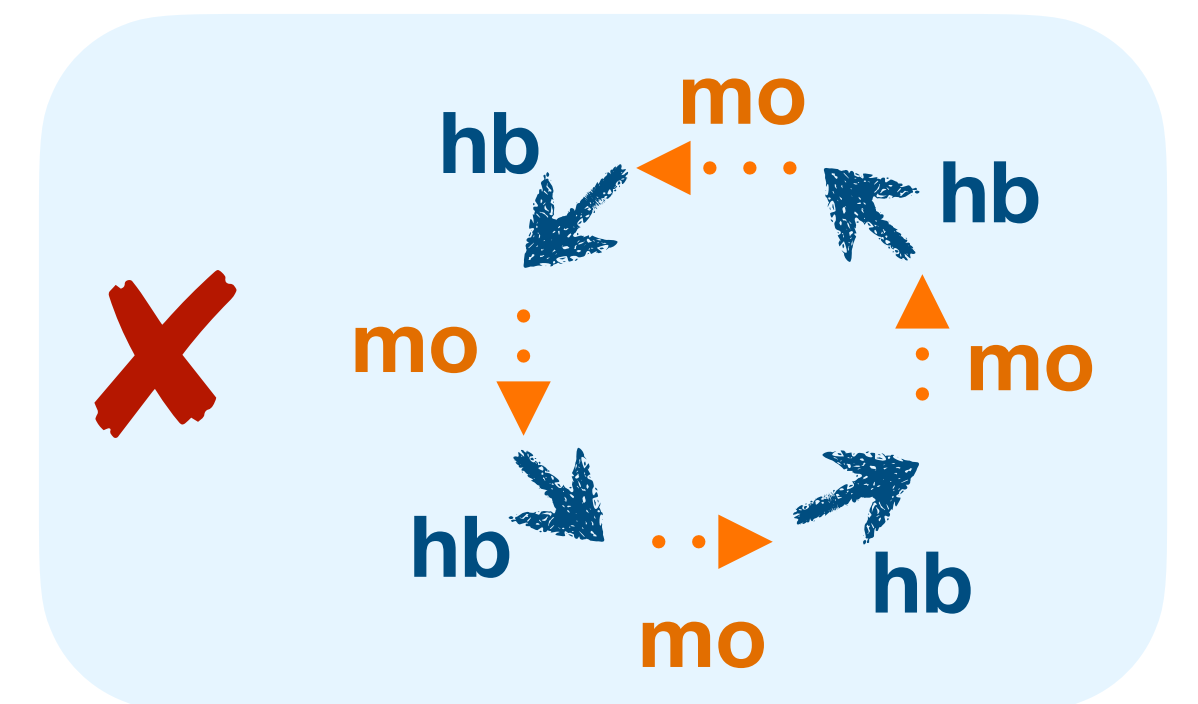
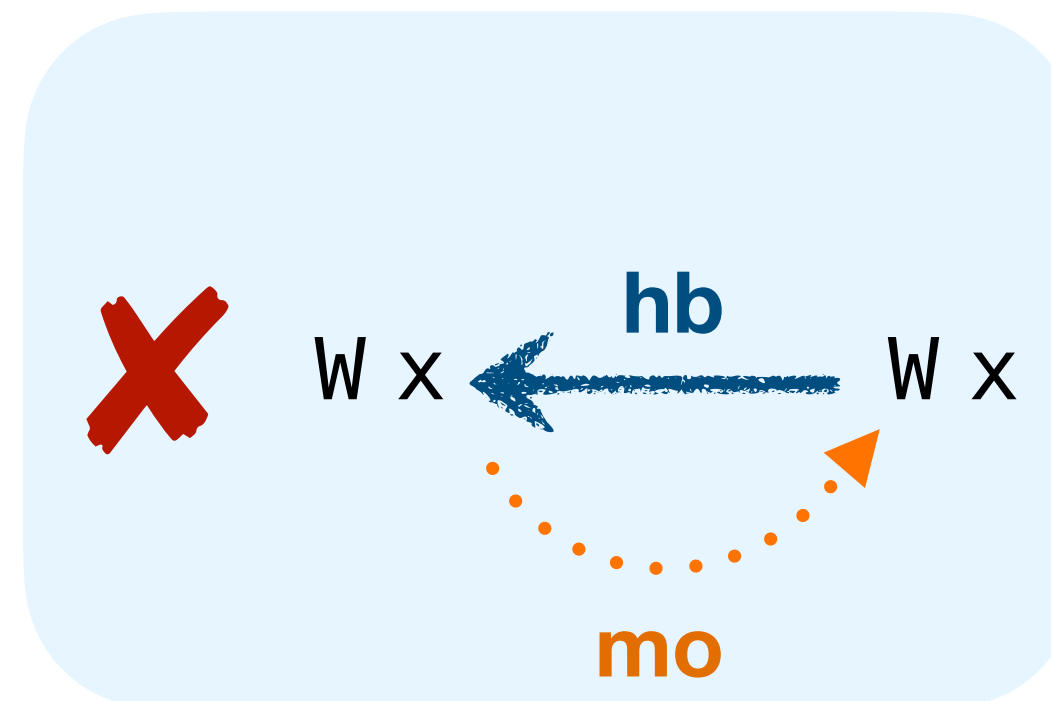


Release/Acquire  
(RA)

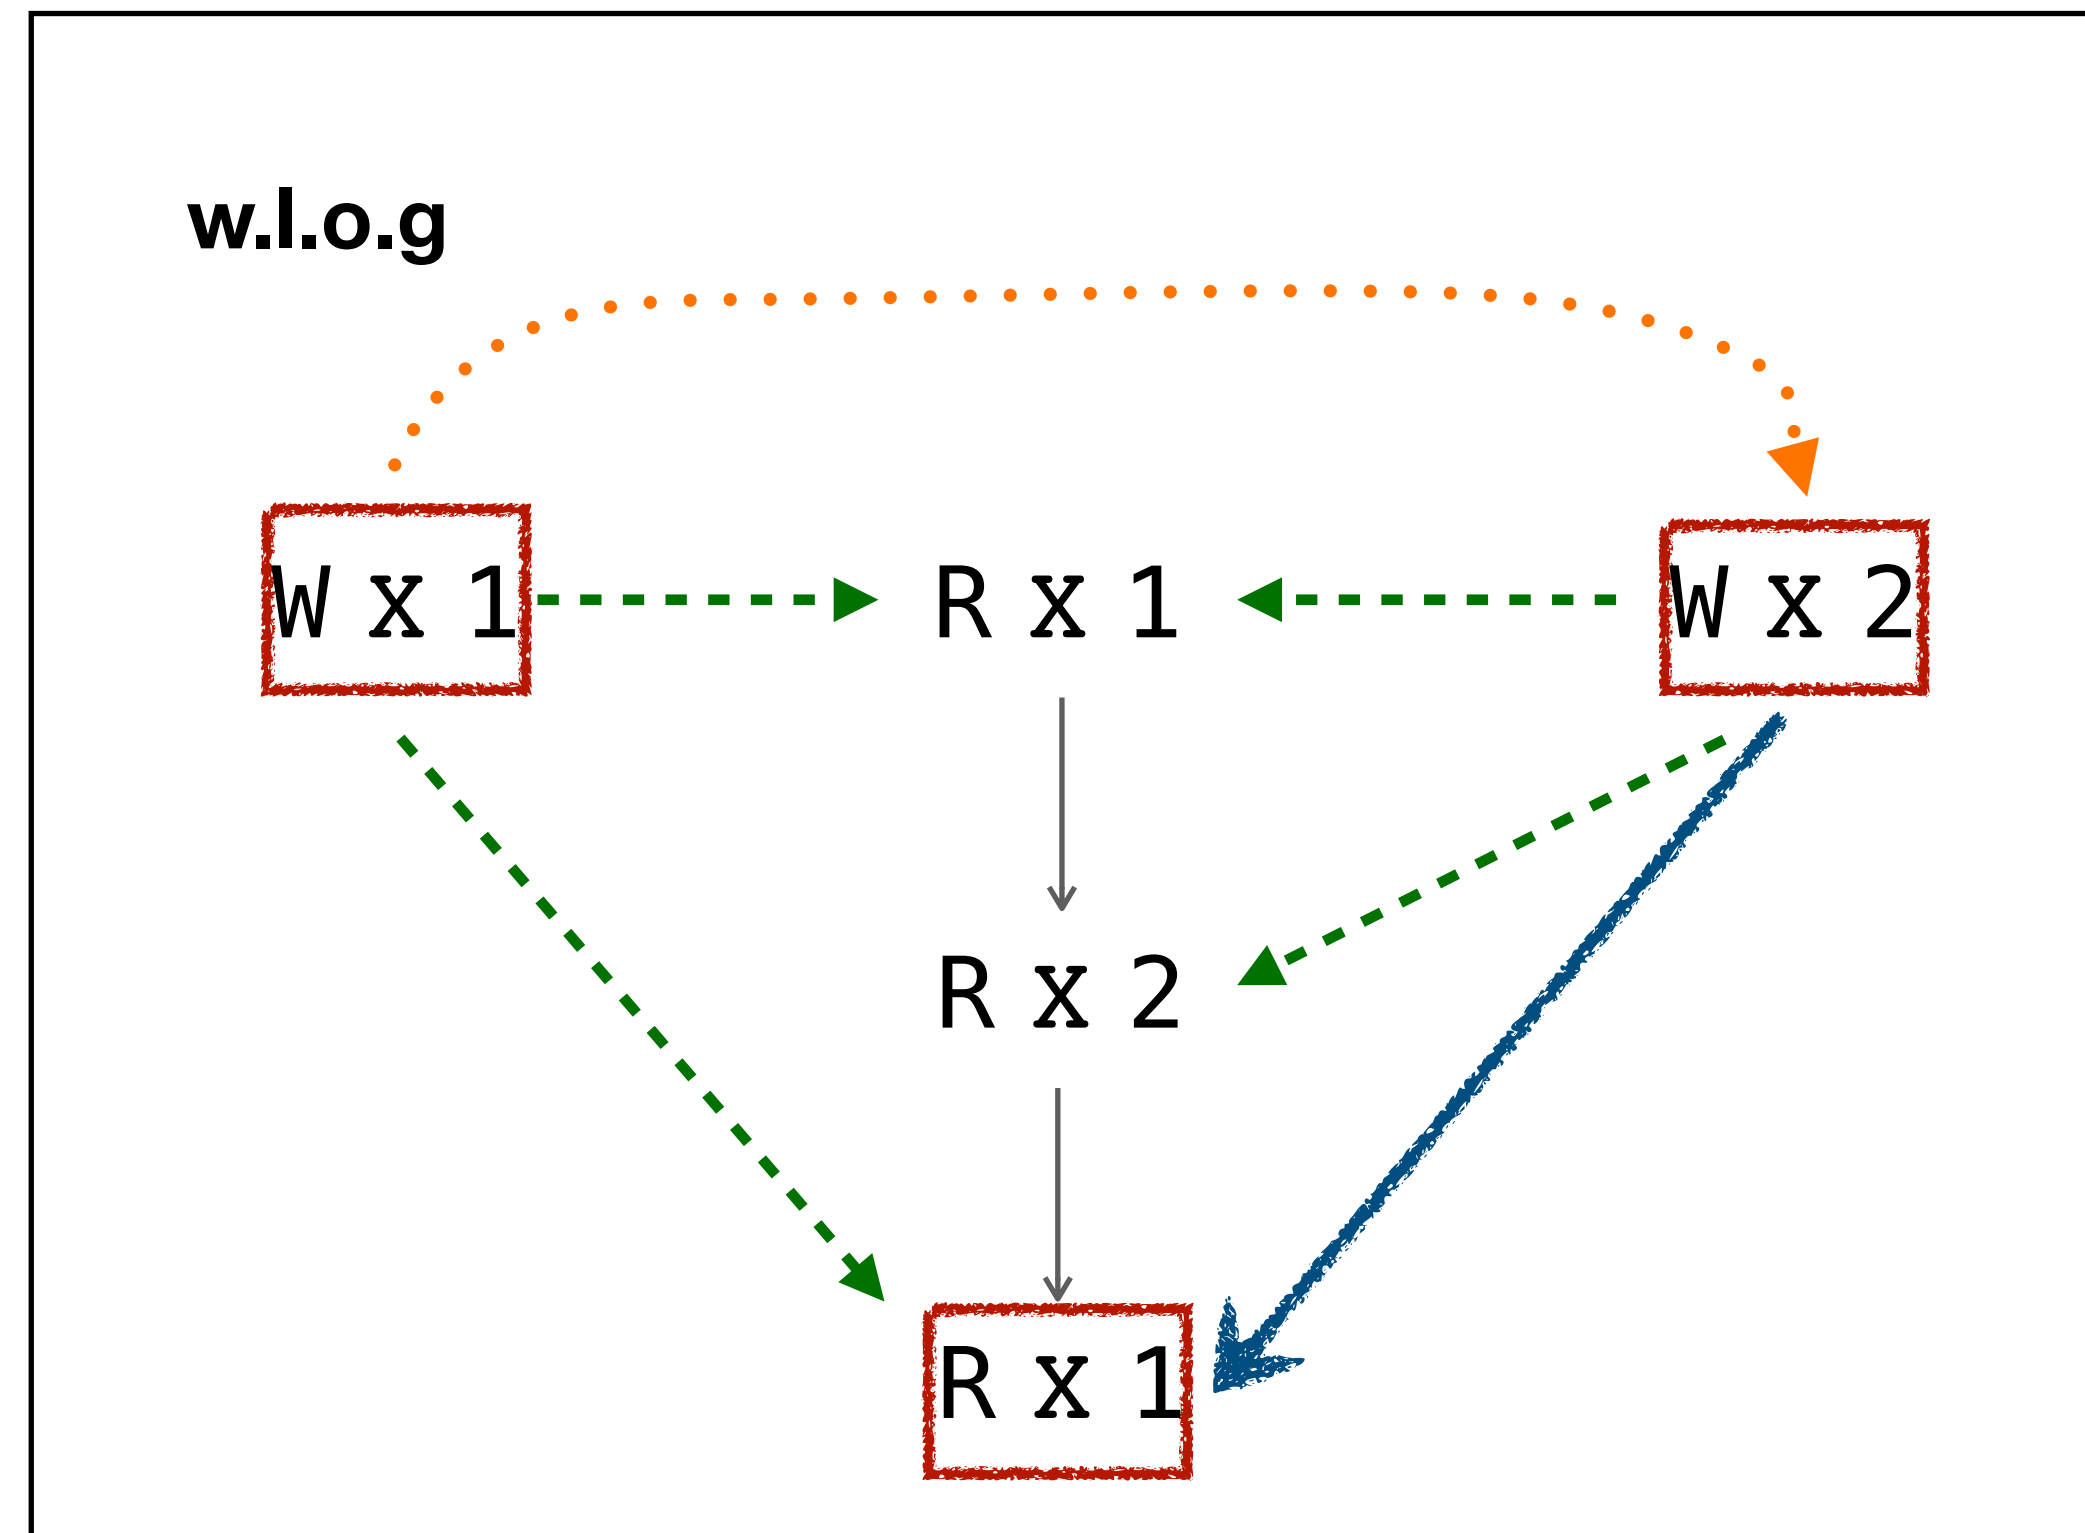
$\exists$  total order on writes to the same location (**modification-order**) s.t.:



