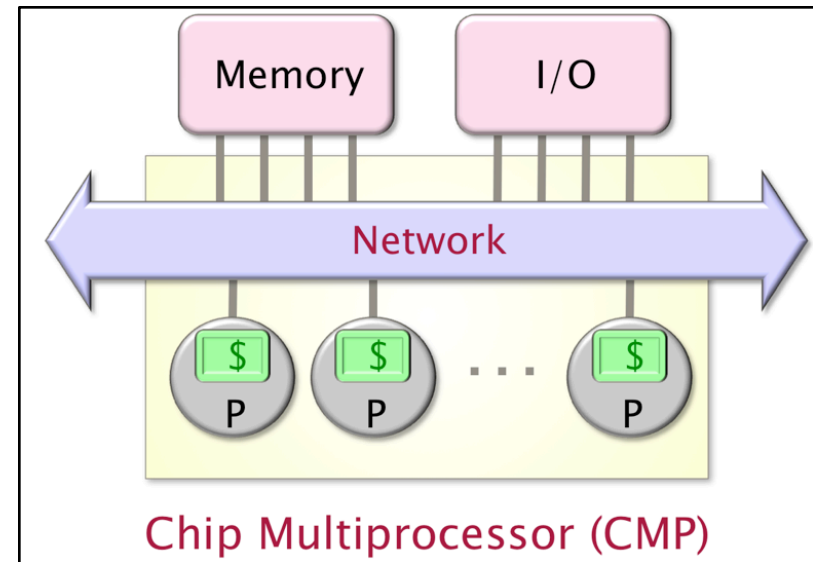



FastTrack: Efficient and Precise Dynamic Race Detection (+ identifying destructive races)

Cormac Flanagan
UC Santa Cruz

Stephen Freund
Williams College

Multithreading and Multicore



- Multithreaded programming is notoriously difficult, in part due to schedule-dependent behavior
 - race conditions, deadlocks, atomicity violations, ...
 - difficult to detect, reproduce, or eliminate

Race Conditions

- Two threads access a shared variable without synchronization, and at least one thread does a write
- Very common