Strong Release/Acquire  
(SRA)



# WRA is strictly weaker than RA



- program-order
- - - → reads-from
- ... → modification-order
- happens-before

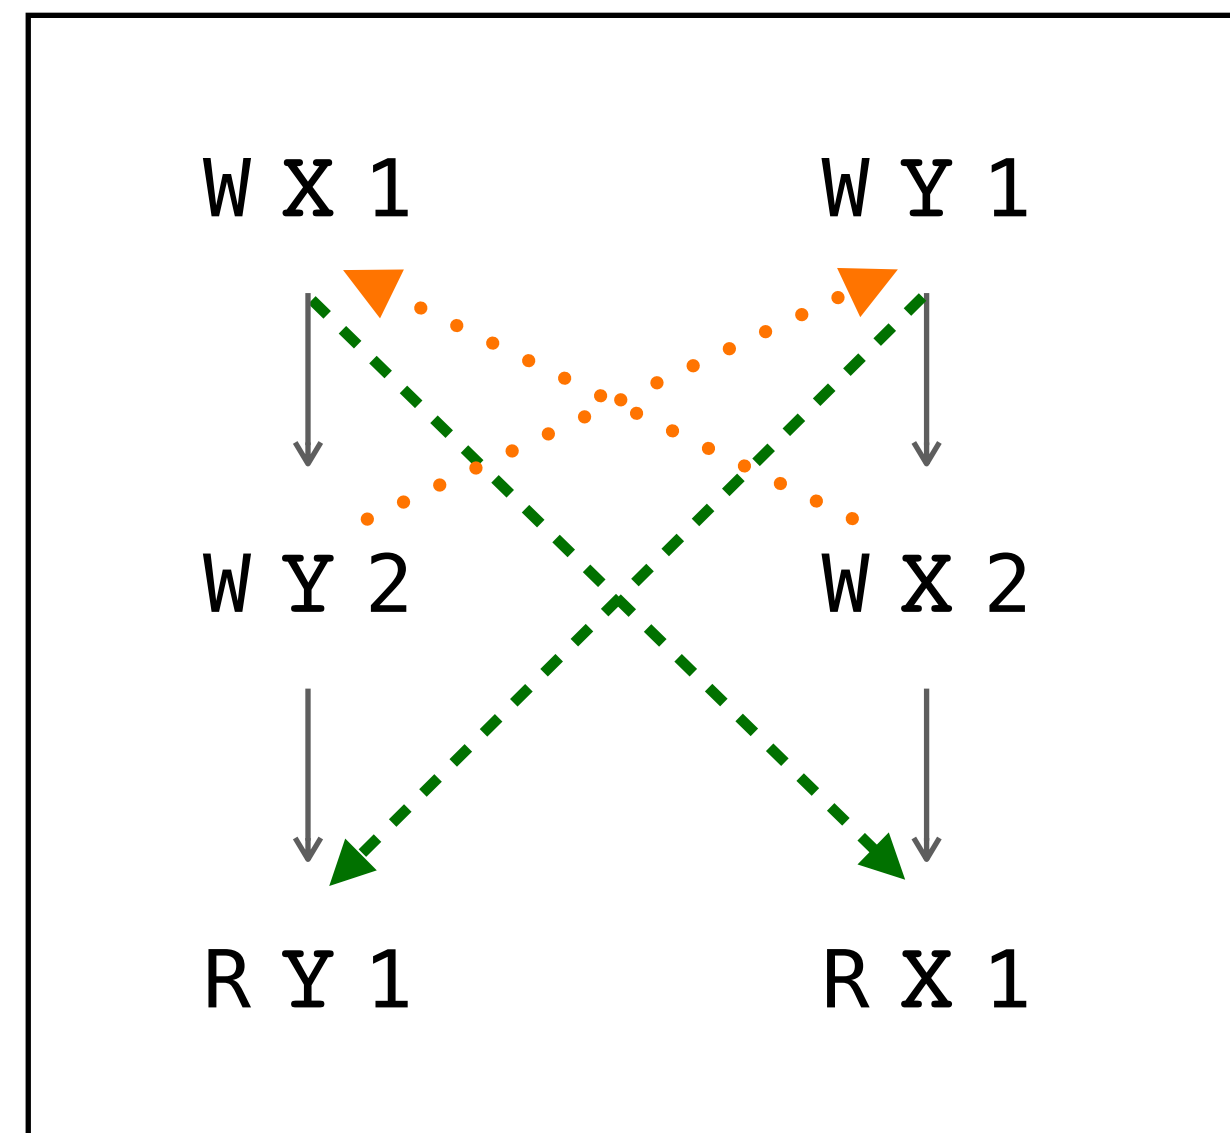
✓ WRA

✗ RA



# RA is strictly weaker than SRA

```
X = 1      || Y = 1
Y = 2      || X = 2
a = Y // 1 || b = X // 1
```



→ program-order  
- - - → reads-from  
... → modification-order

✓ RA

✗ SRA

# Causal Consistency

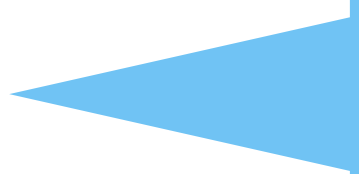
Sequential  
Consistency



x86-TSO



**Causal  
Consistency**



SRA



RA



WRA

**Distributed data-stores**

**POWER architecture [L, Giannarakis, Vafeiadis. POPL'16]**

**C/C++11**

**Java 9**

**CC in [Bouajjani, Enea, Guerraoui, Hamza. POPL'17]**

**[Kokologiannakis, L, Sagonas, Vafeiadis. POPL'18]**

# Write-Write-Race Freedom

## Theorem

*Prog* has no WW-races under **SRA**  $\implies \llbracket Prog \rrbracket_{WRA} = \llbracket Prog \rrbracket_{RA} = \llbracket Prog \rrbracket_{SRA}$

# Results

## Theorem

The verification problems under **SRA** and **WRA** are **decidable**.

- In contrast with **RA** [Abdulla, Arora, Atig, Krishna. PLDI'2019]
- To obtain this result we develop a **new semantics** for SRA and WRA.

## Corollary

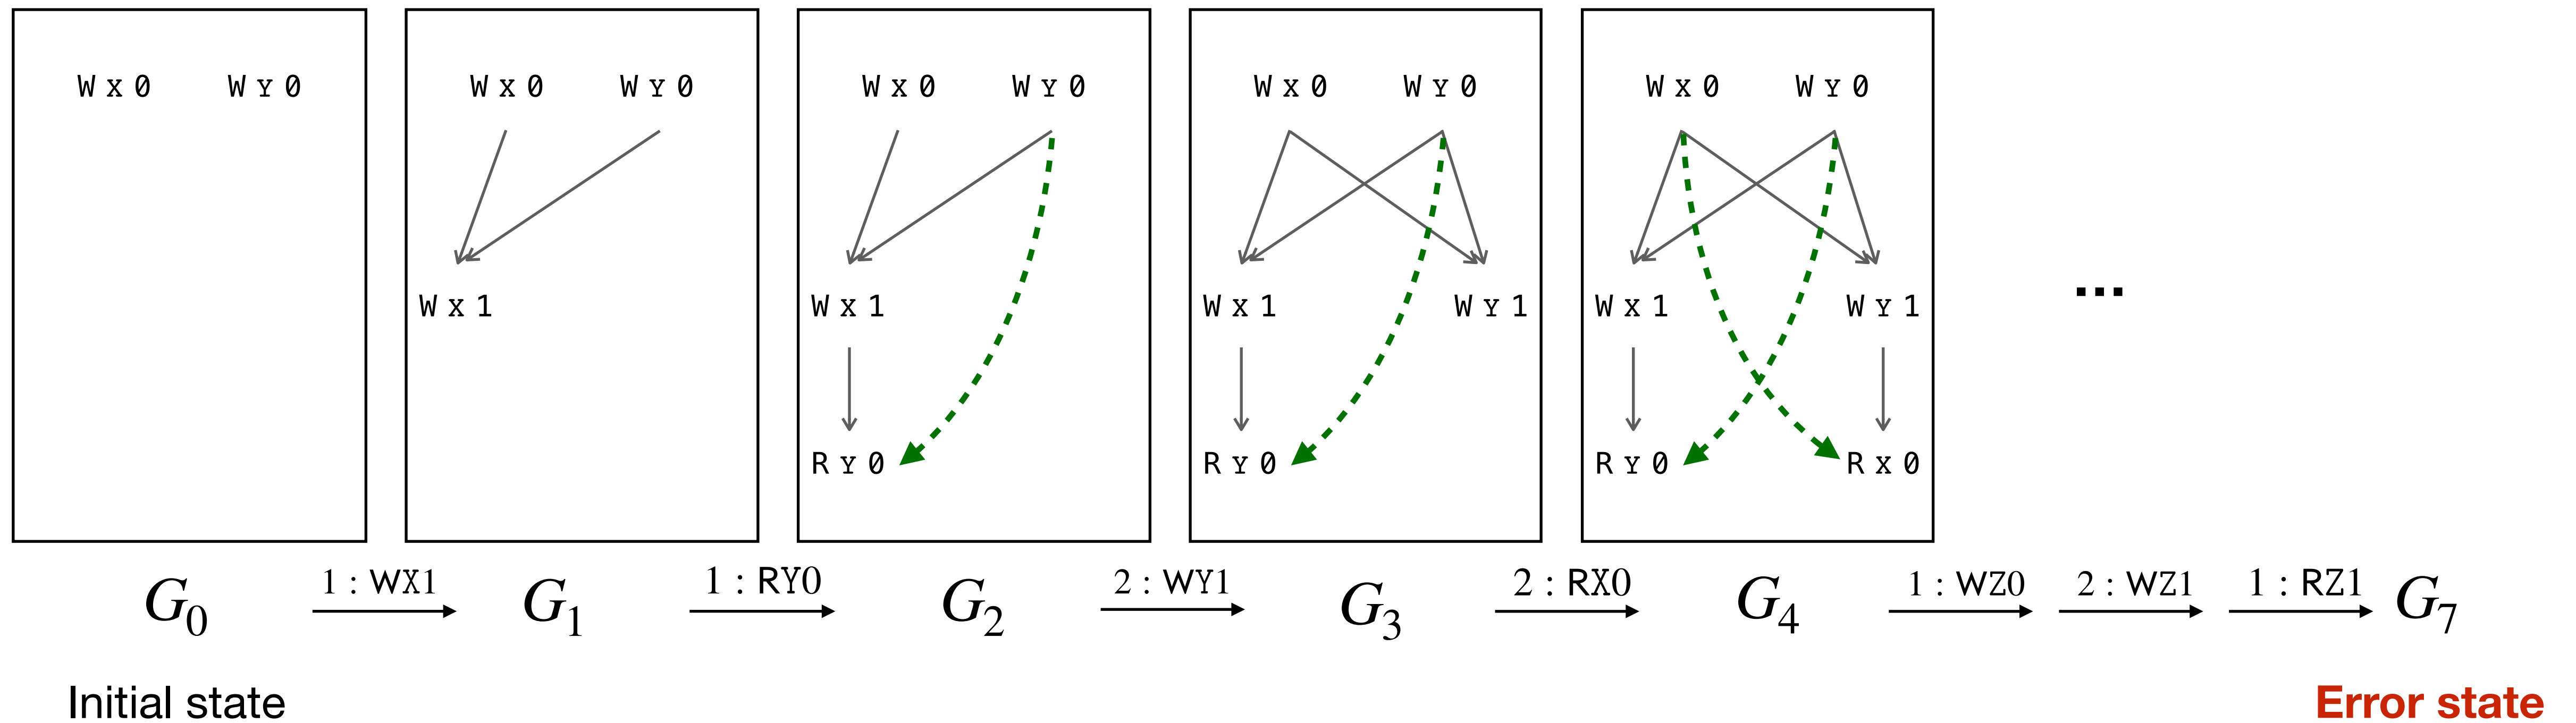
The verification problem under **RA** is **decidable** for write-write-race-free programs.

# Lower Complexity Bound

- For causal consistency, the problem is **non-primitive recursive**
- We can simulate a lossy FIFO channel machine
  - as for x86-TSO [Atig, Bouajjani, Burckhardt, Musuvathi. POPL'2010]



Even when the program is finite state,  
its synchronization with a causally consistent memory is **infinite** state.