2003 Blackout (\$6 Billion)

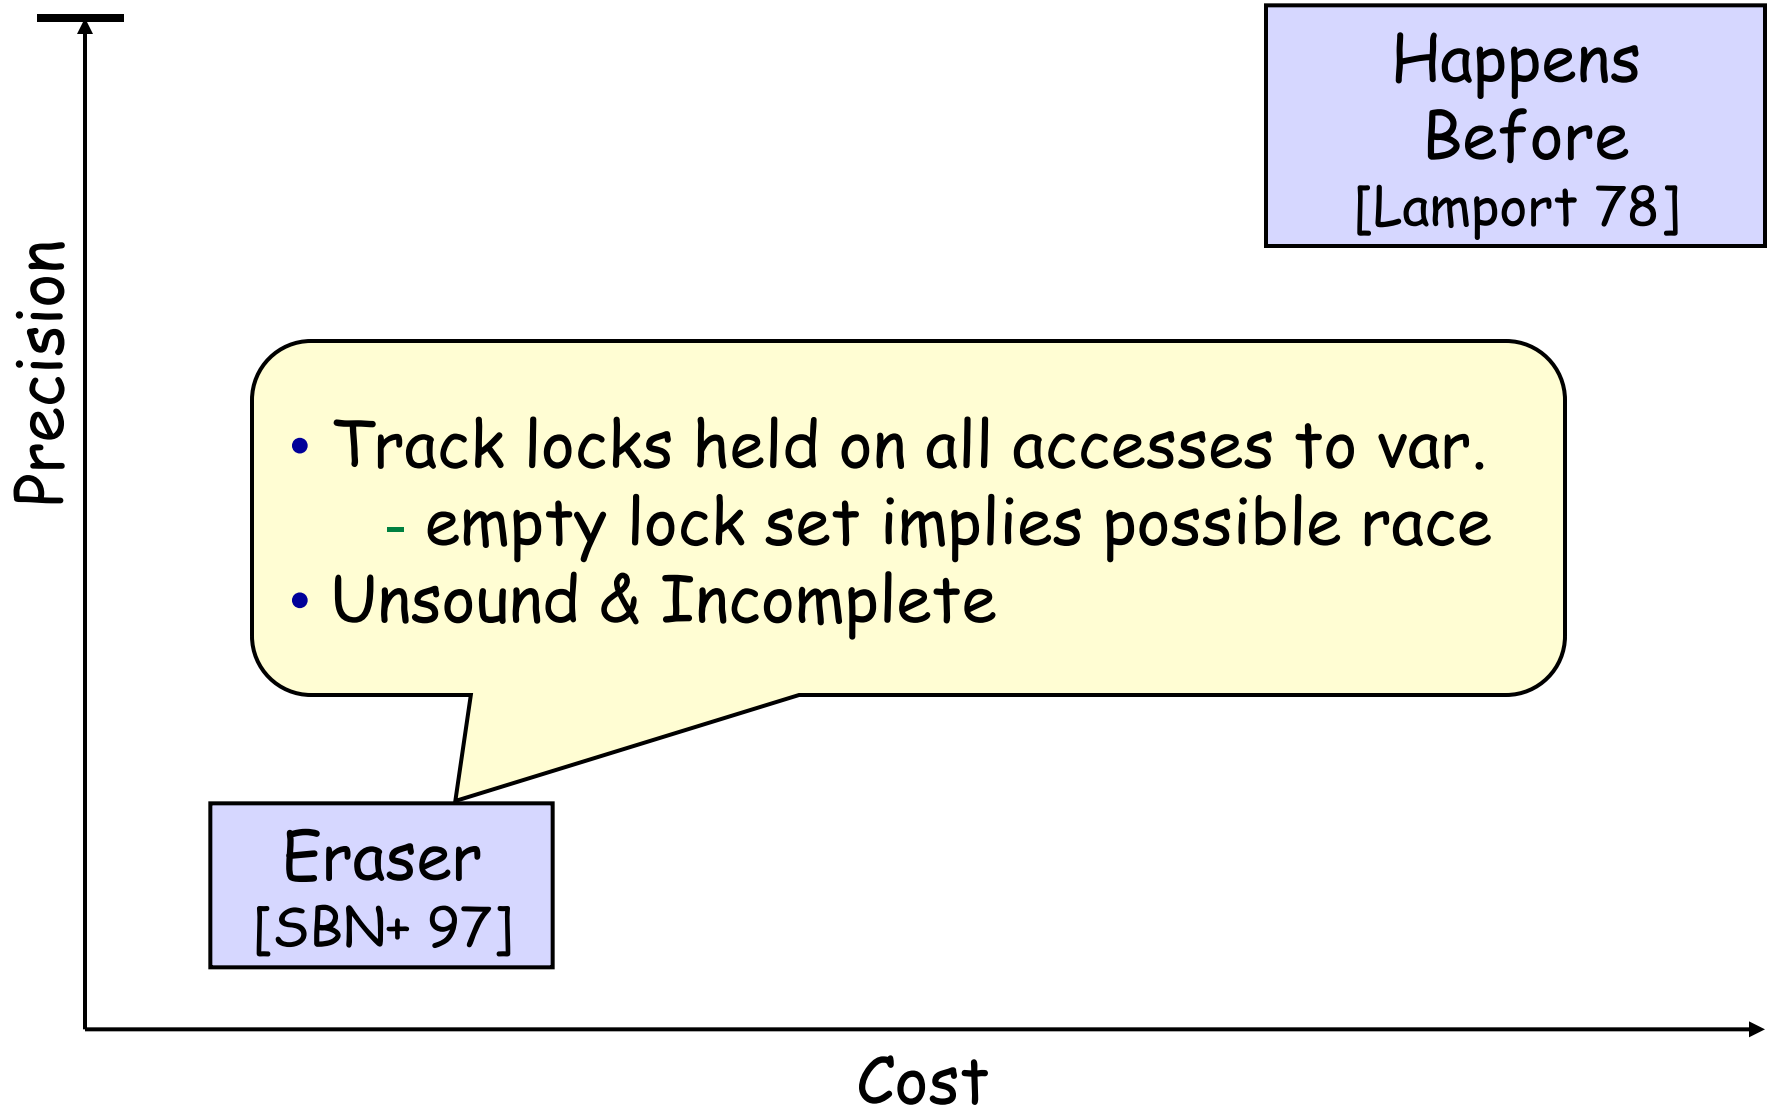


Therac-25

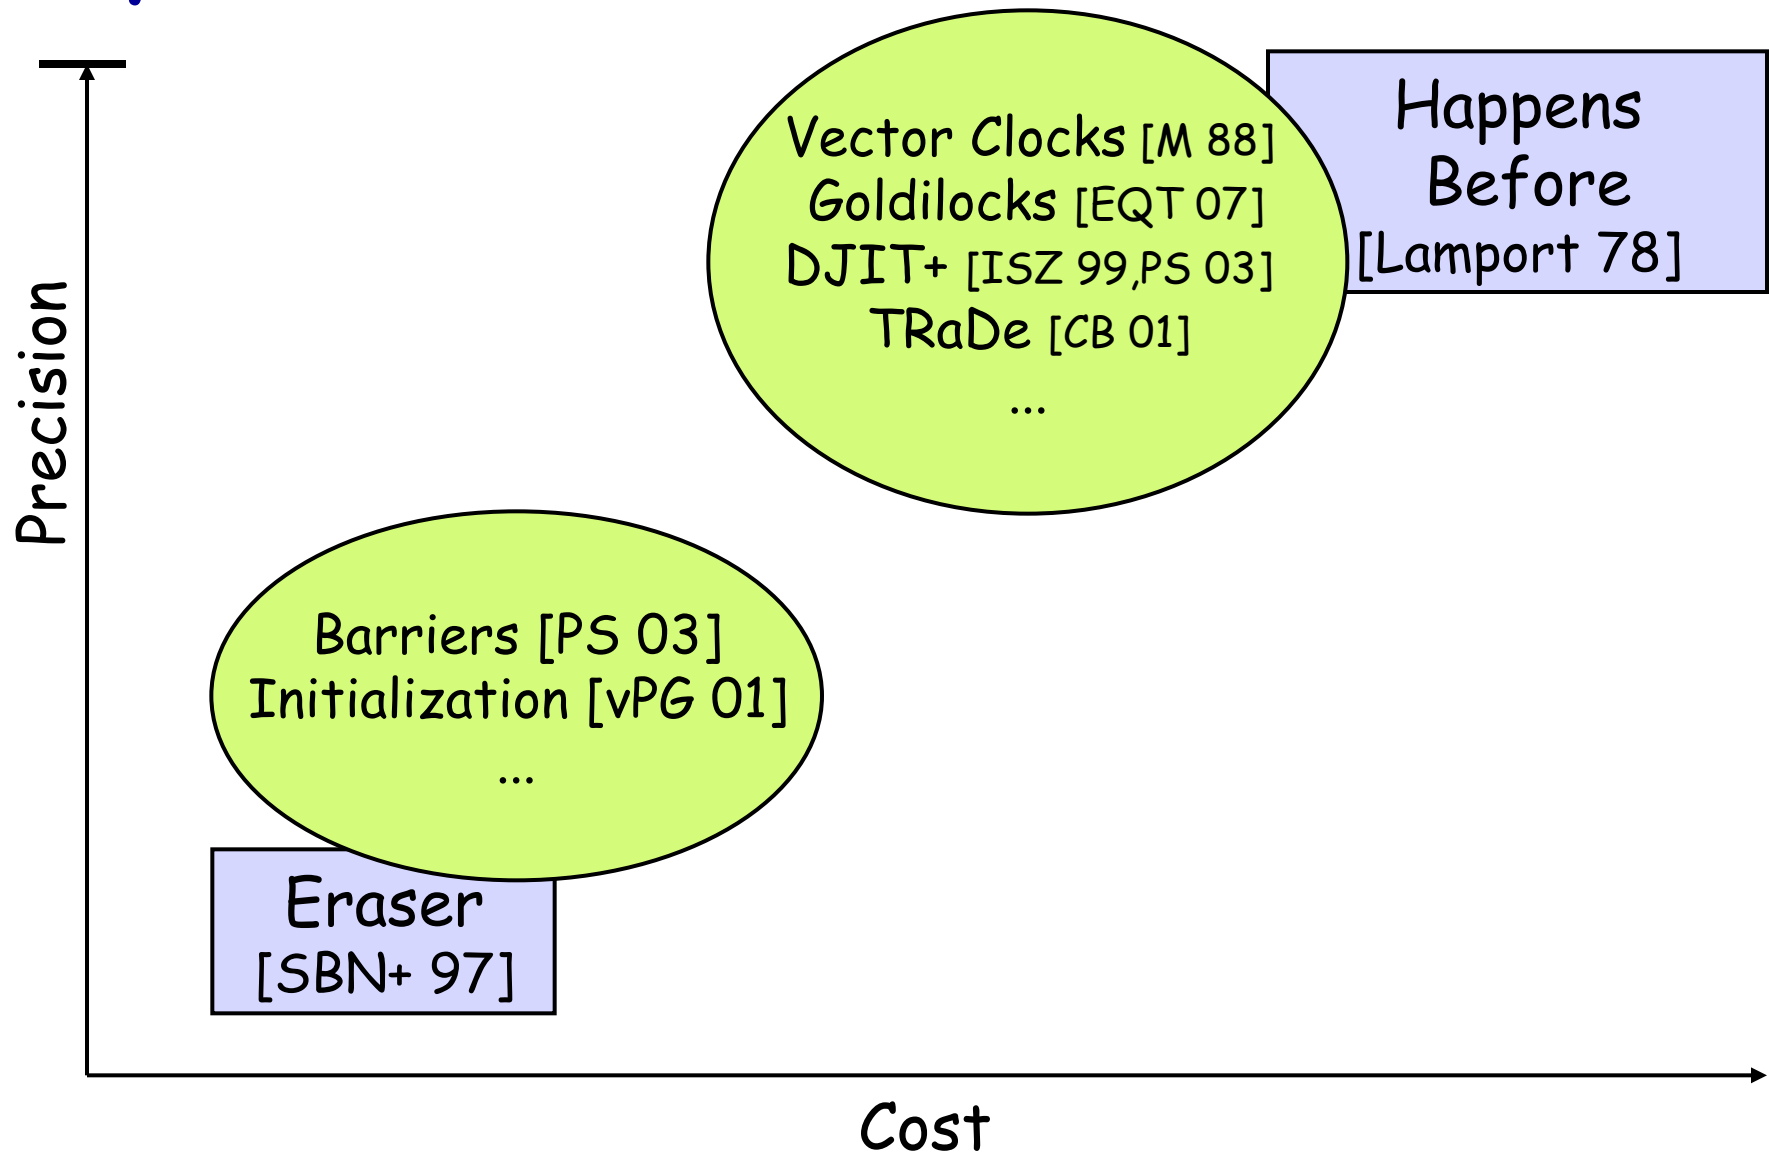
Dynamic Race Detection



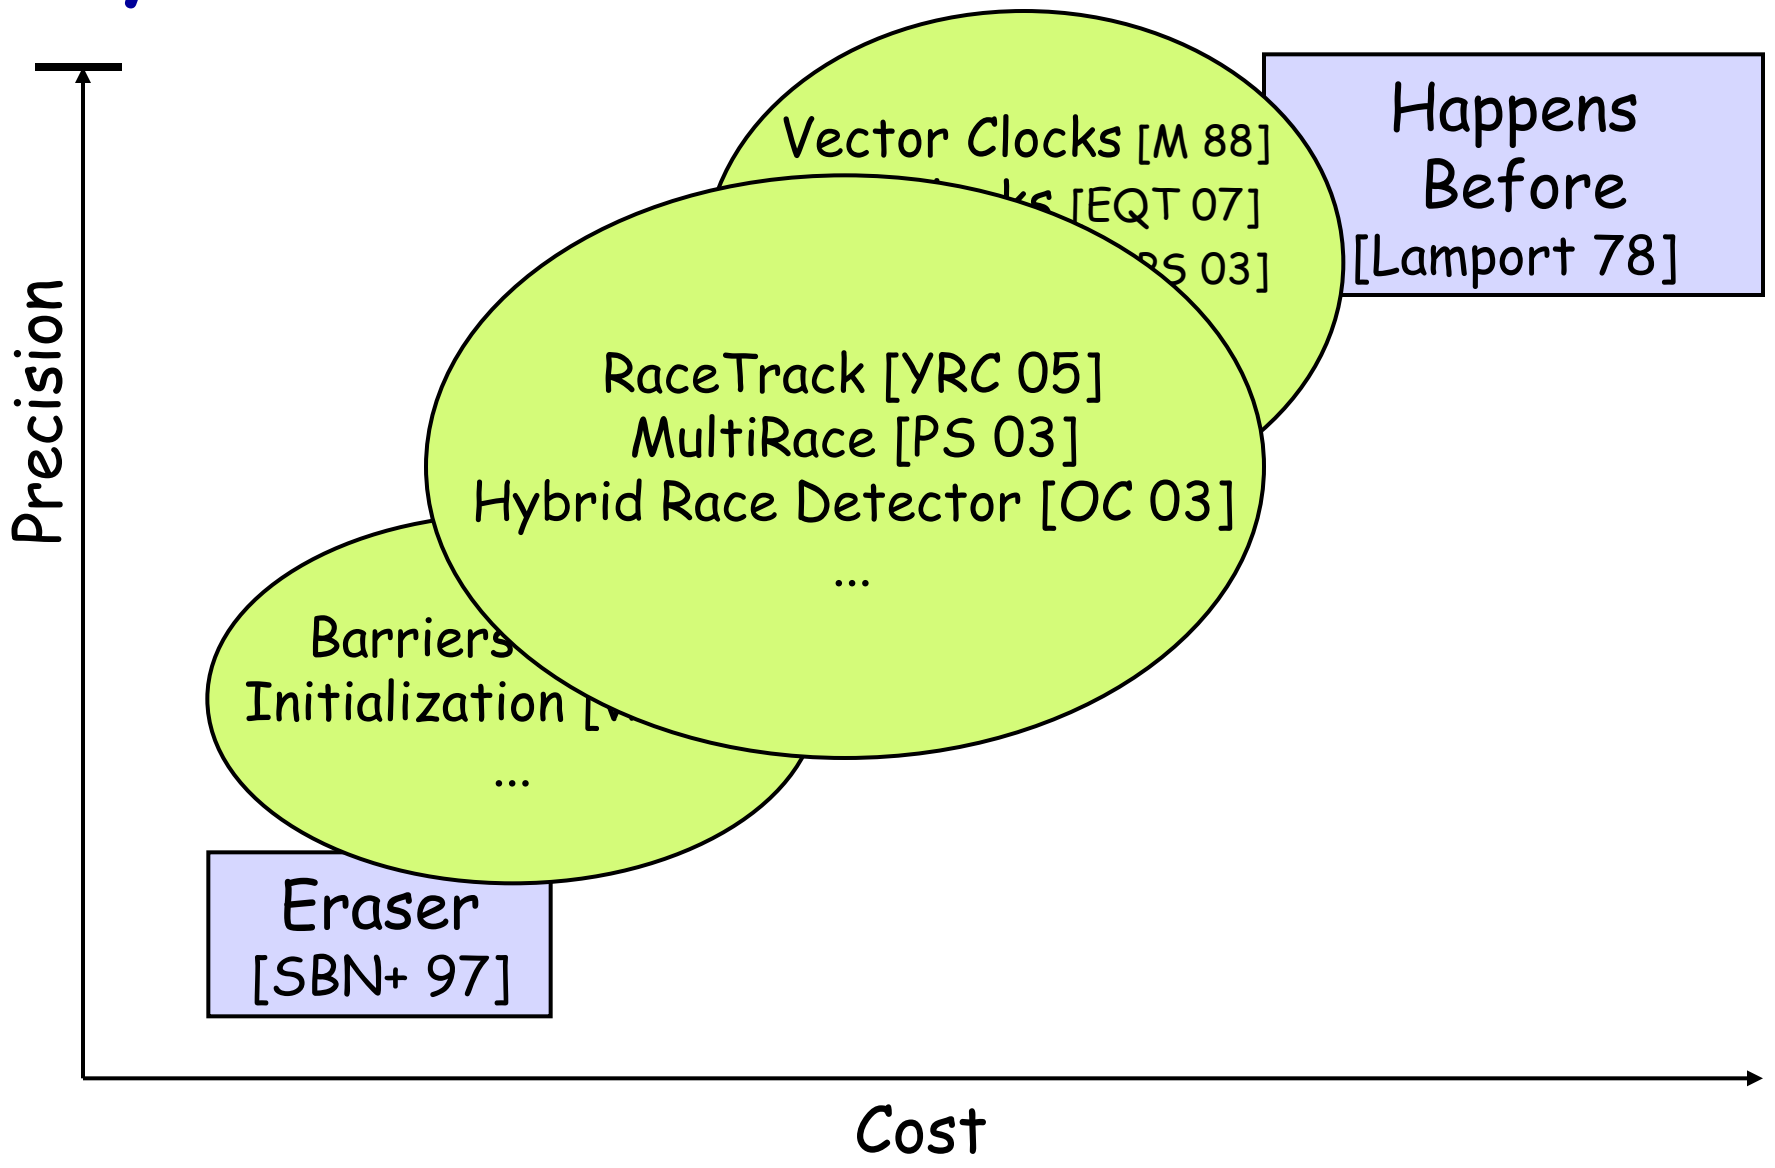
Dynamic Race Detection



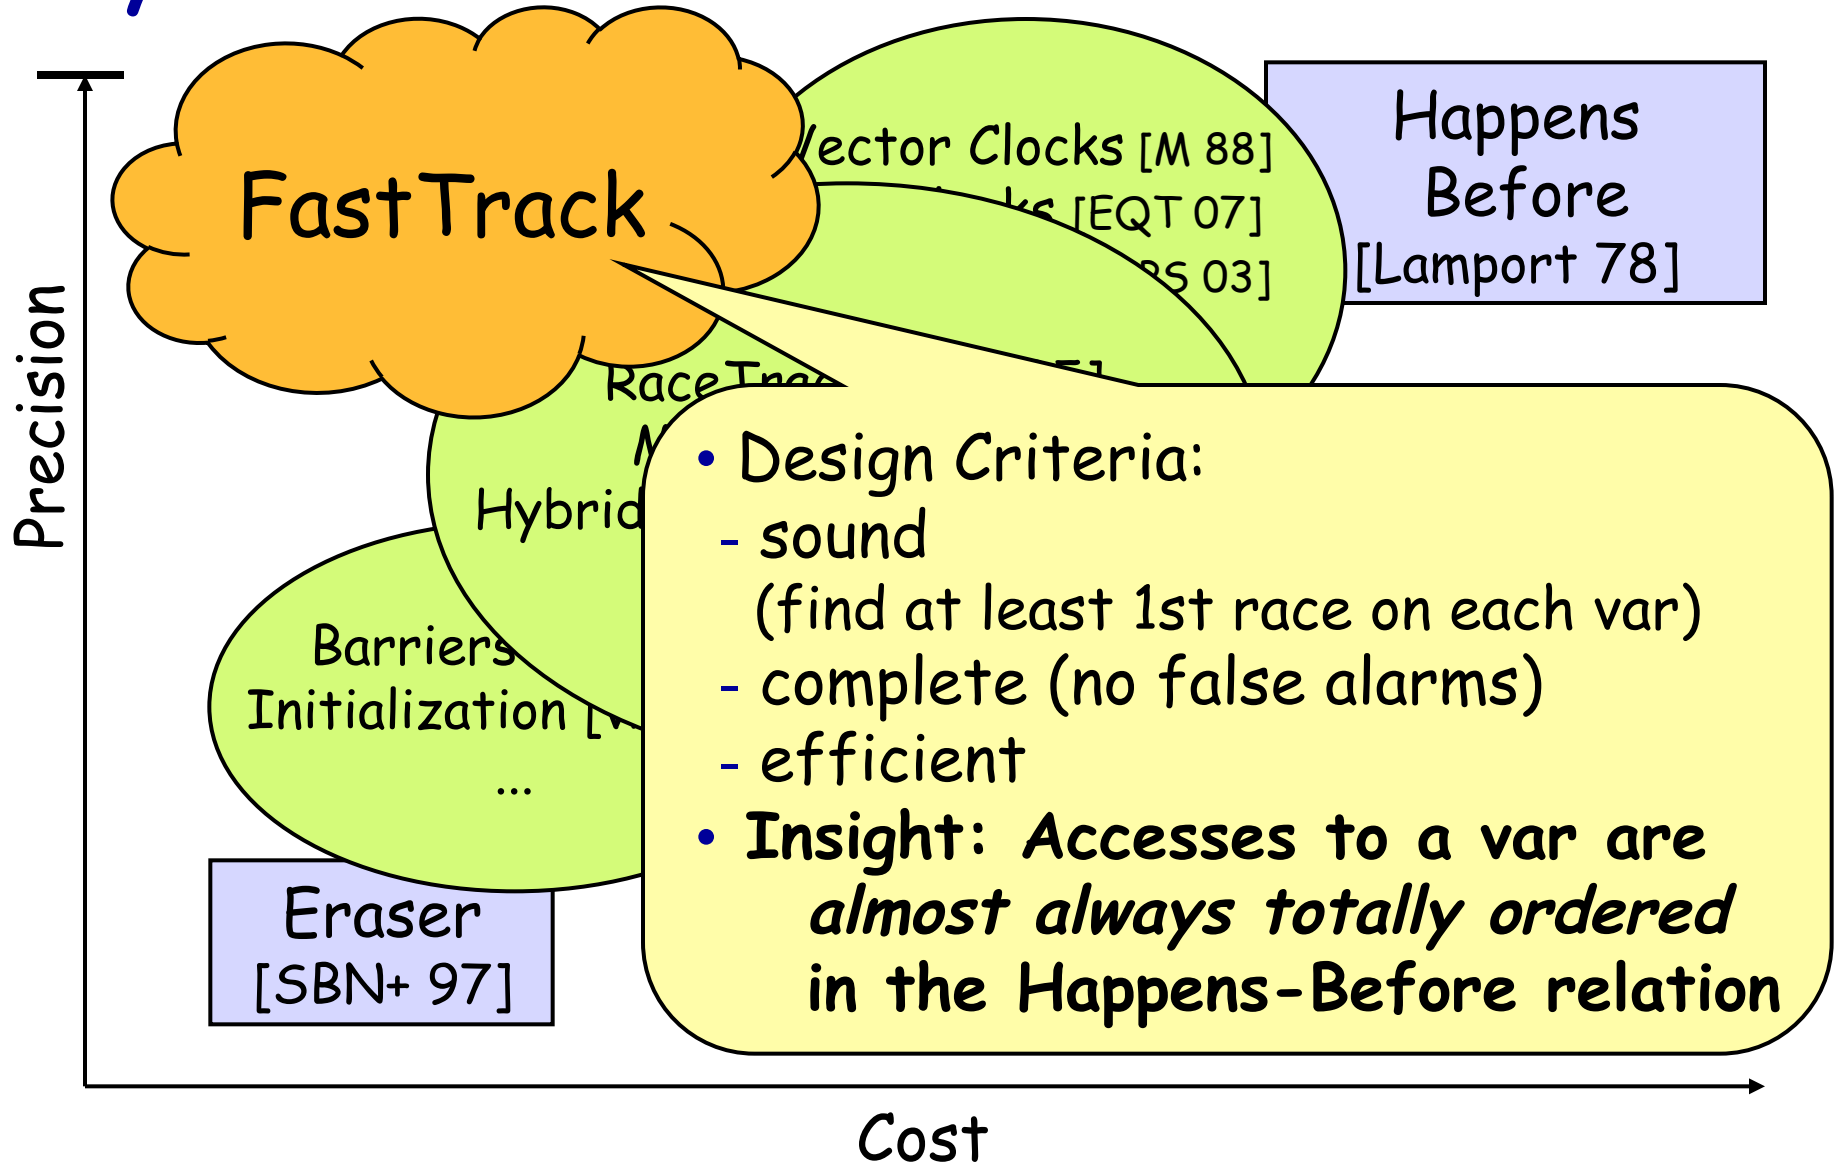
Dynamic Race Detection



Dynamic Race Detection

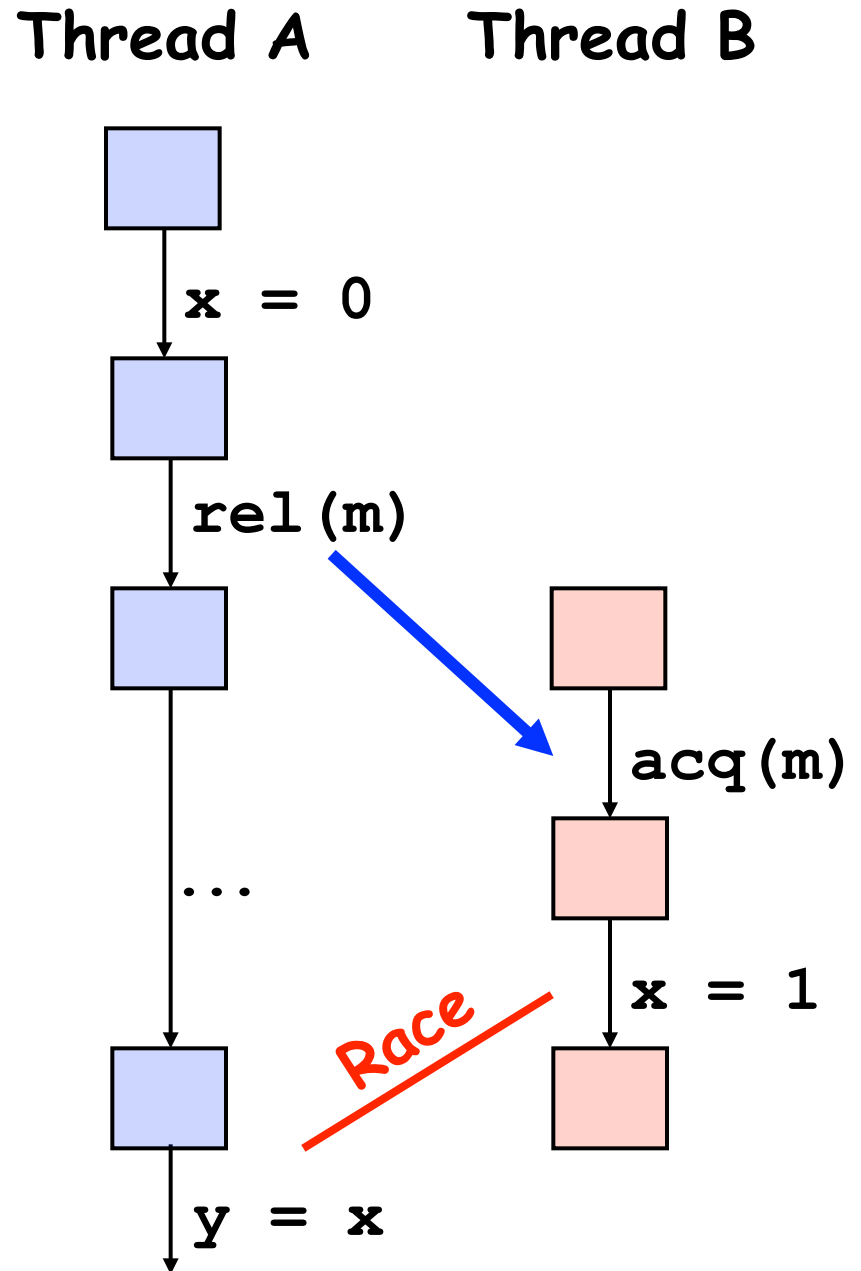


Dynamic Race Detection



Happens-Before

- Event Ordering:
 - program order
 - synchronization order
- Types of Races:
 - Write-Write
 - (write before read)
 - Write-Read
 - (write before read)
 - Read-Write
 - (read before write)



VC_A

4	1
---	---

A B

VC_B

2	8
---	---

A B

L_m

2	1
---	---

A B

W_x






3	0
---	---

A B

R_x

0	1
---	---

A B



Write-Write Check: $W_x \sqsubseteq VC_A$?