- To establish **decidability**:
- We use the framework of *well-structured transition systems* (WSTS)  
[Abdullah] [Finkel, Schnoebelen] ...
- **Challenge**: find a WSTS equivalent to a casually consistent memory

# Backward Reachability

**Input:** LTS  $(Q, Q_0, \rightarrow)$ , state  $q_{\text{bad}}$

**Output:** is  $q_{\text{bad}}$  reachable?

$S := \{q_{\text{bad}}\}$

repeat

$S_{\text{prev}} := S$

$S := S \cup \text{pre}(S)$

until  $(S = S_{\text{prev}})$

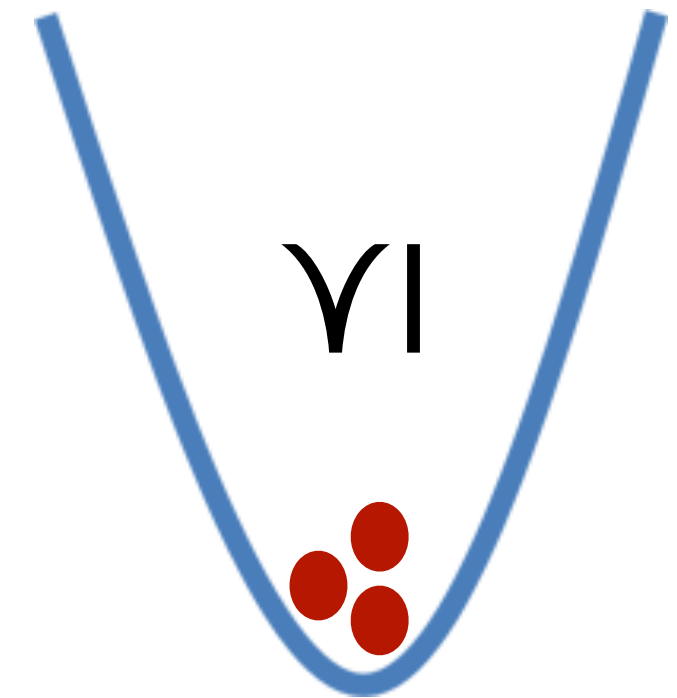
return  $(Q_0 \cap S \neq \emptyset)$

To make this an algorithm we  
need to work with  
*finitely representable sets* &  
*guarantee termination*



# Well-Structured Transition Systems

- Equip the transition system with a **well quasi-order**  $\preceq$
- Work with **upward closed** sets of states represented by their finite basis



$\preceq$  should be compatible with  $\rightarrow$

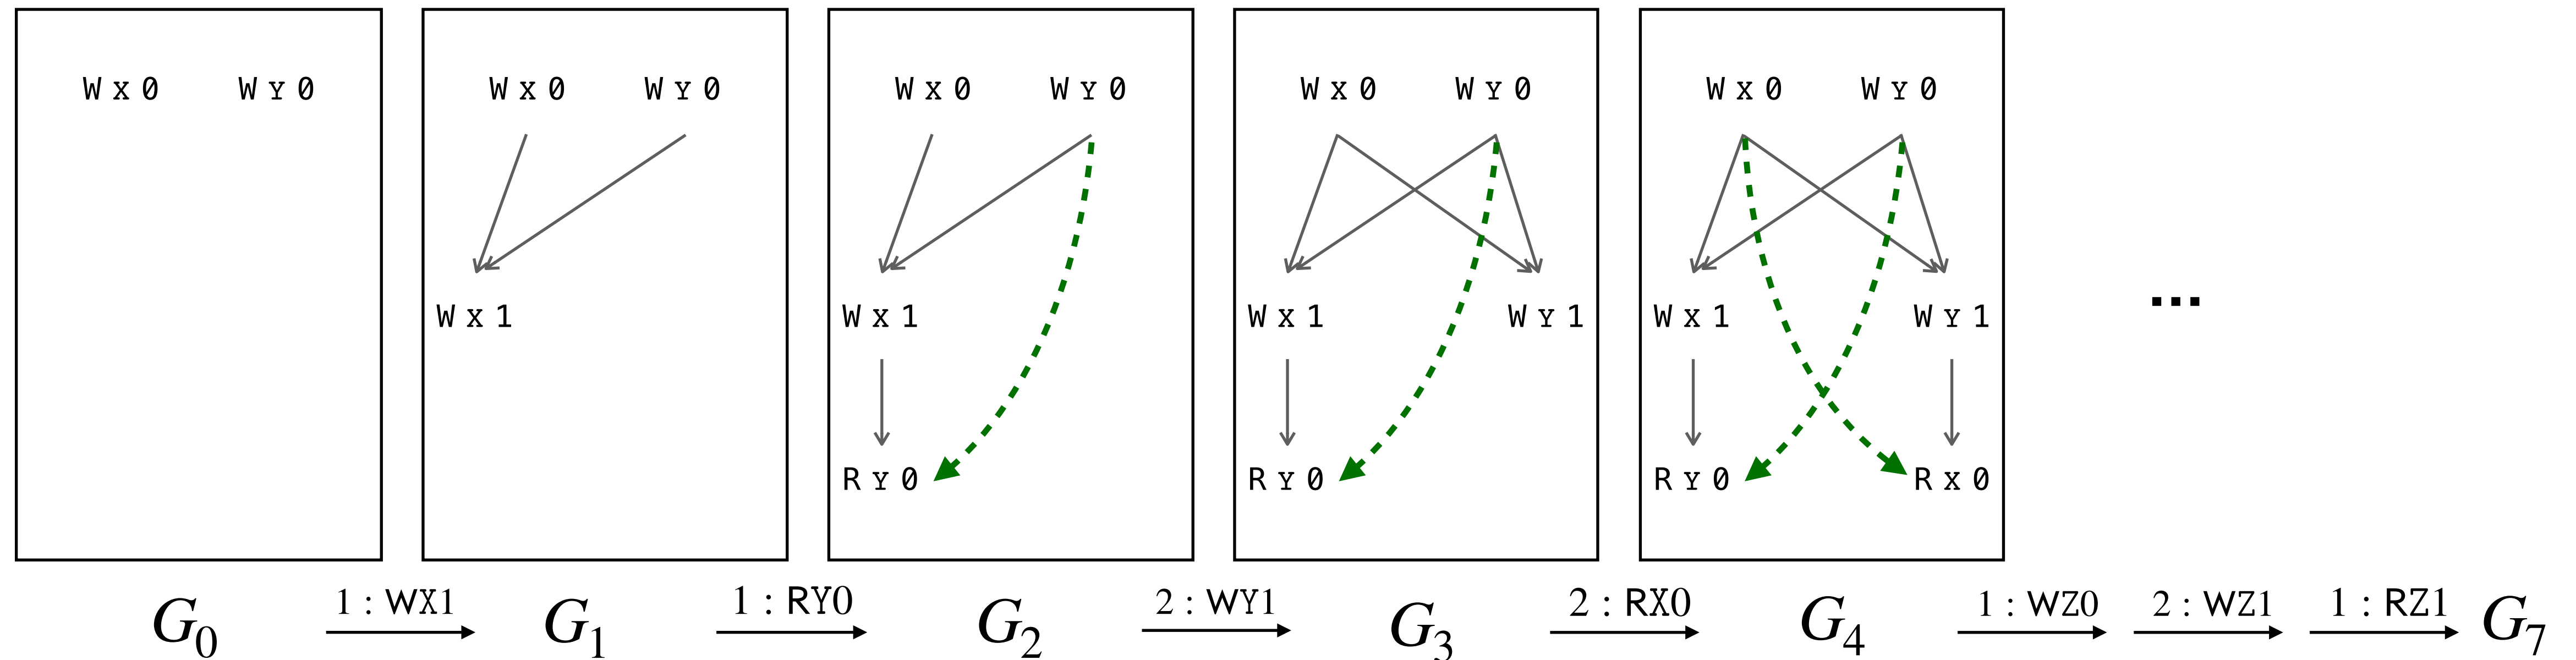
$$\begin{array}{ccc} q'_1 & \xrightarrow{*} & q'_2 \\ \forall I & & \forall I \\ q_1 & \longrightarrow & q_2 \end{array}$$

Compatibility is guaranteed in “lossy” systems

$$\text{(lose)} \frac{q' \succeq q}{q' \rightarrow q}$$

- **Challenge:** characterize casually consistent shared memory as a **lossy** system

# Causal Consistency as a WSTS



Two critical obstacles:

- Partial order embedding is **not** a well-quasi-order
- Execution histories are **not** lossy

# Key Idea

 Share

Why do I keep focusing on the past instead of the future?

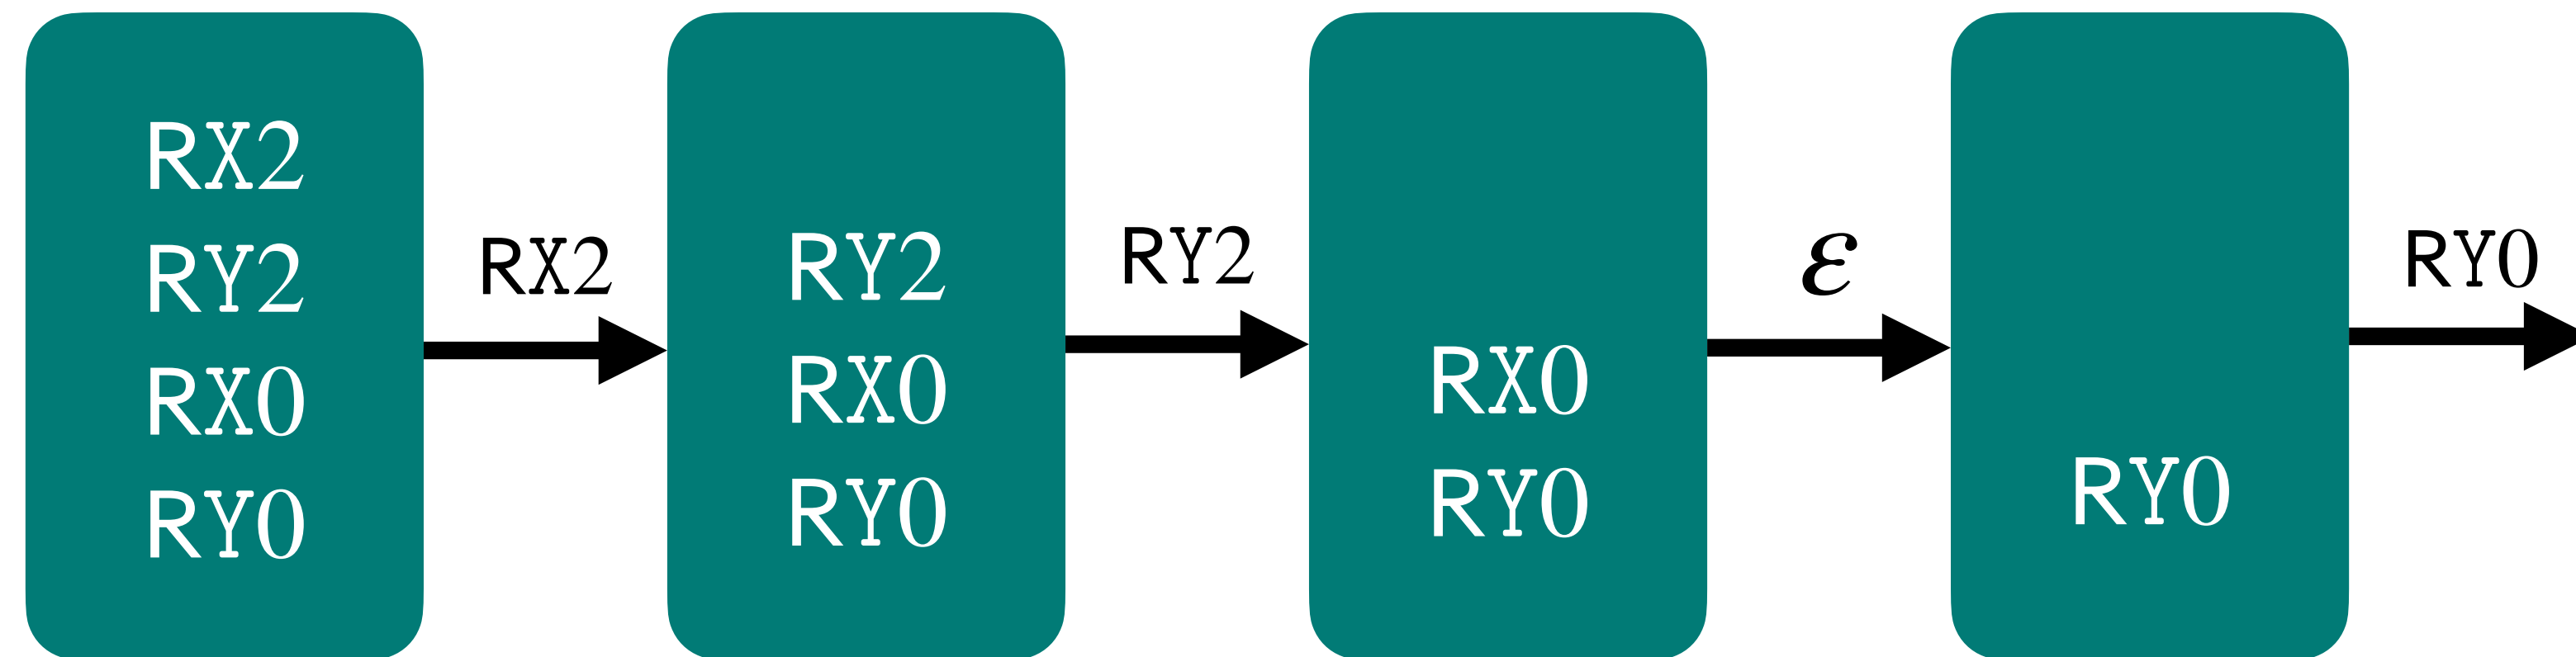
 16 Answers

 Last Updated: 05/01/2018 at 12:47pm

Record the **threads' potentials** in memory states:

***what possible sequences of reads each thread can execute?***

# Lossy SRA



$Q = \text{Threads} \rightarrow \{Rxv \mid x \in \text{Var}, v \in \text{Val}\}^*$

$Q_0 = \text{Threads} \rightarrow \{Rx0 \mid x \in \text{Var}\}^*$

$q \preceq q' \iff \forall \tau \in \text{Threads} . q(\tau) \sqsubseteq q'(\tau)$