3	0
---	---

 \sqsubseteq

4	1
---	---

? **Yes**

Read-Write Check: $R_x \sqsubseteq VC_A$?


0	1
---	---

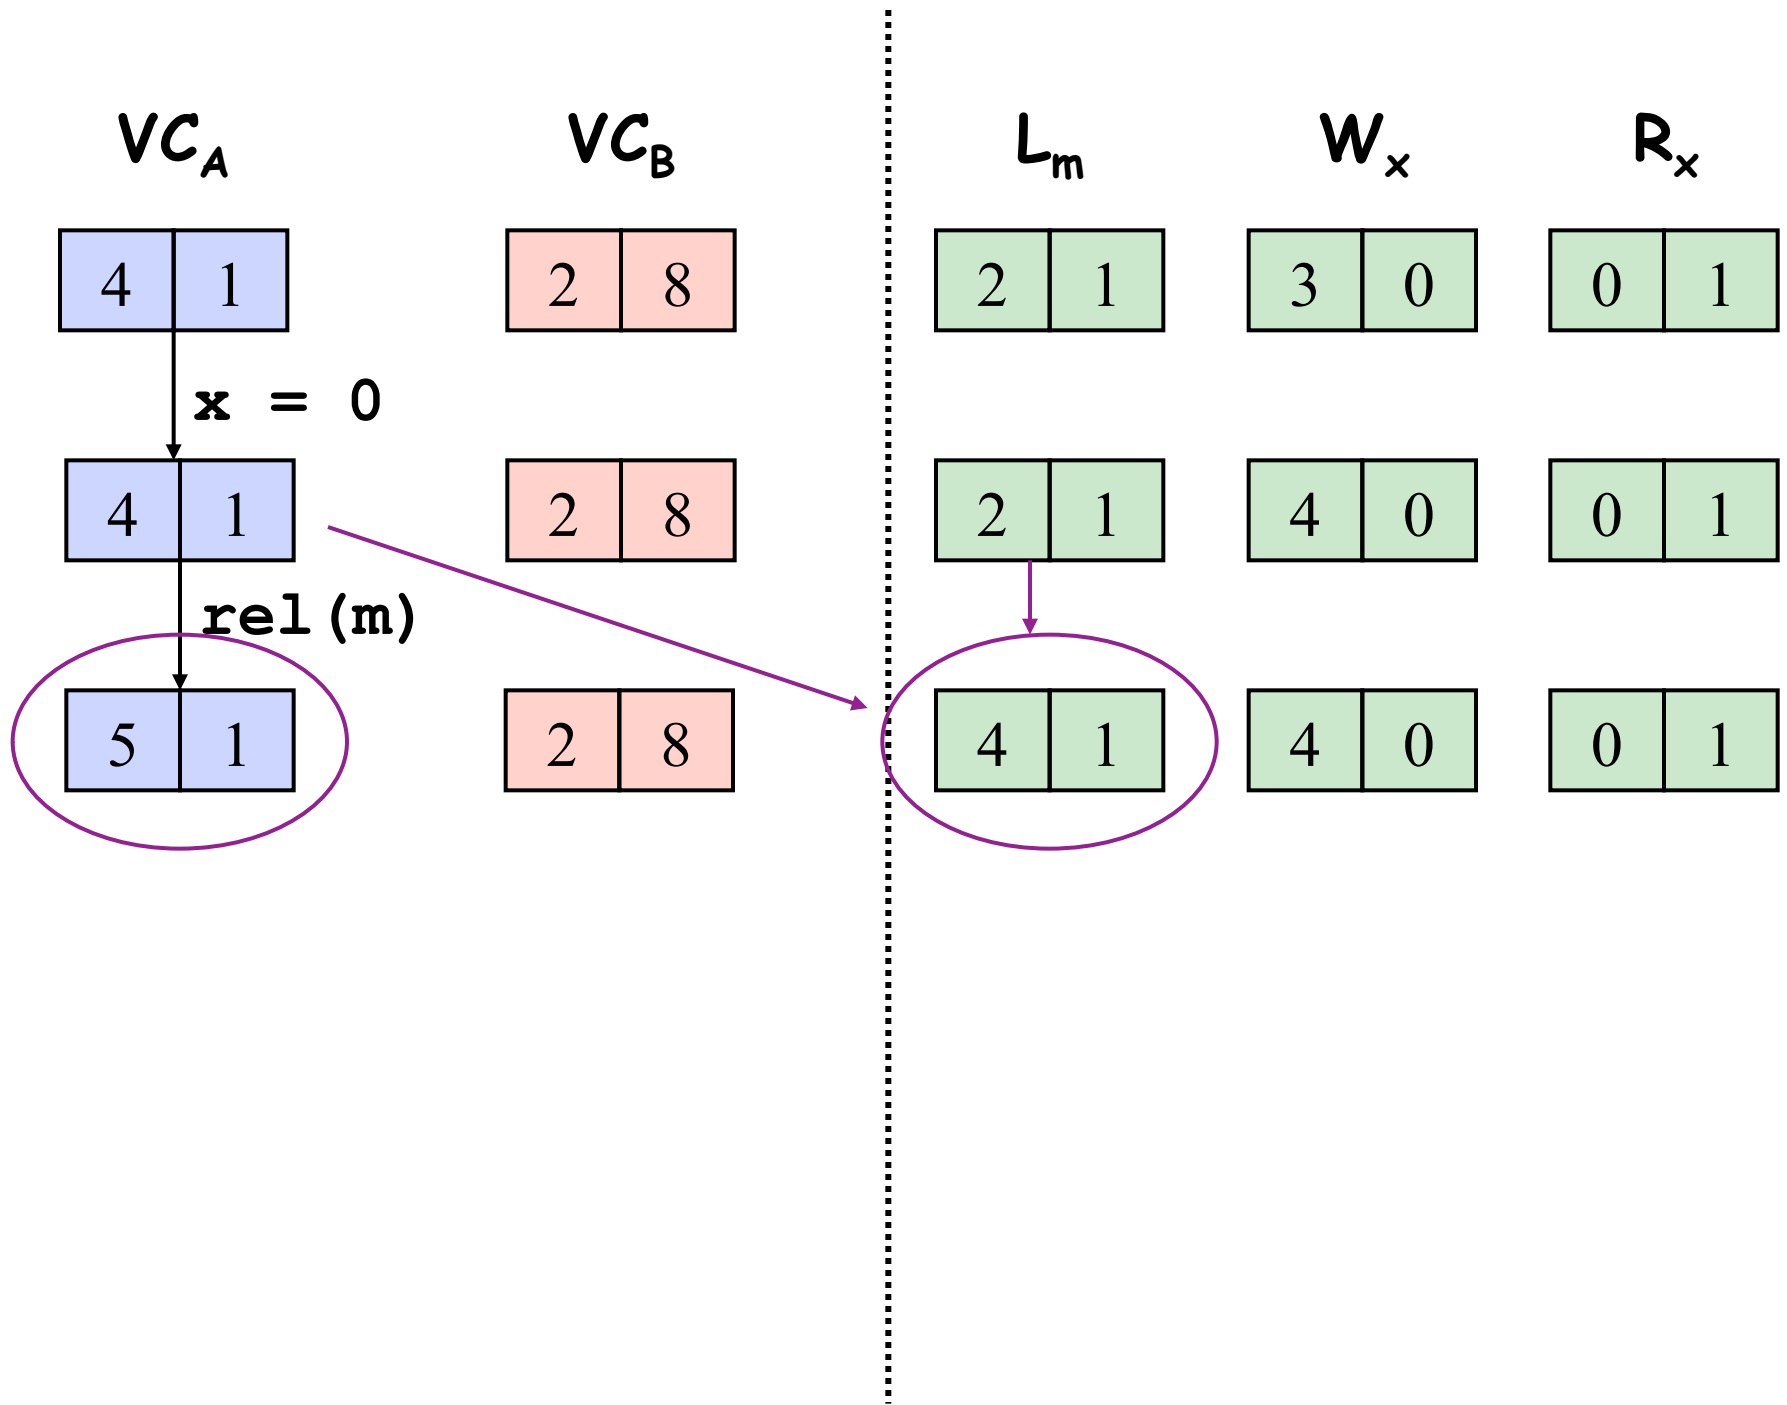
 \sqsubseteq

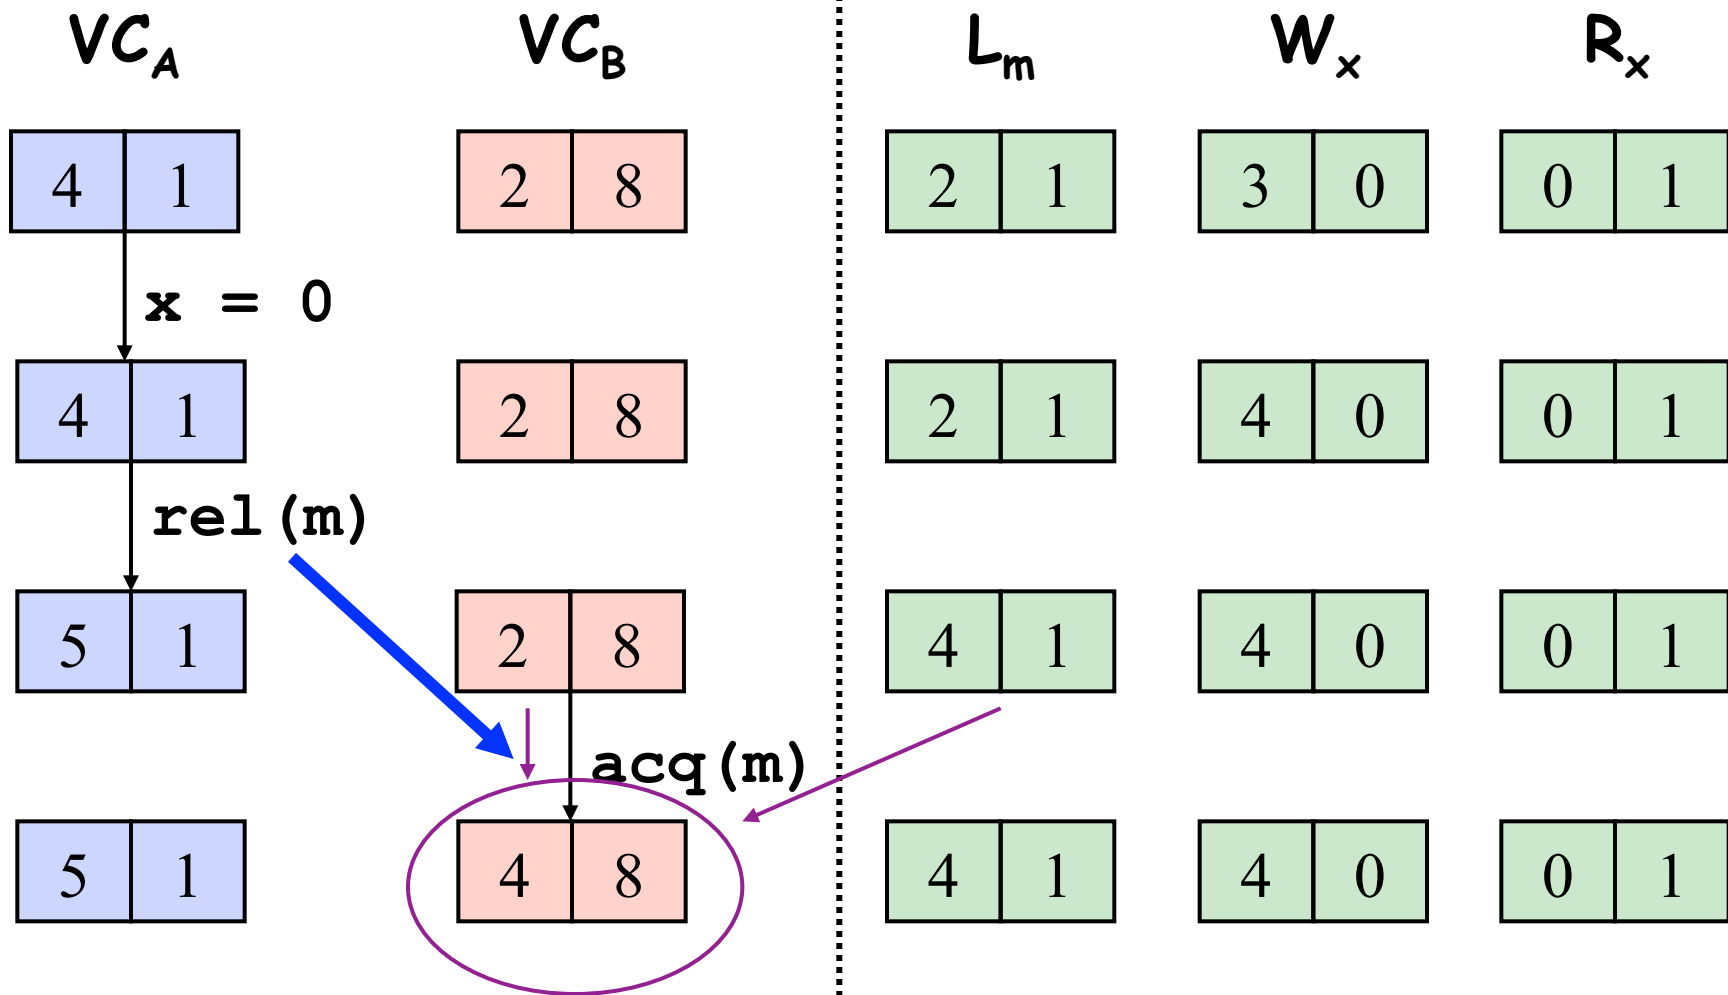
4	1
---	---

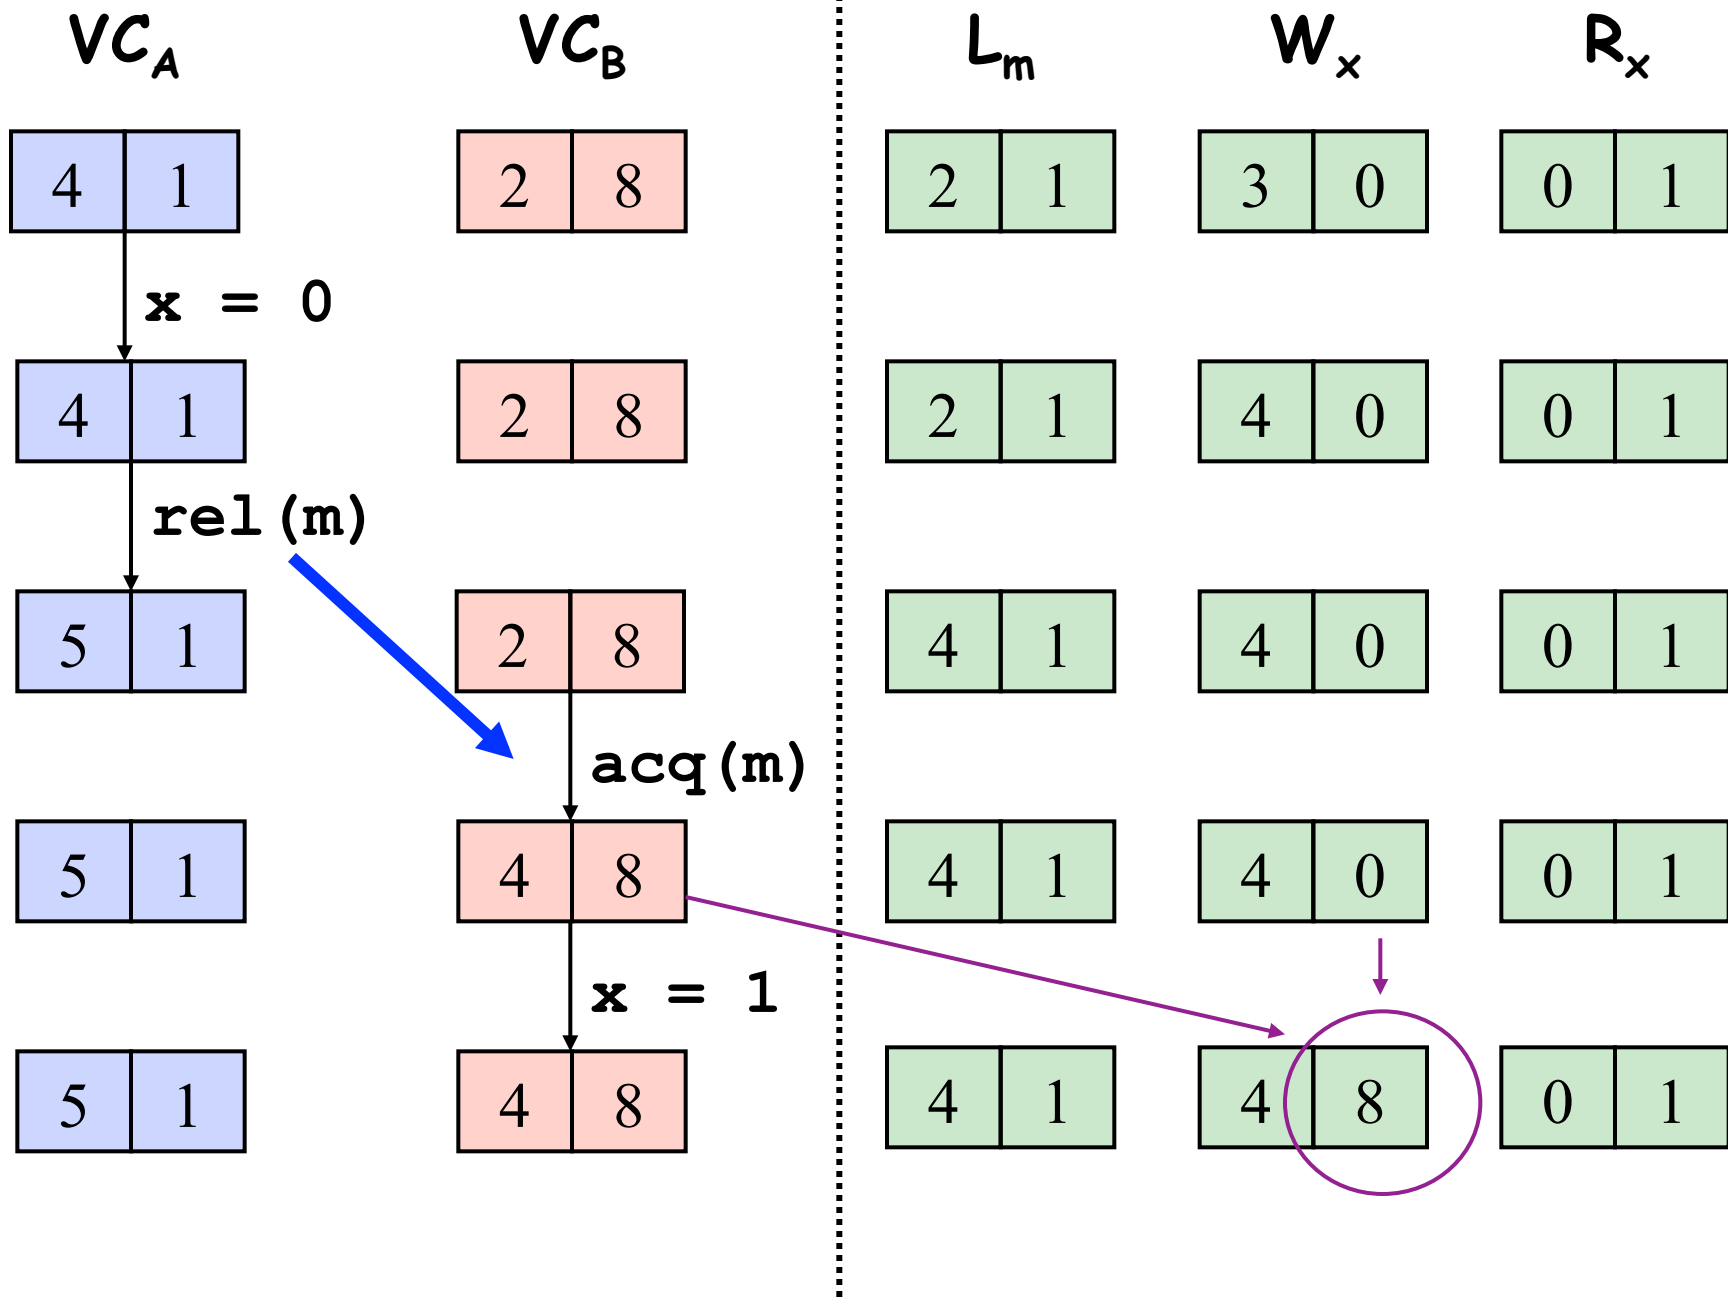
? **Yes**

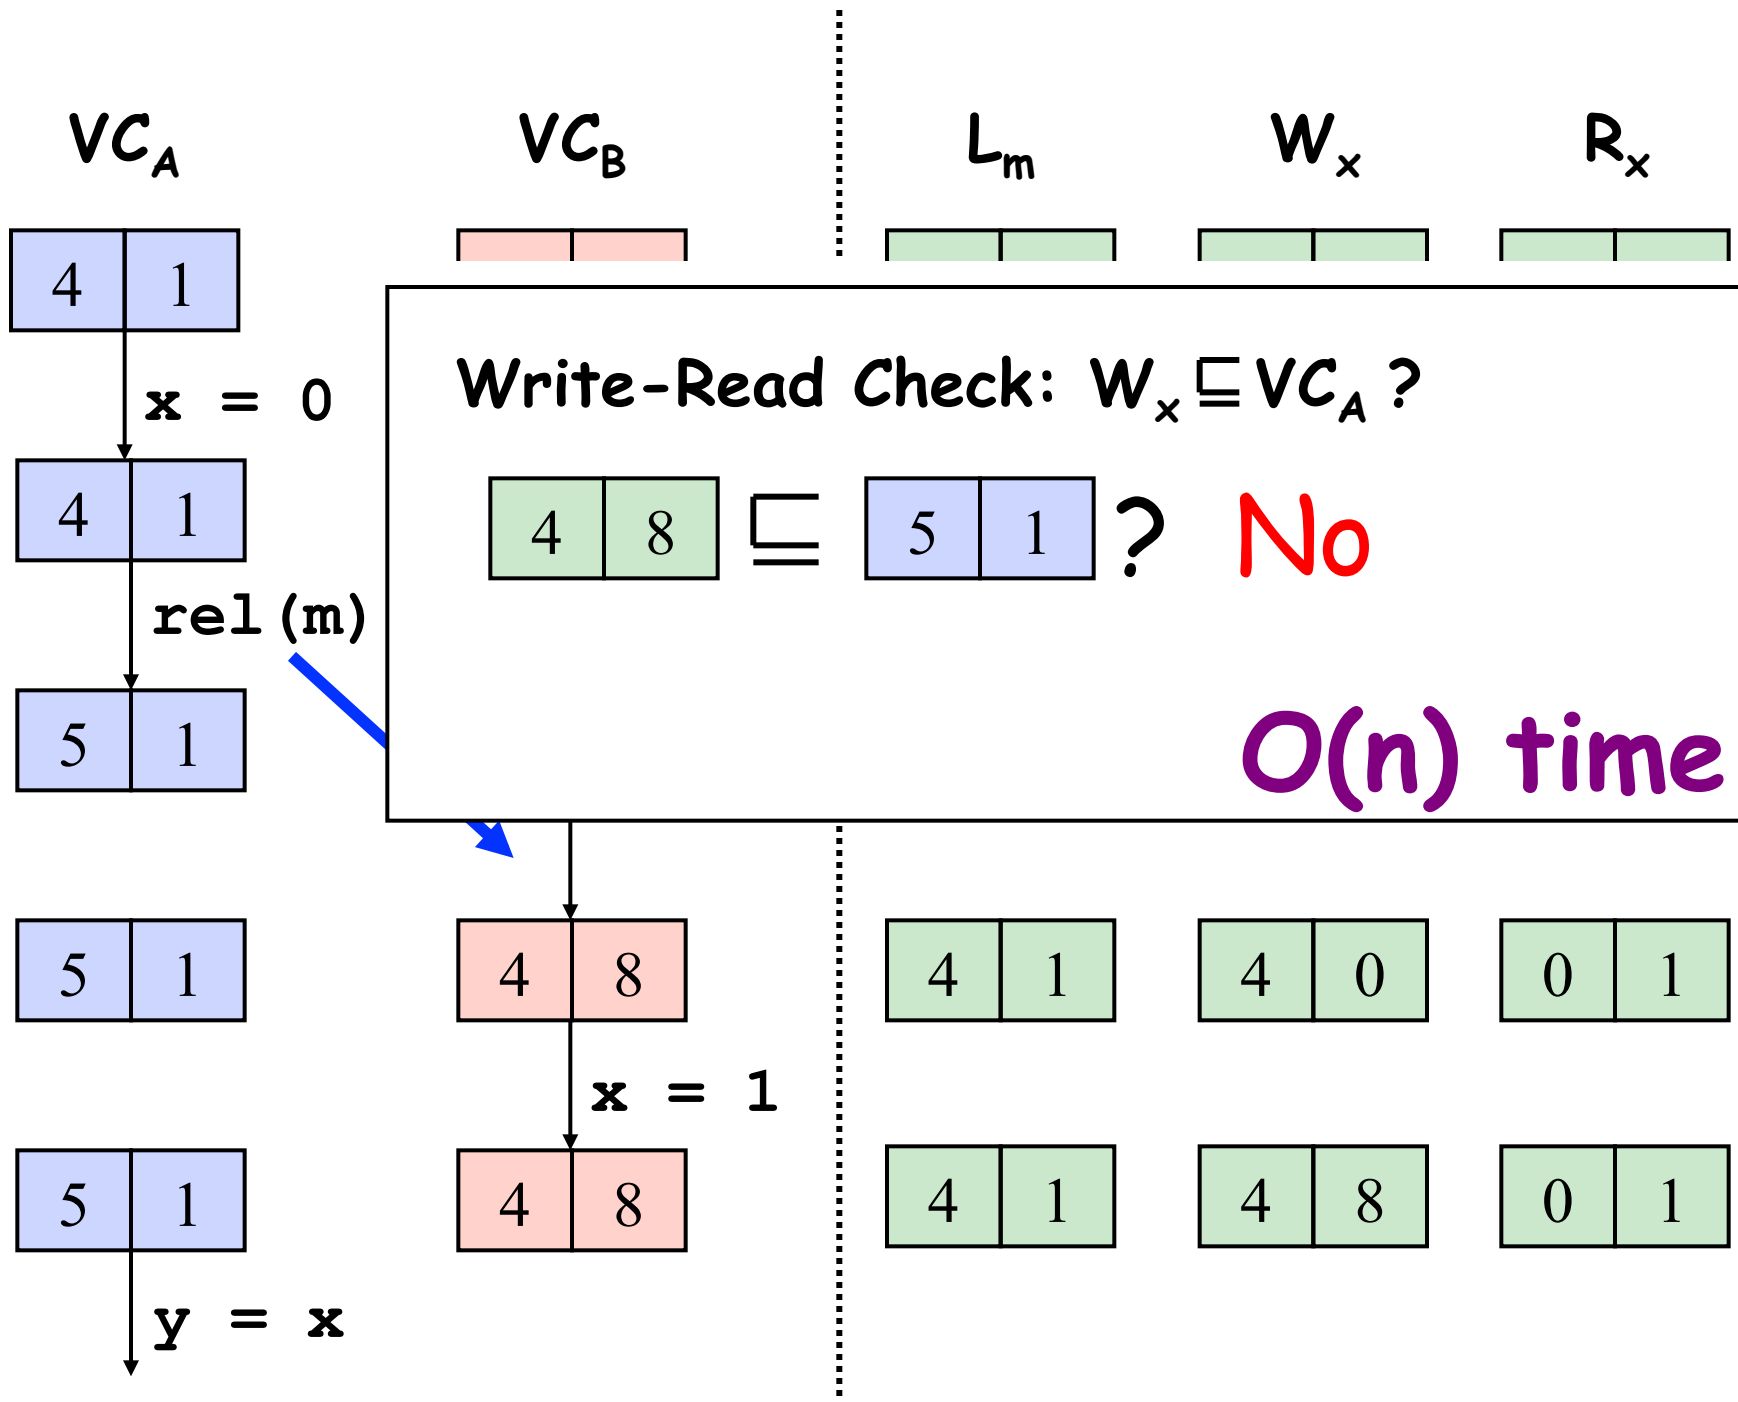
$O(n)$ time












Write-Write and Write-Read Races

Thread A

Thread B

Thread C

Thread D




No Races Yet: Writes Totally Ordered!

Thread A

Thread B

Thread C

Thread D




No Races Yet: Writes Totally Ordered!