↑  
subsequence

```

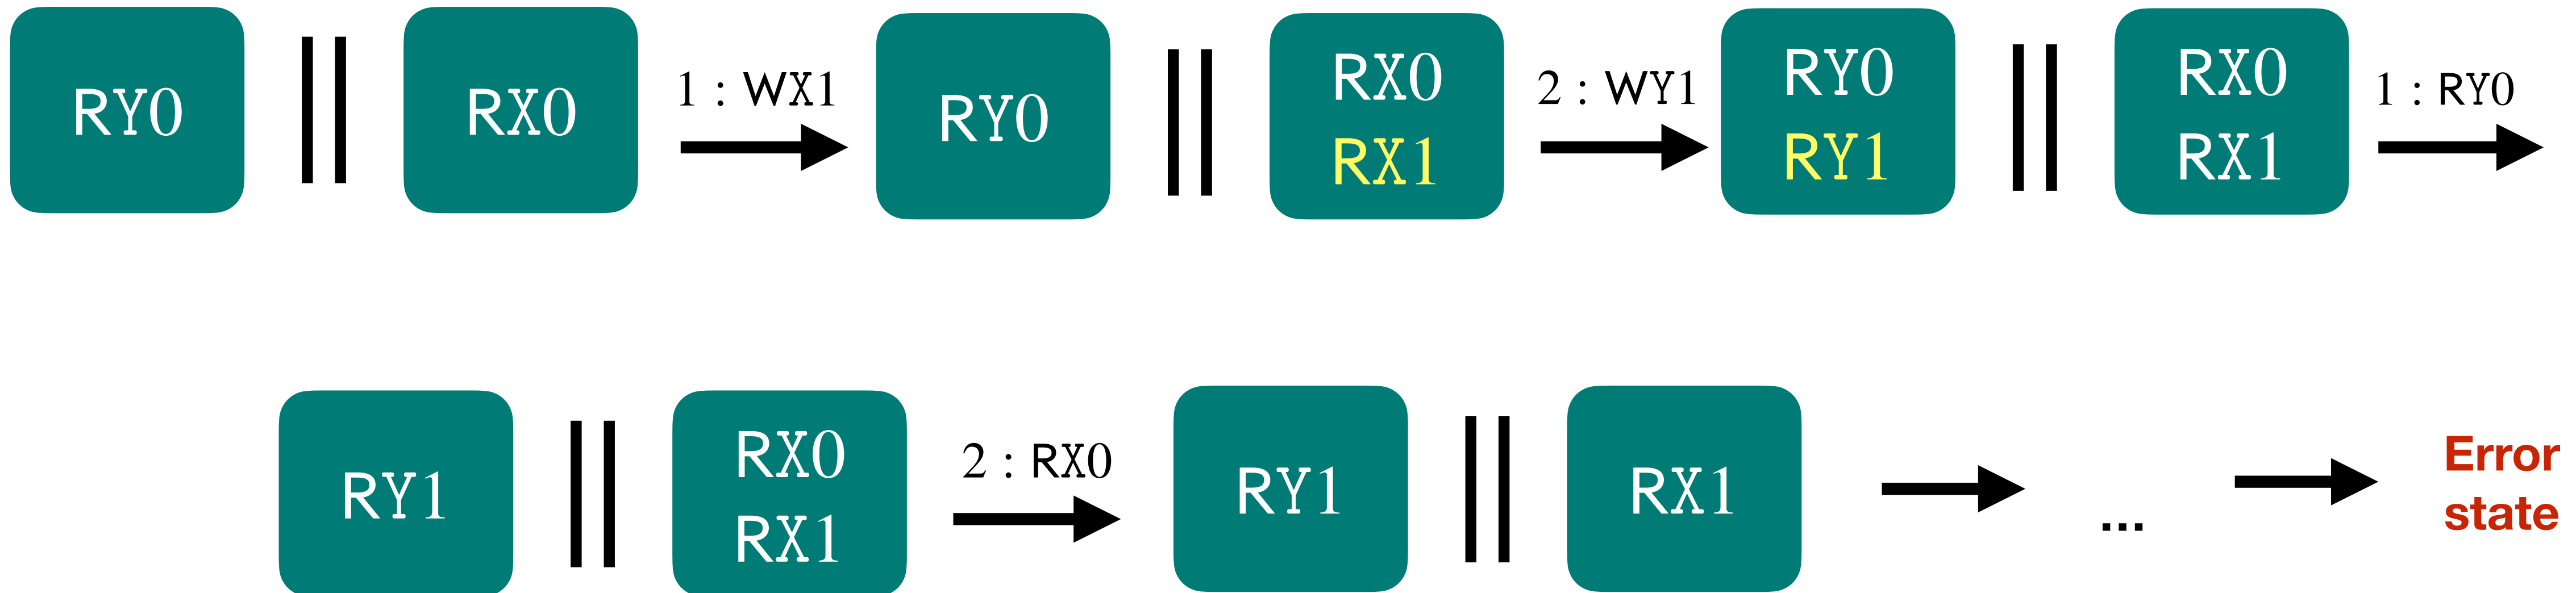
X := 1;
repeat
  a := Y;
until (a = 0);
Z := 0;
assert (Z = 0);

```

```

Y := 1;
repeat
  b := X;
until (b = 0);
Z := 1;
assert (Z = 1);

```



# Potential Maintenance for SRA

$\varepsilon$   
→

## Lose step

Remove some elements from the potentials

$\tau : R_{xv}$

→

## Read steps

Precondition: first element in  $\tau$ 's potential is  $R_{xv}$

**Deterministic**

$\tau : W_{xv}$

→

## Write steps

Precondition: no  $R_{x\_}$  in  $\tau$ 's potential

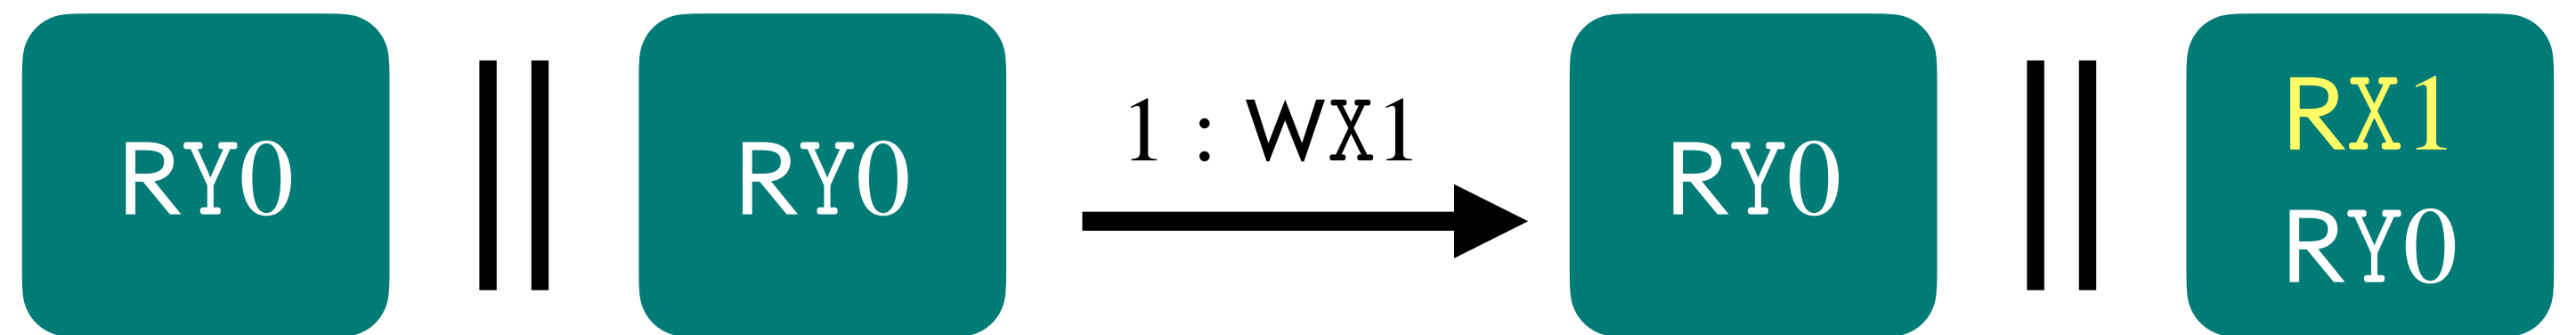
All threads may get new options  $R_{xv}$

**Non-Deterministic**

→ **where?**

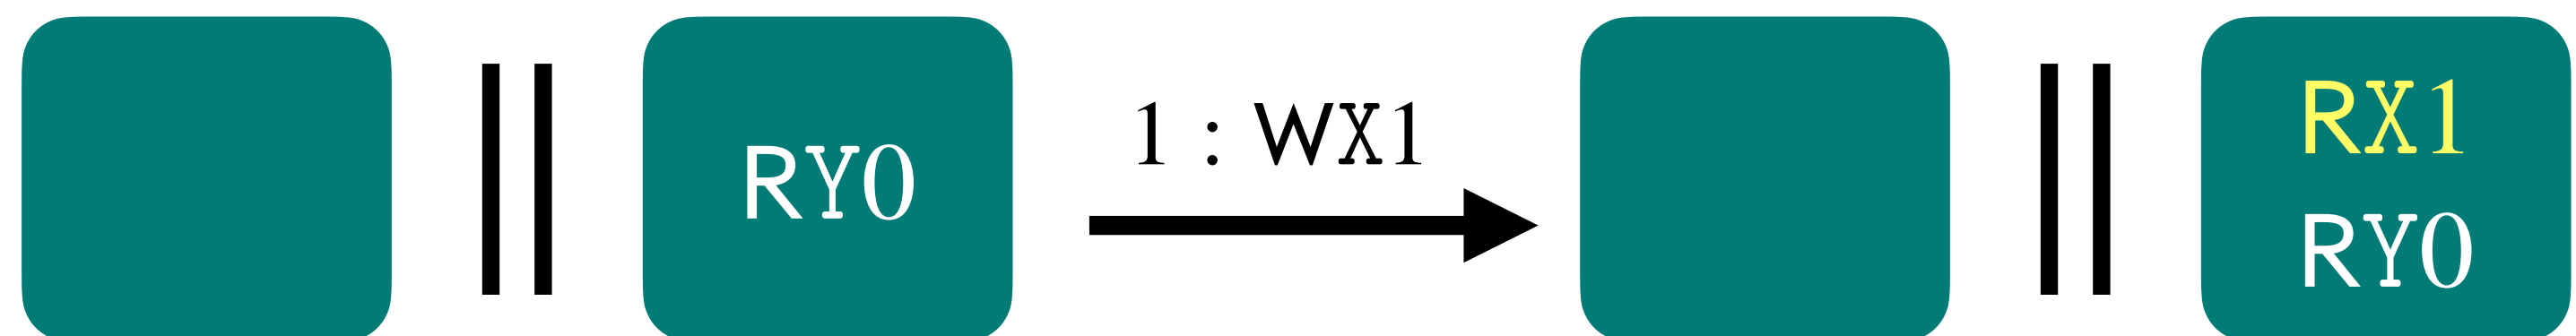
# Where?

```
X = Y = 0  
||  
a = X // 1  
b = Y // 0  
X = 1
```



This transition **should be allowed**

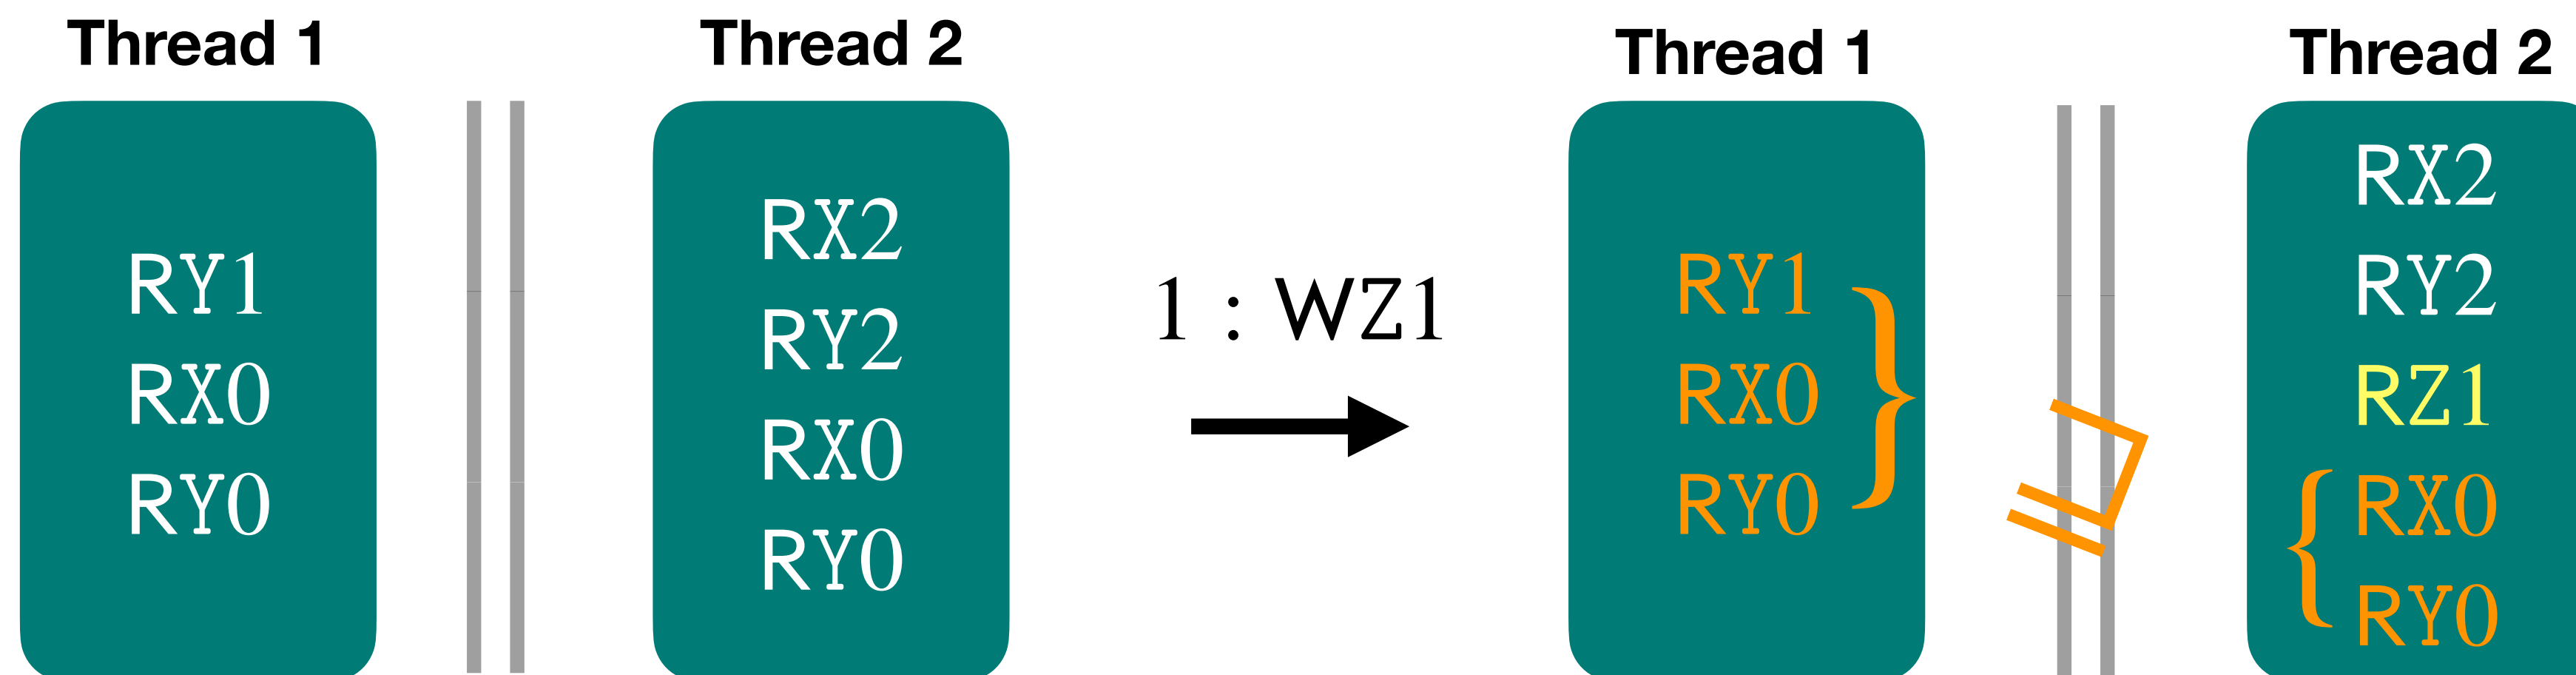
```
X = Y = 0  
Y = 1  
||  
a = X // 1  
b = Y // 0  
X = 1
```