Thread A

Thread B

Thread C

Thread D





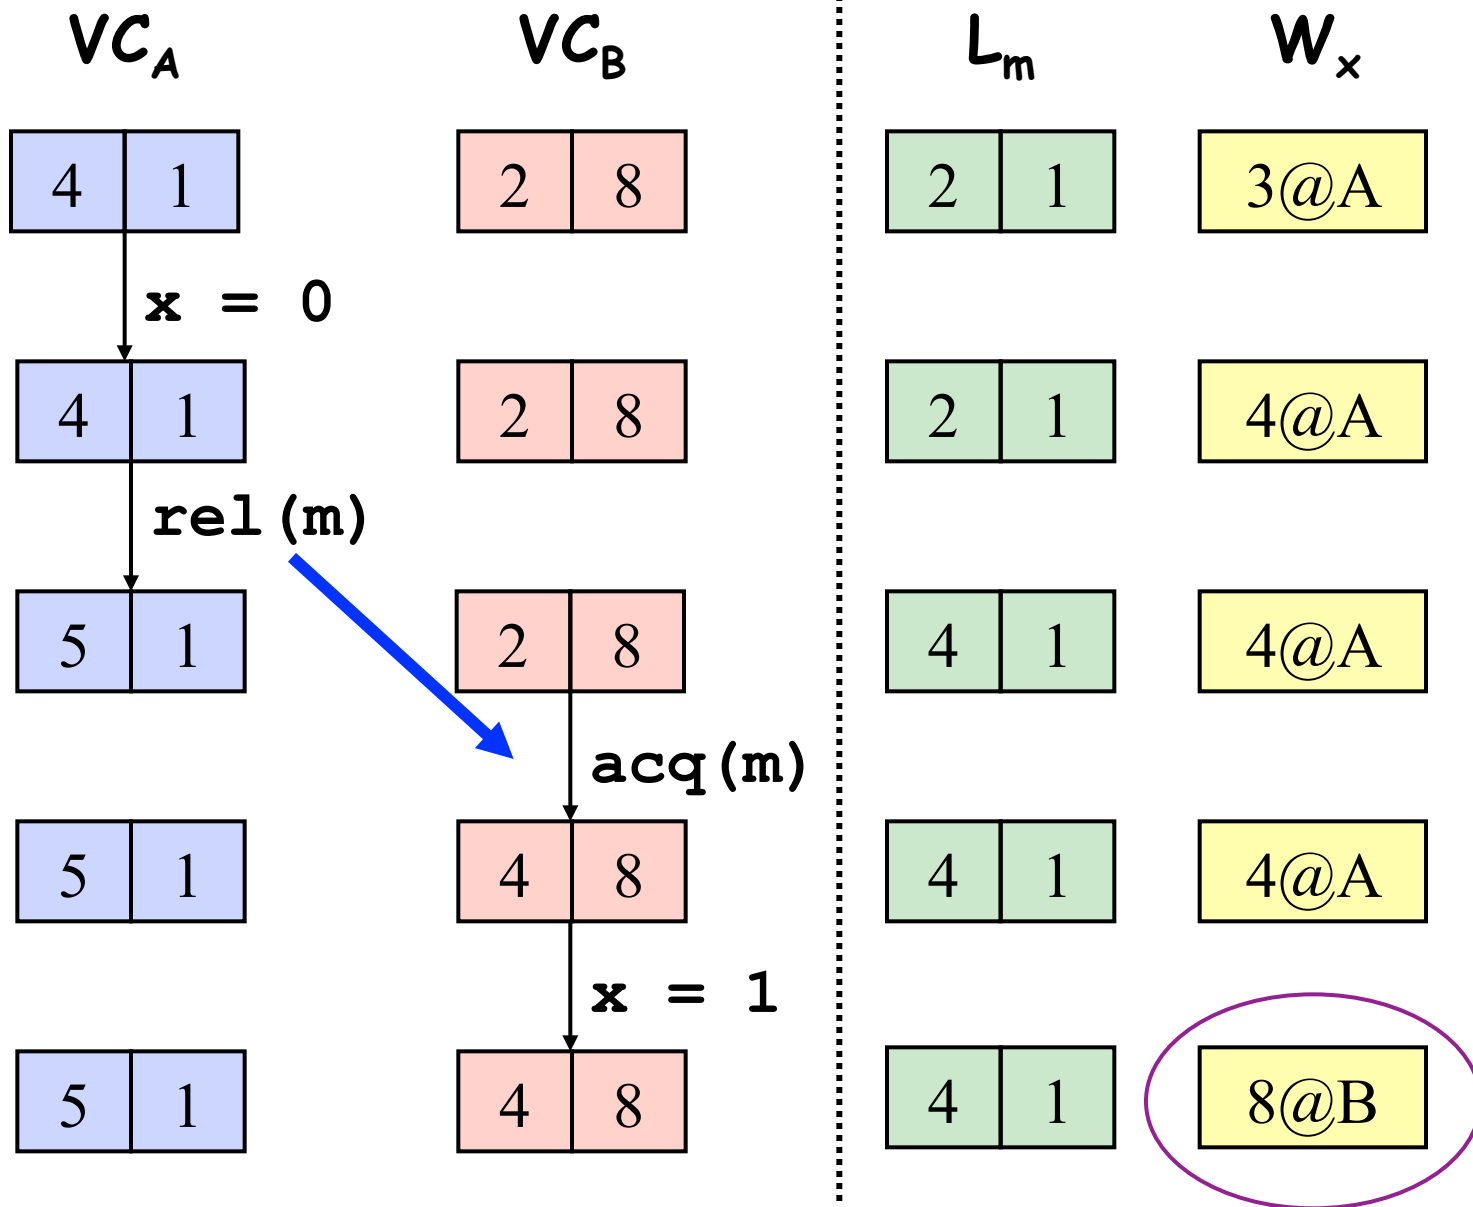
Last Write
"Epoch"

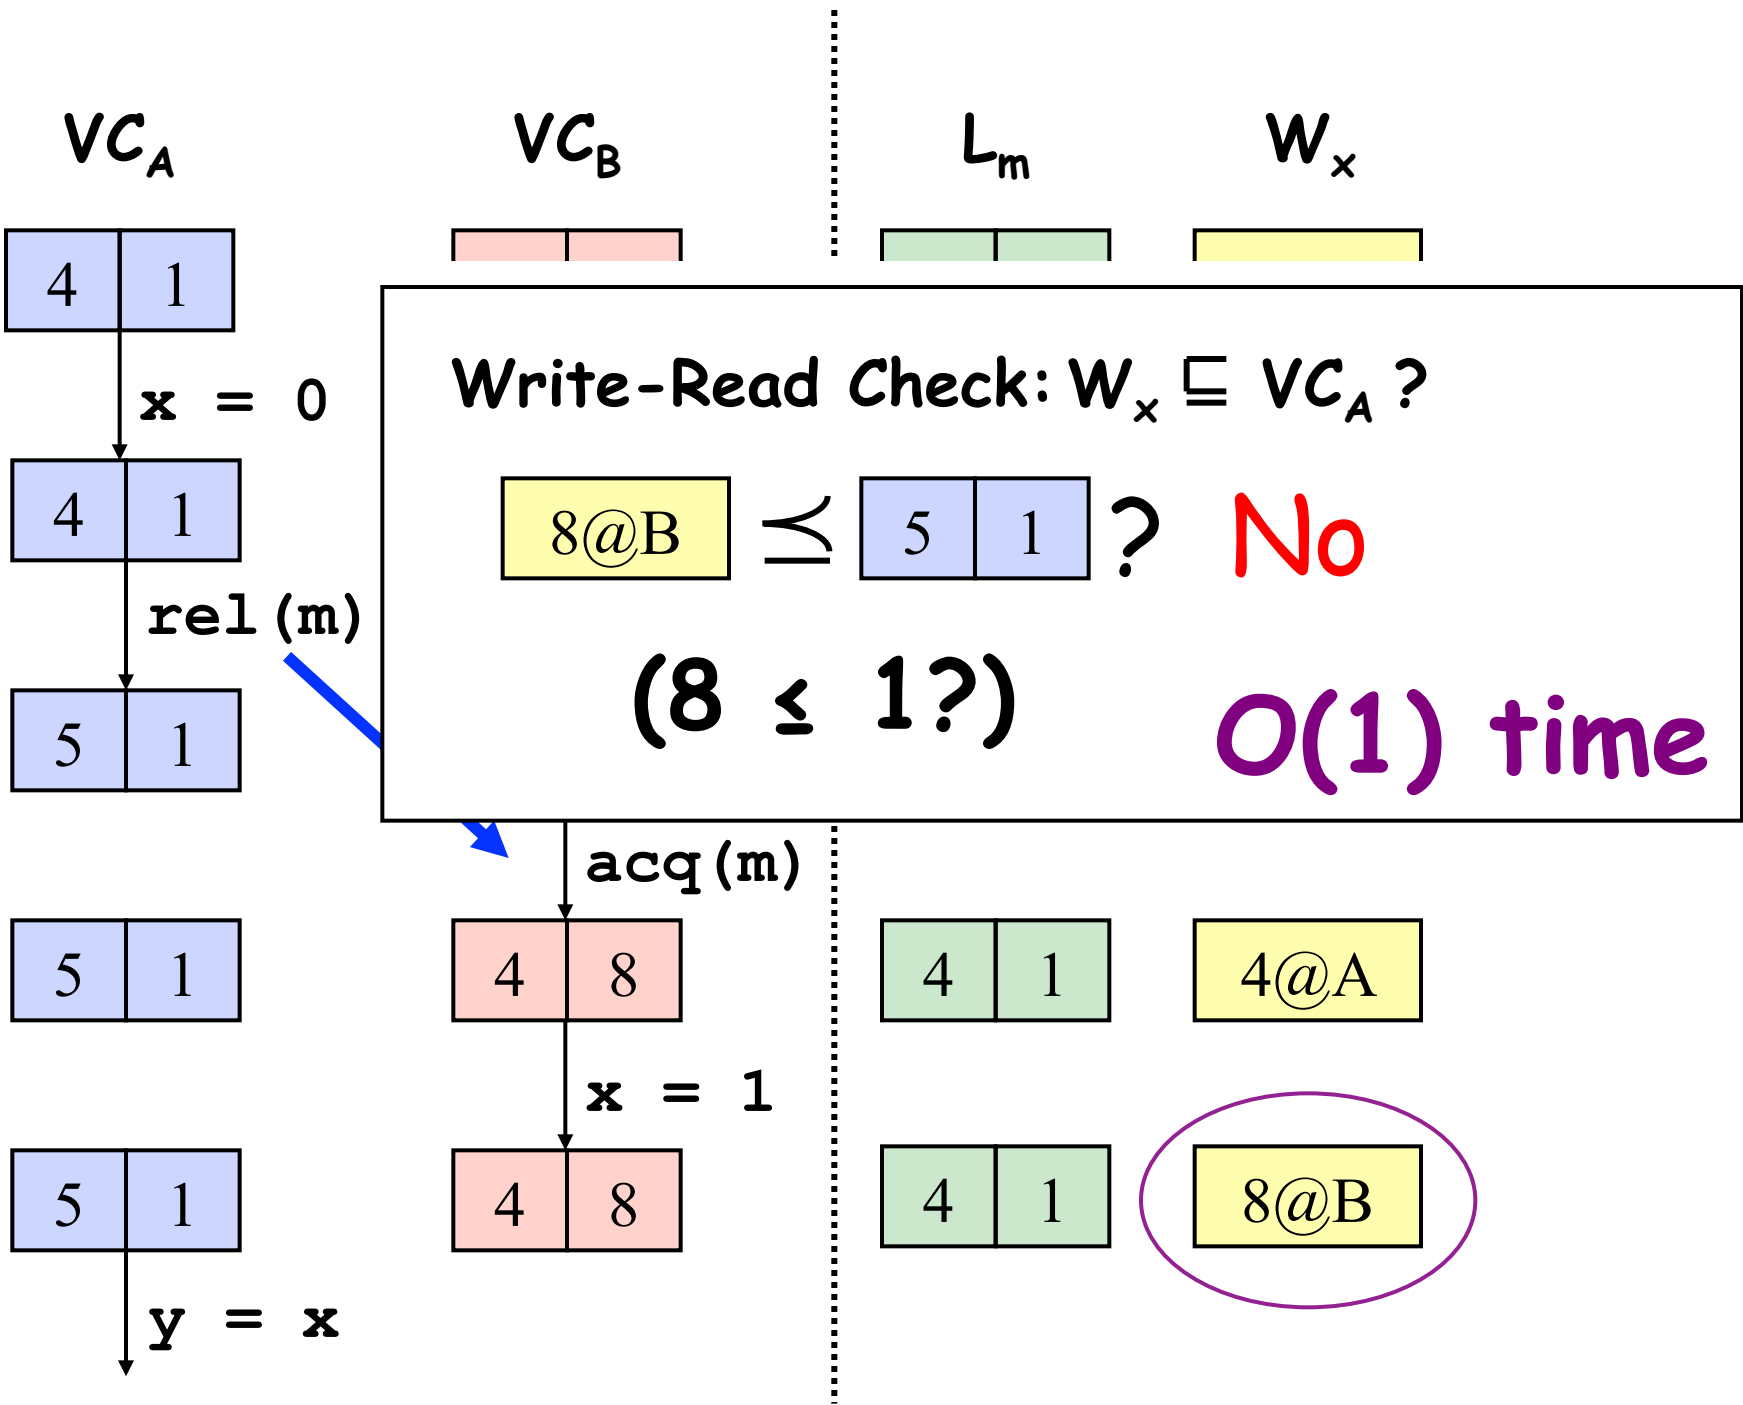
Write-Write Check: $W_x \sqsubseteq VC_A$?