This transition **should not be allowed**

# Shared-Memory Causality Principle

*Every sequence of reads that thread  $\pi$  can perform after reading from a certain write executed by thread  $\tau$  could be performed by thread  $\tau$  immediately after it executed the write.*

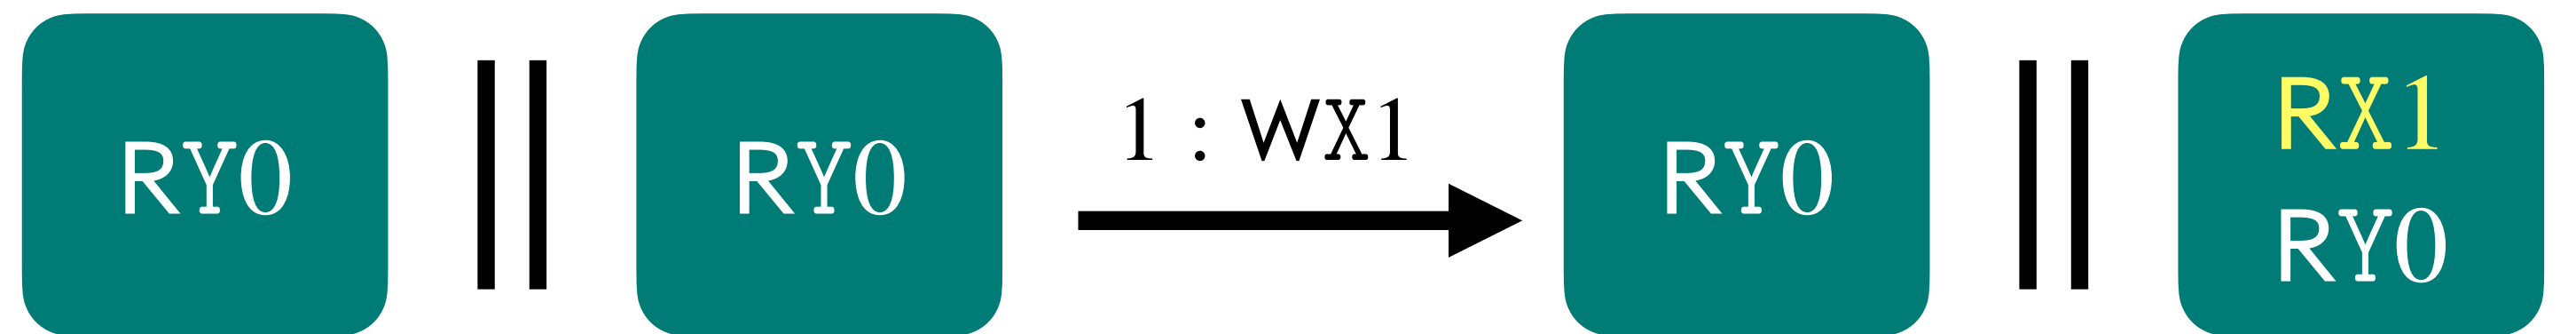




# Where?

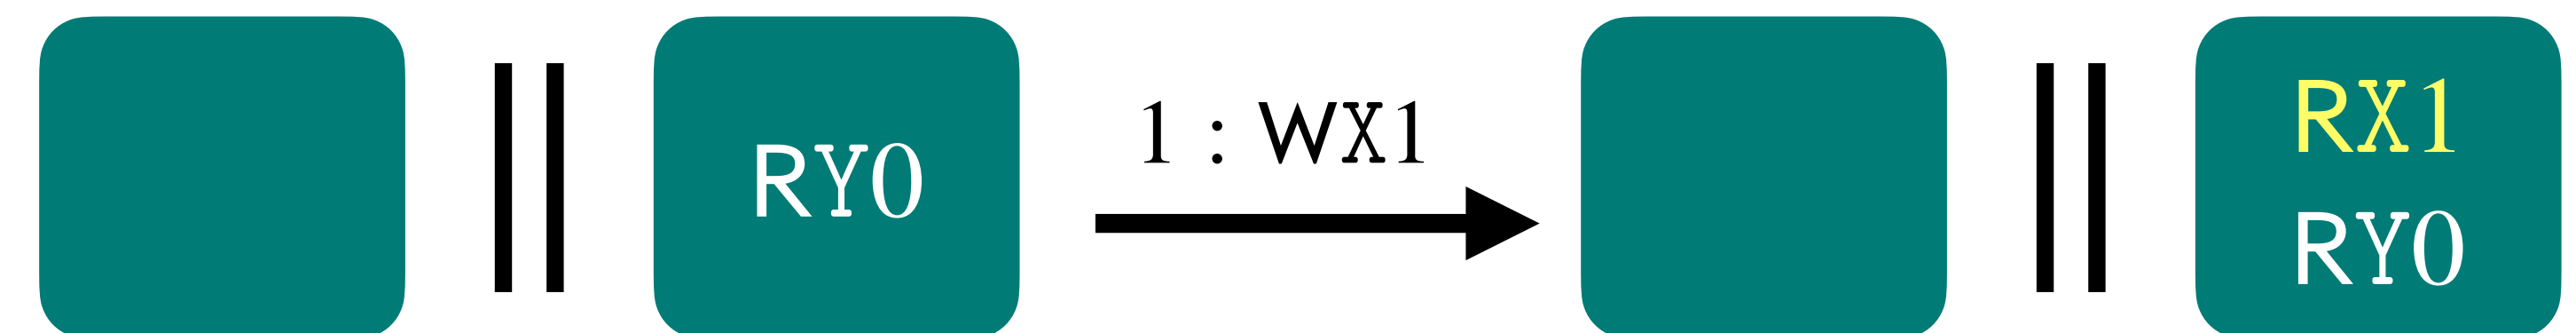
```
X = Y = 0
|||
X = 1
|||
a = X // 1
b = Y // 0
```

```
X = Y = 0
|||
Y = 1
|||
X = 1
|||
a = X // 1
b = Y // 0
```



Thread 1  
can read  
Y=0

This transition **should be allowed**



Thread 1  
cannot read  
Y=0

This transition **should not be allowed**

# More Details



- *Multiple lists* per thread
- **Writer thread's id** in “read options”
- Additional flags to handle **RMWs** (atomic Read-Modify-Writes)

$$\begin{array}{l}
 \text{WRITE} \\
 \forall \pi \in \text{Tid}, L' \in \mathcal{B}'(\pi). \exists n \geq 0, u_1, \dots, u_n, L_0, \dots, L_n. \\
 L' = L_0 \cdot \langle \tau, x, v_W, u_1 \rangle \cdot L_1 \cdot \dots \cdot \langle \tau, x, v_W, u_n \rangle \cdot L_n \\
 \wedge L_0 \cdot \dots \cdot L_n \in \mathcal{B}(\pi) \wedge L_1 \cdot \dots \cdot L_n \in \mathcal{B}(\tau) \\
 \wedge \forall o \in L_1 \cdot \dots \cdot L_n. \text{loc}(o) \neq x \\
 \wedge \forall o \in L_0. \text{loc}(o) = x \implies \pi \neq \tau \wedge \text{rmw}(o) = \text{R} \\
 \hline
 \mathcal{B} \xrightarrow{\tau, W(x, v_W)}_{\text{loSRA}} \mathcal{B}'
 \end{array}$$

$$\begin{array}{l}
 \text{READ} \\
 \text{loc}(o) = x \quad \text{val}(o) = v_R \quad \mathcal{B} = \mathcal{B}'[\tau \mapsto o \cdot \mathcal{B}'(\tau)] \\
 \hline
 \mathcal{B} \xrightarrow{\tau, R(x, v_R)}_{\text{loSRA}} \mathcal{B}'
 \end{array}$$

$$\begin{array}{l}
 \text{RMW} \\
 \text{loc}(o) = x \quad \text{val}(o) = v_R \\
 \text{rmw}(o) = \text{RMW} \\
 \mathcal{B} = \mathcal{B}_{\text{mid}}[\tau \mapsto o \cdot \mathcal{B}_{\text{mid}}(\tau)] \\
 \mathcal{B}_{\text{mid}} \xrightarrow{\tau, W(x, v_W)}_{\text{loSRA}} \mathcal{B}' \\
 \hline
 \mathcal{B} \xrightarrow{\tau, \text{RMW}(x, v_R, v_W)}_{\text{loSRA}} \mathcal{B}'
 \end{array}$$

$$\begin{array}{l}
 \text{LOWER} \\
 \mathcal{B}' \sqsubseteq \mathcal{B} \\
 \hline
 \mathcal{B} \xrightarrow{\varepsilon}_{\text{loSRA}} \mathcal{B}'
 \end{array}$$

# Multiple Lists per Thread

$x := 0$ $x := 1$ $a_1 := z // 1$ $a_2 := y // 0$	$y := 0$ $y := 1$ $b_1 := x // 1$ $b_2 := z // 0$	$z := 0$ $z := 1$ $c_1 := y // 1$ $c_2 := x // 0$	$d_1 := x // 1$ $d_2 := y // 1$ $d_3 := z // 0$	$e_1 := y // 1$ $e_2 := z // 1$ $e_3 := x // 0$	$f_1 := z // 1$ $f_2 := x // 1$ $f_3 := y // 0$
--	--	--	---	---	---

Ry1  
Rx0

Rx1  
Ry0

Rx1  
Ry0

# Writer Thread in “Read Options”

<u><math>x := 0</math></u>	<u><math>y := 0</math></u>	<u><math>a := z // 1</math></u>	<u><math>c := w // 1</math></u>
$x := 1$	$y := 1$	$w := 1$	$d := y // 0$
<u><math>z := 1</math></u>	<u><math>z := 1</math></u>	$b := x // 0$	

$R_z1$	$R_z1$
$R_y0$	$R_x0$

1 : $R_z1$	2 : $R_z1$
$R_y0$	$R_x0$

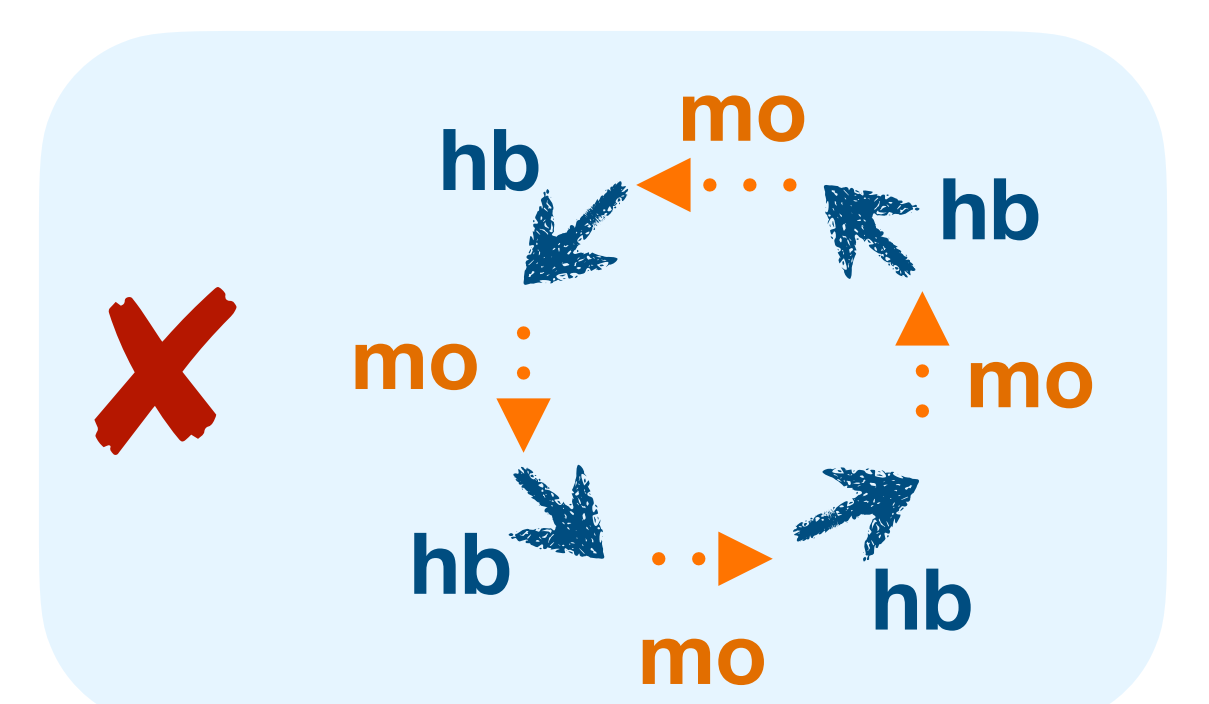
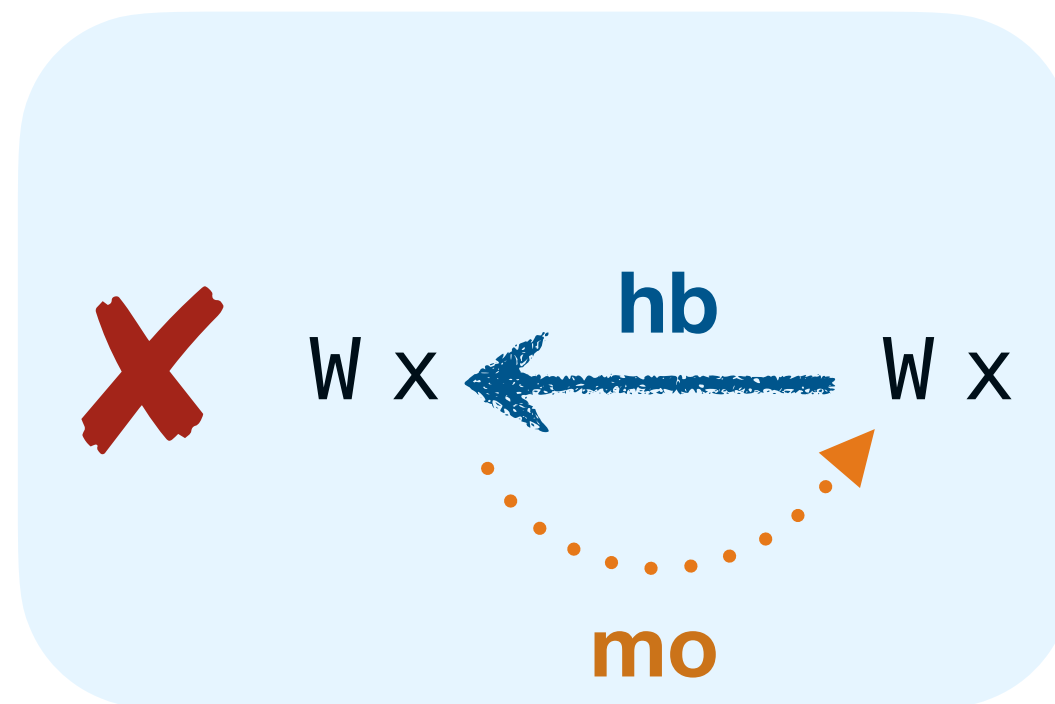
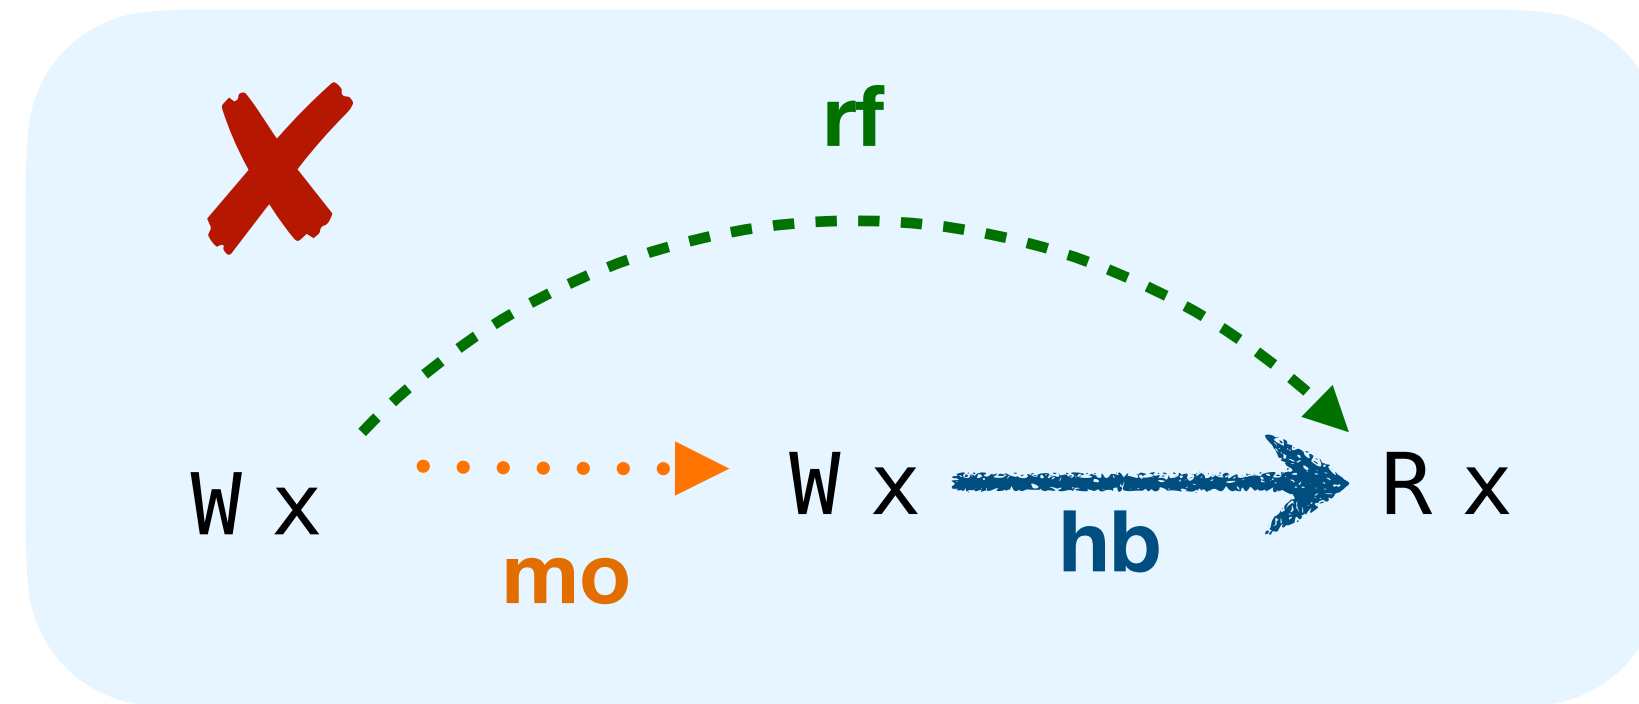
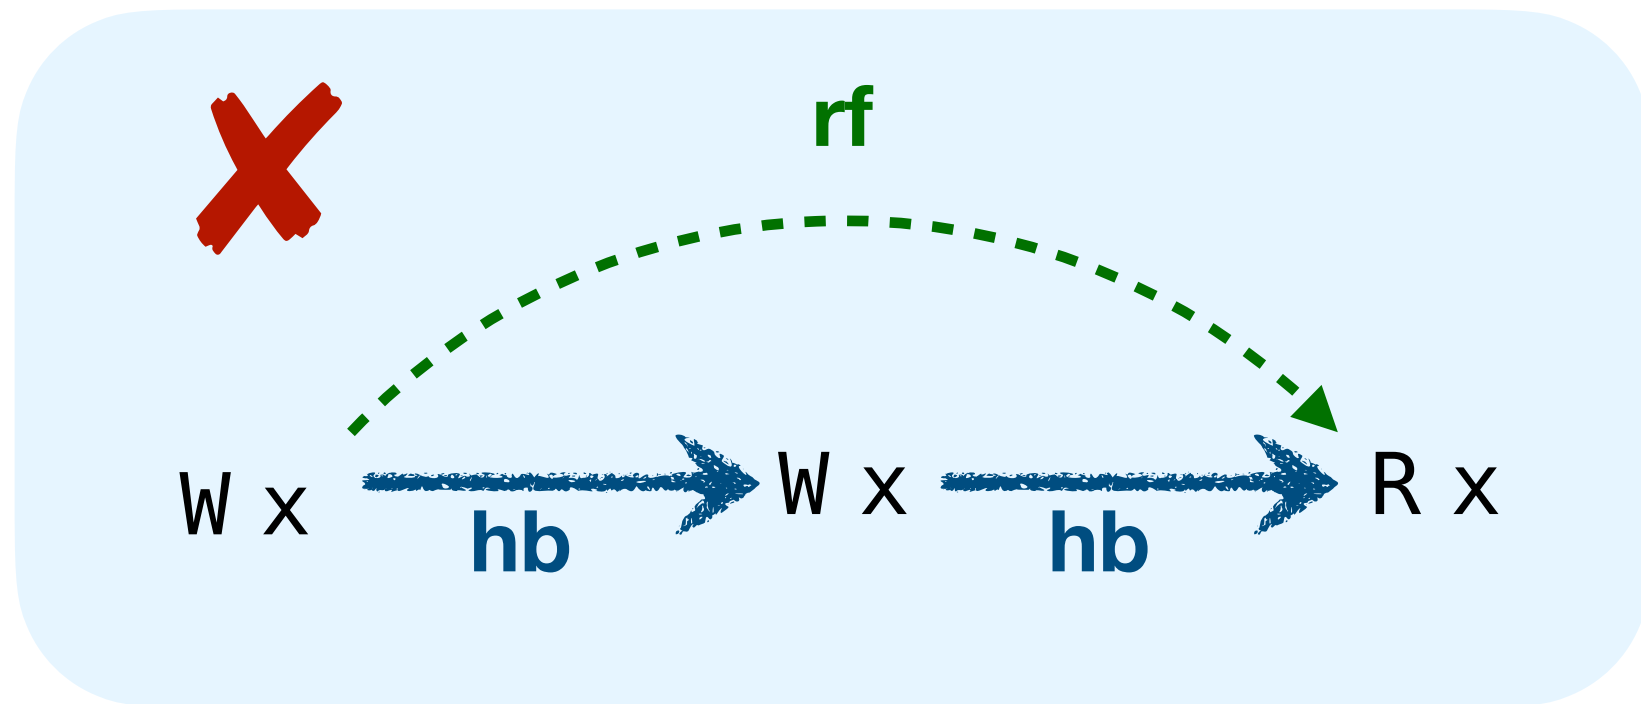
# Three Variants

Weak Release/Acquire  
(WRA)


Release/Acquire  
(RA)

Strong Release/Acquire  
(SRA)

$\exists$  total order on writes to the same location (**modification-order**) s.t.:



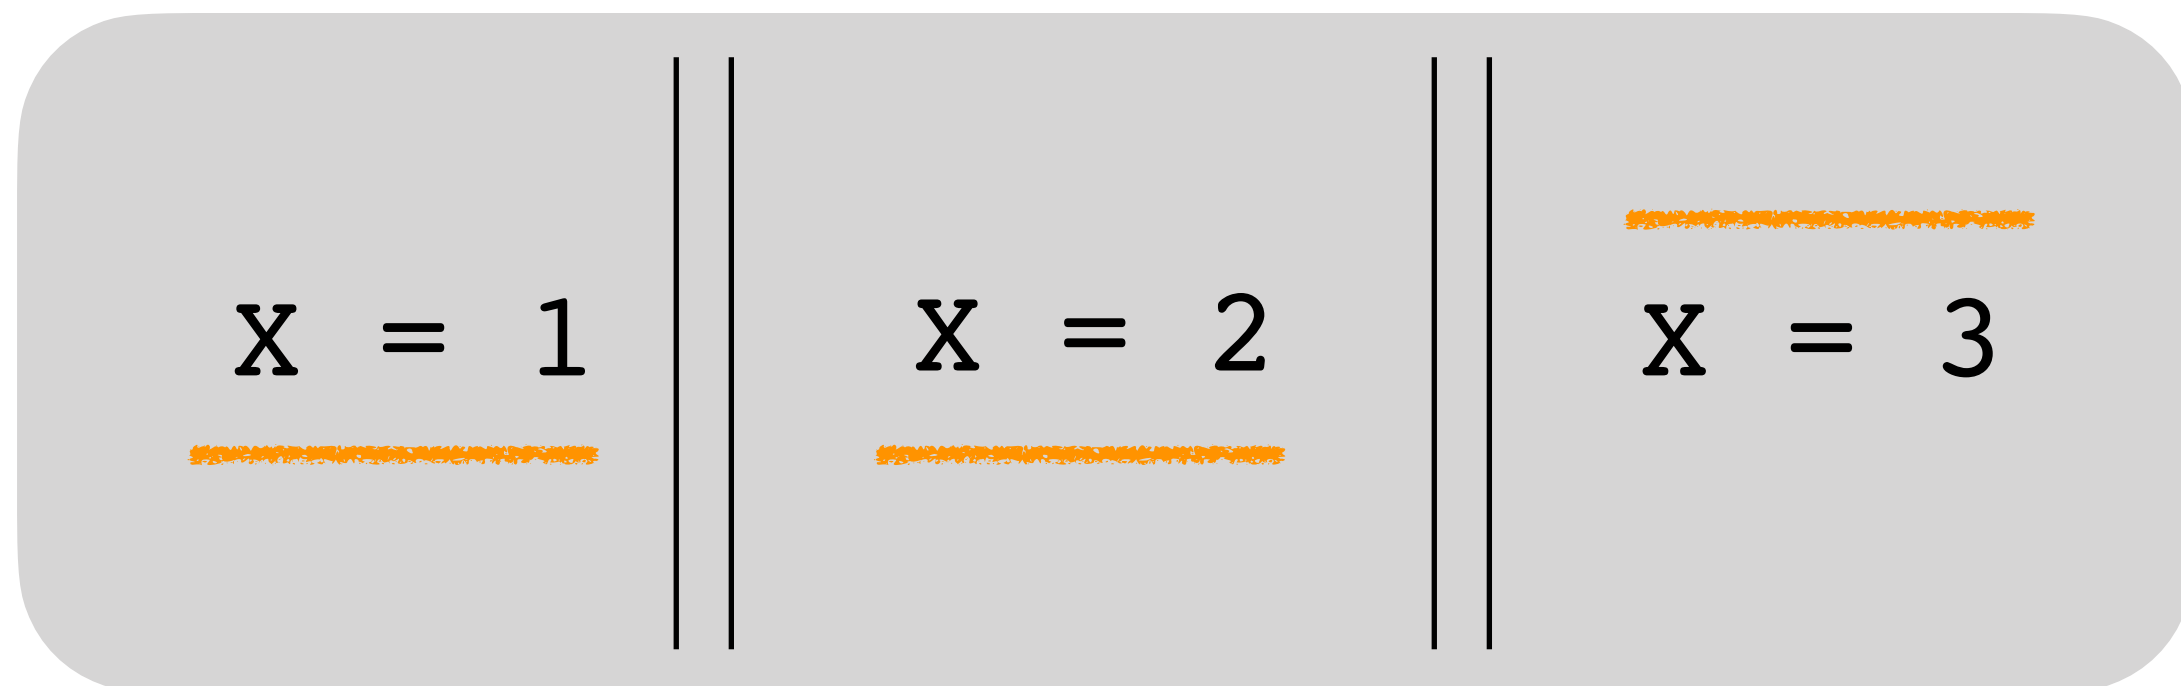
# Potential-Based System for WRA

$\tau : W_{xv}$   


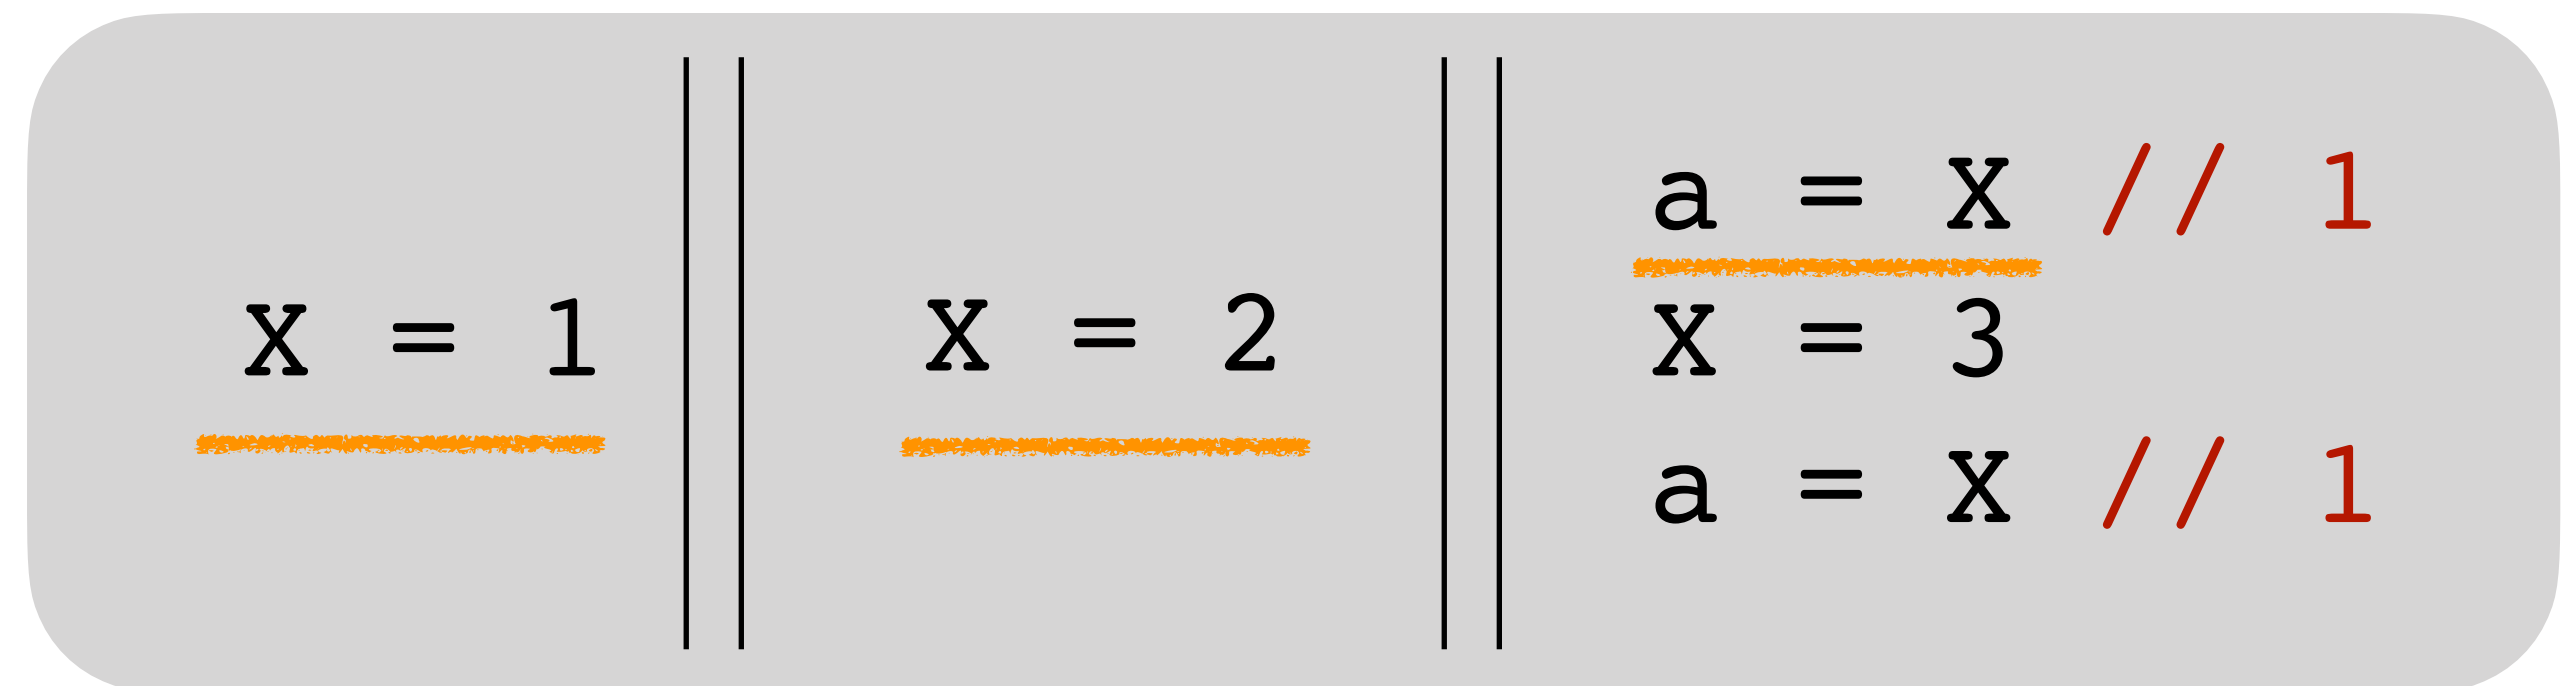
## Write steps

Precondition: no  $Rx\_$  in  $\tau$ 's potential 

All threads may get new options  $R_{xv}$



RX1  
 RX2  
 RX1

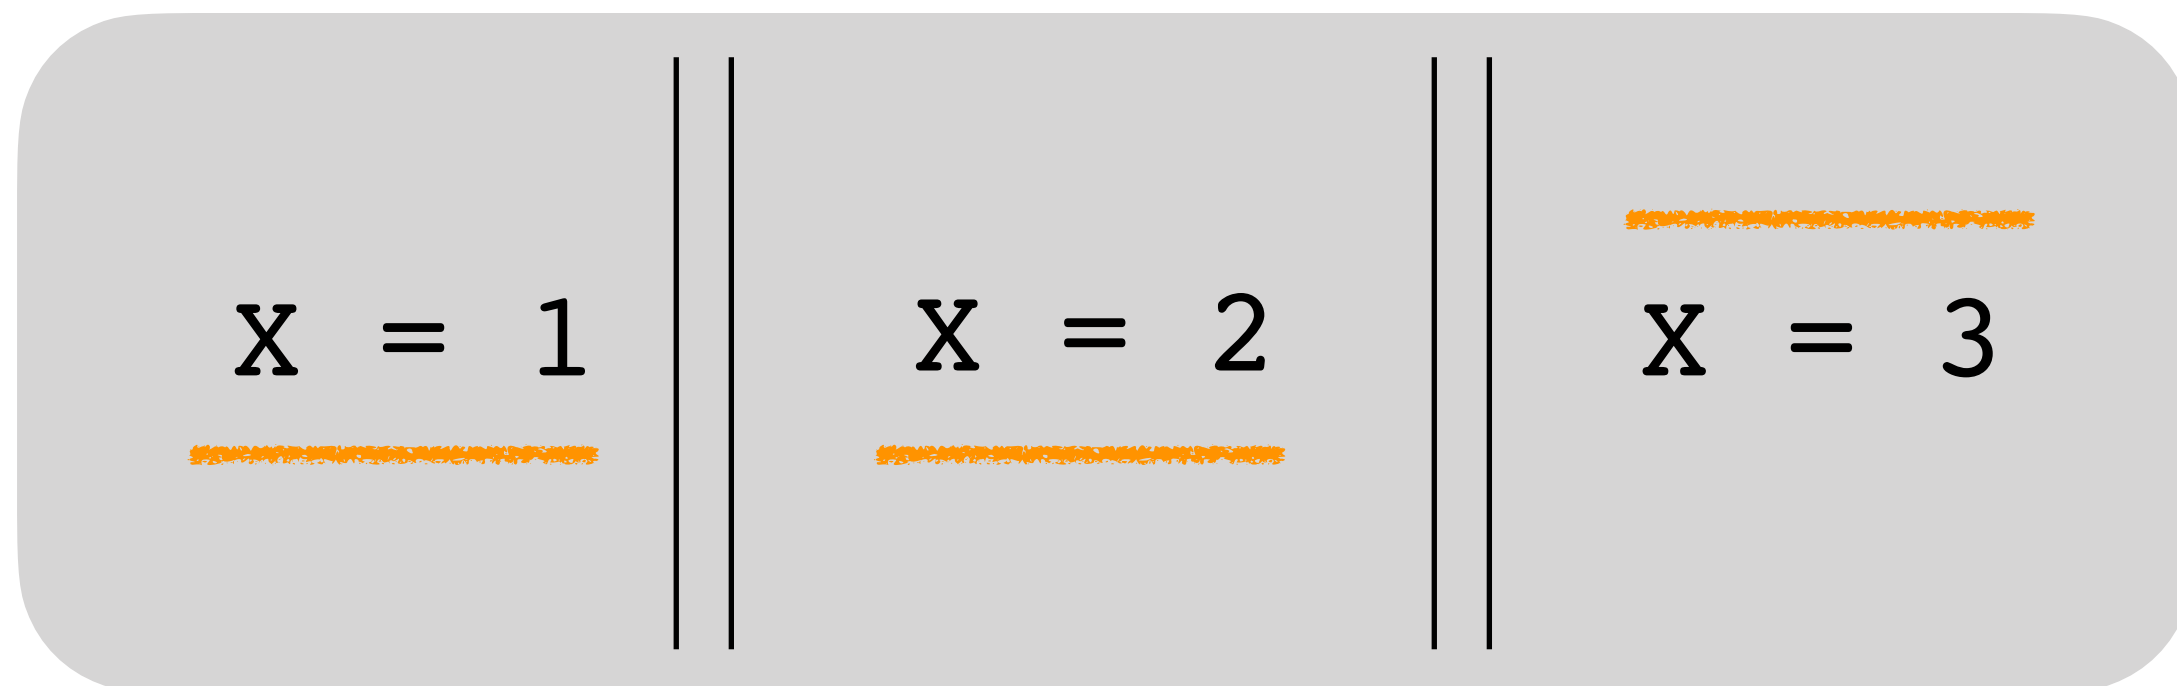


~~RX1~~  
 RX2  
 RX1

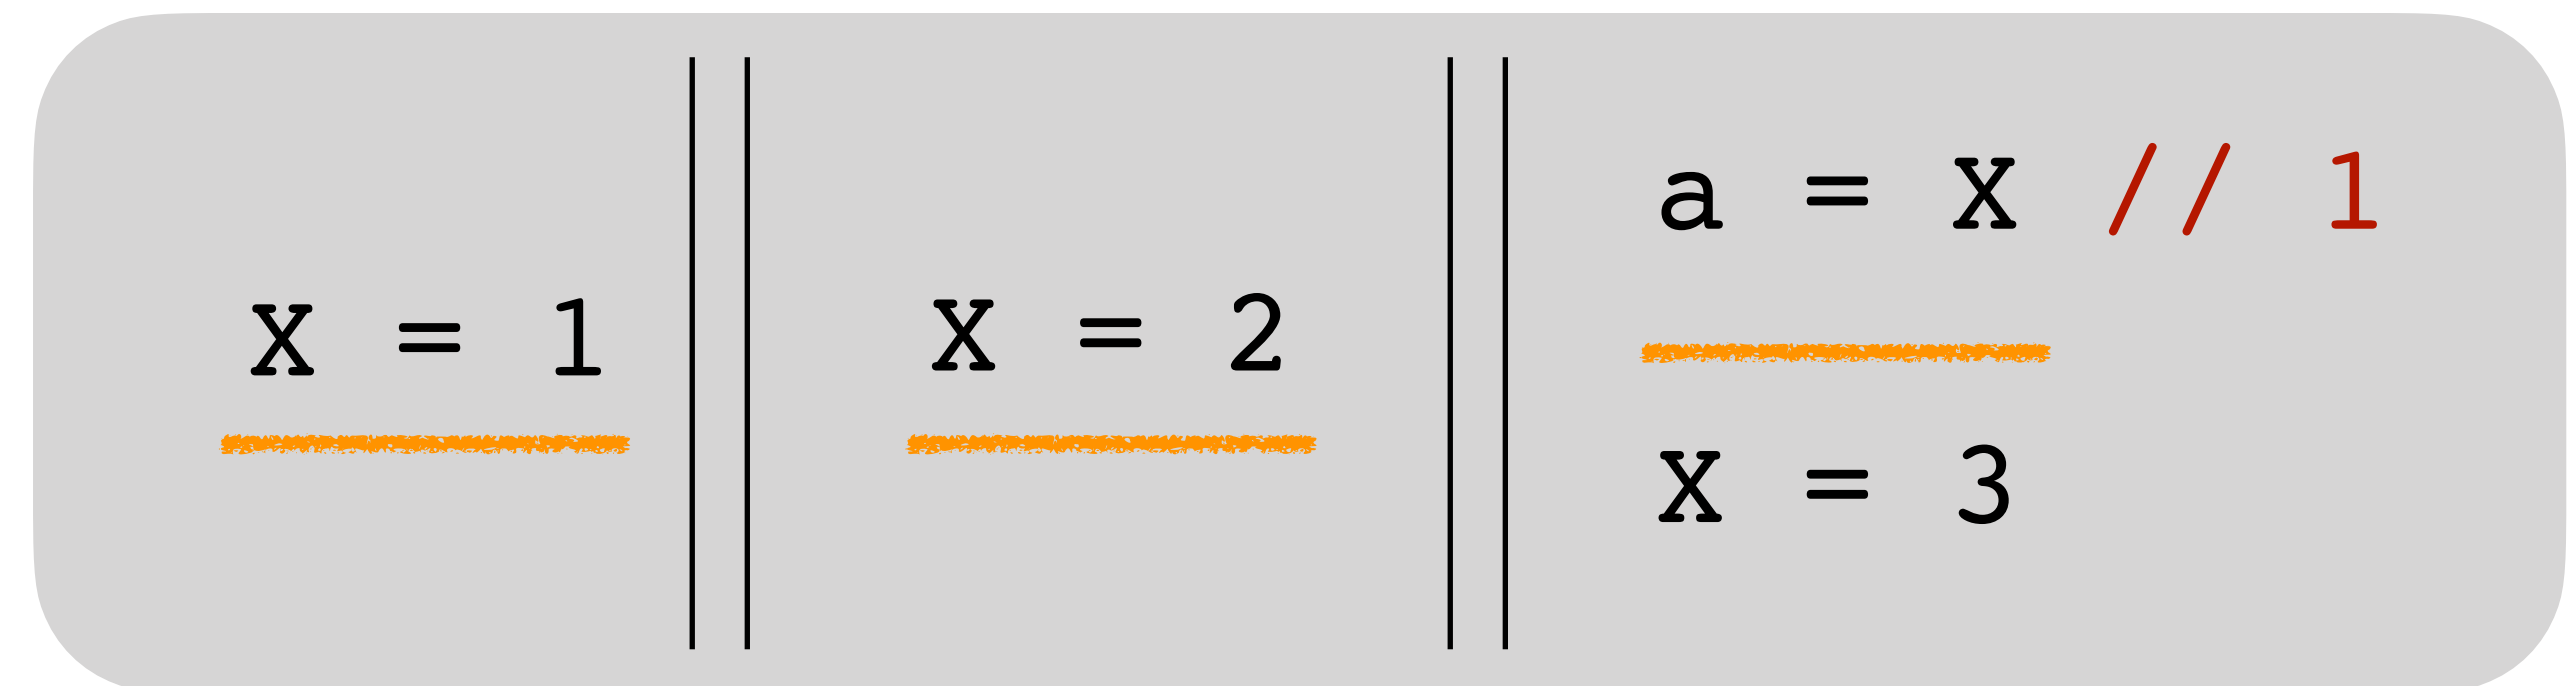
# Potential-Based System for WRA

Use "**write options**" to mark when writes are allowed

+ Simple constraints on where read options are added wrt write options



WX  
RX1  
RX2  
RX1

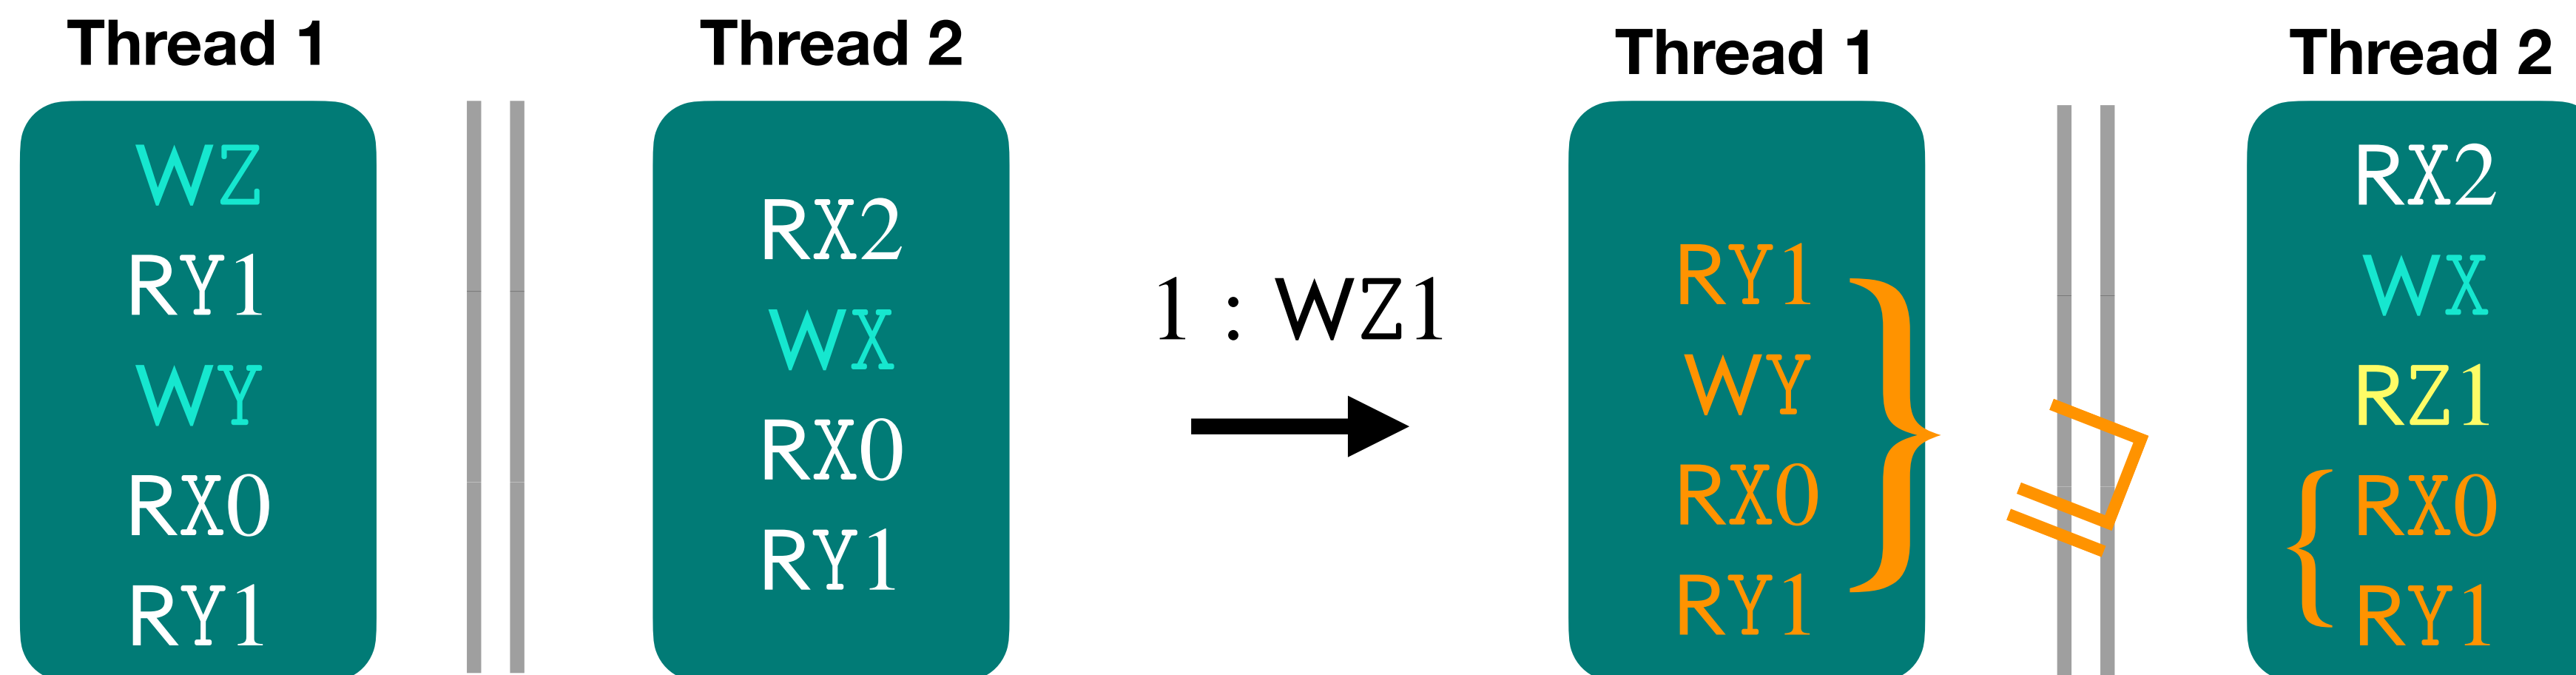


~~RX1~~  
RX2  
RX1



# Shared-Memory Causality Principle

*Every sequence of reads and writes that thread  $\pi$  can perform after reading from a certain write executed by thread  $\tau$  could be performed by thread  $\tau$  immediately after it executed the write.*

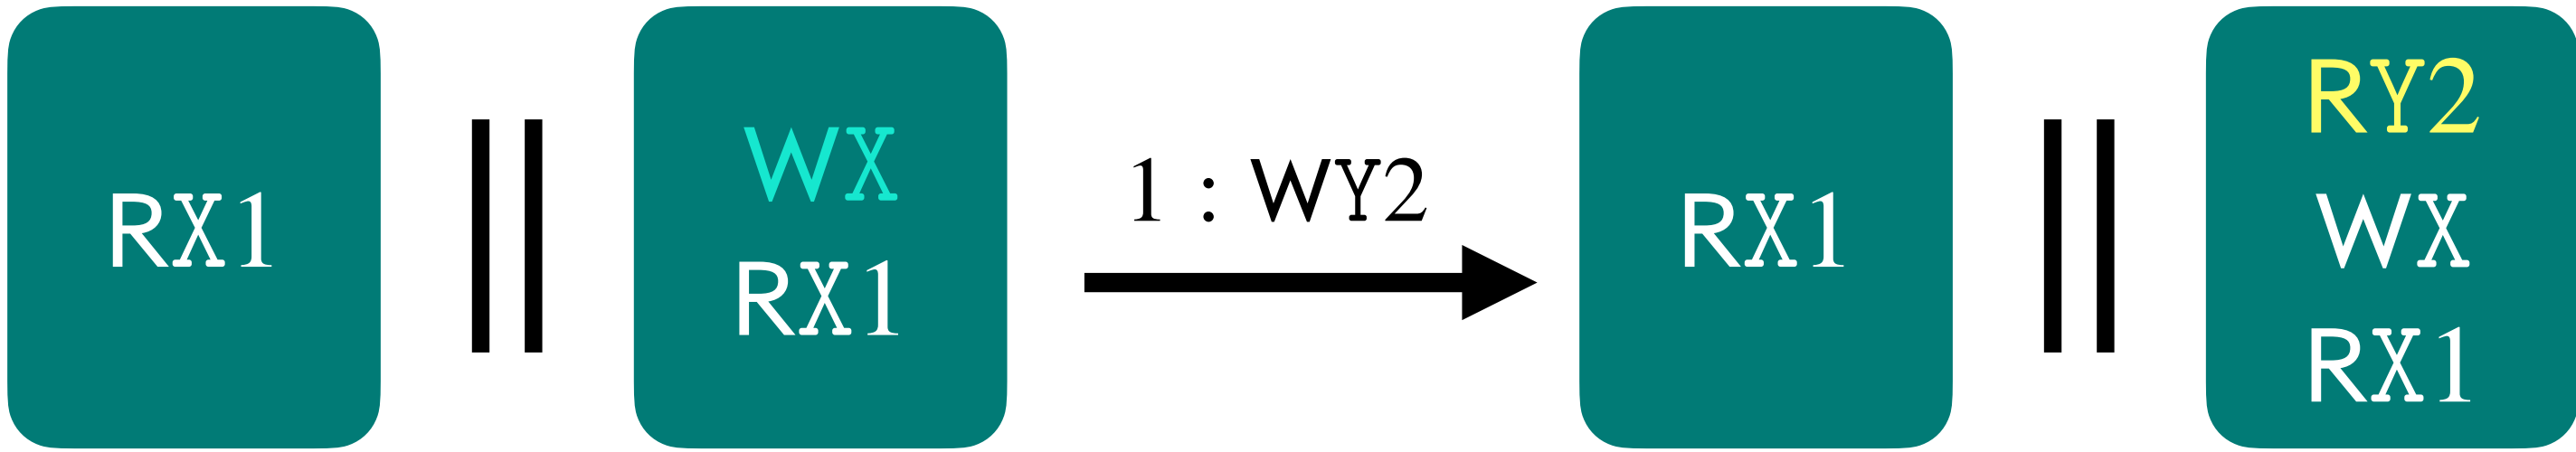




```

X = Y = 0
||
X = 1
Y = 2
||
a = Y // 2
X = 3
b = X // 1

```



**Thread 1  
cannot write  
to X and then  
read X=1**

This transition **is disallowed**

# Results

## Theorem

The potential-based memory systems are **equivalent** to the SRA/WRA systems.



## Theorem

When synchronized with a (finite-state) concurrent program, the potential-based memory systems form **WSTS**.

# Results

## Theorem

The verification problems under **SRA** and **WRA** are **decidable**.

## Corollary

The verification problem under **RA** is **decidable** for write-write-race-free programs.

# Research Questions

- Useful **implementation**
- **RA** without RMWs?
- **Other models** and extensions of causal consistency (get closer to RA?)
- **Parametrized** programs
- Use the potential-based semantics for **other verification approaches**





**Interested in concurrency & verification?  
I'm looking for students / postdocs!**