1@B	⊆	<table border="1" style="border-collapse: collapse; text-align: center;"> <tr> <td style="width: 30px; height: 30px;">4</td> <td style="width: 30px; height: 30px;">1</td> </tr> </table>	4	1	?	Yes
4	1					

(1 ≤ 1?)

O(1) time





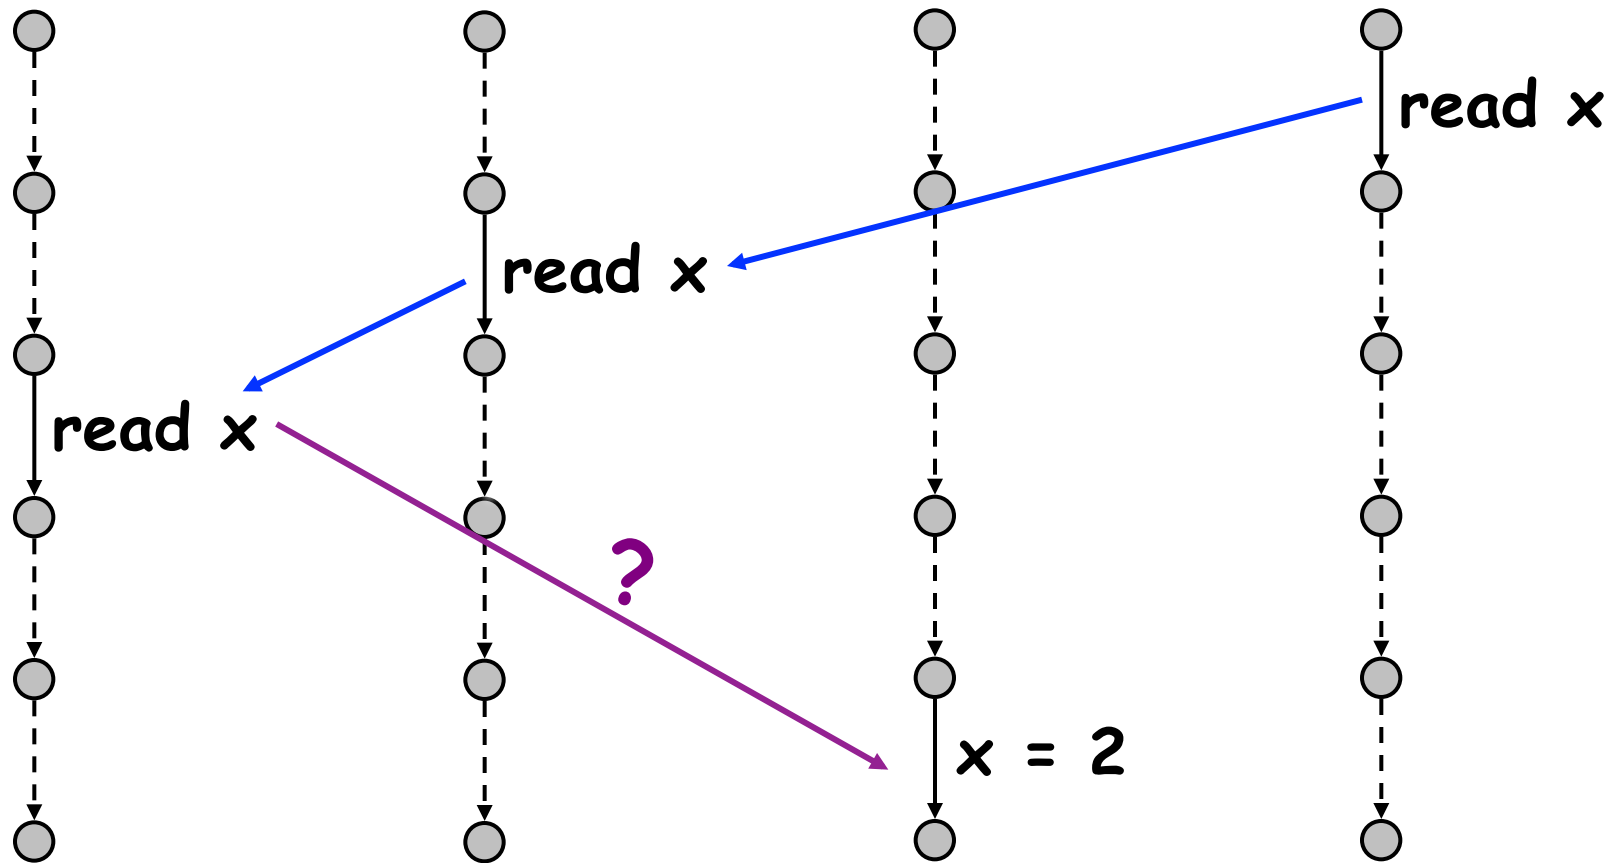
Read-Write Races -- Ordered Reads

Thread A

Thread B

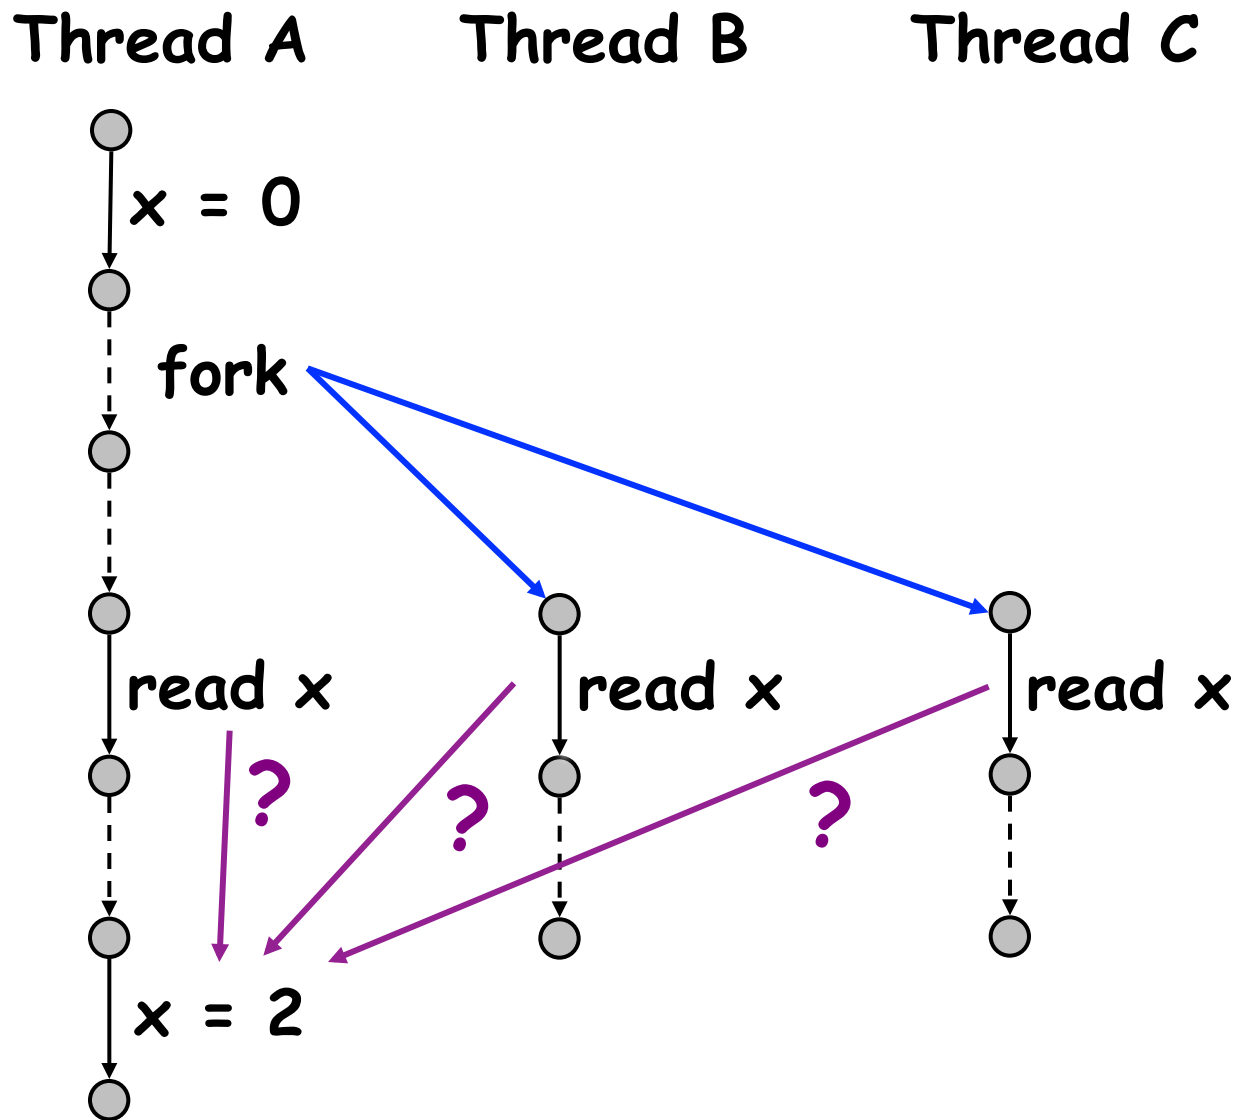
Thread C

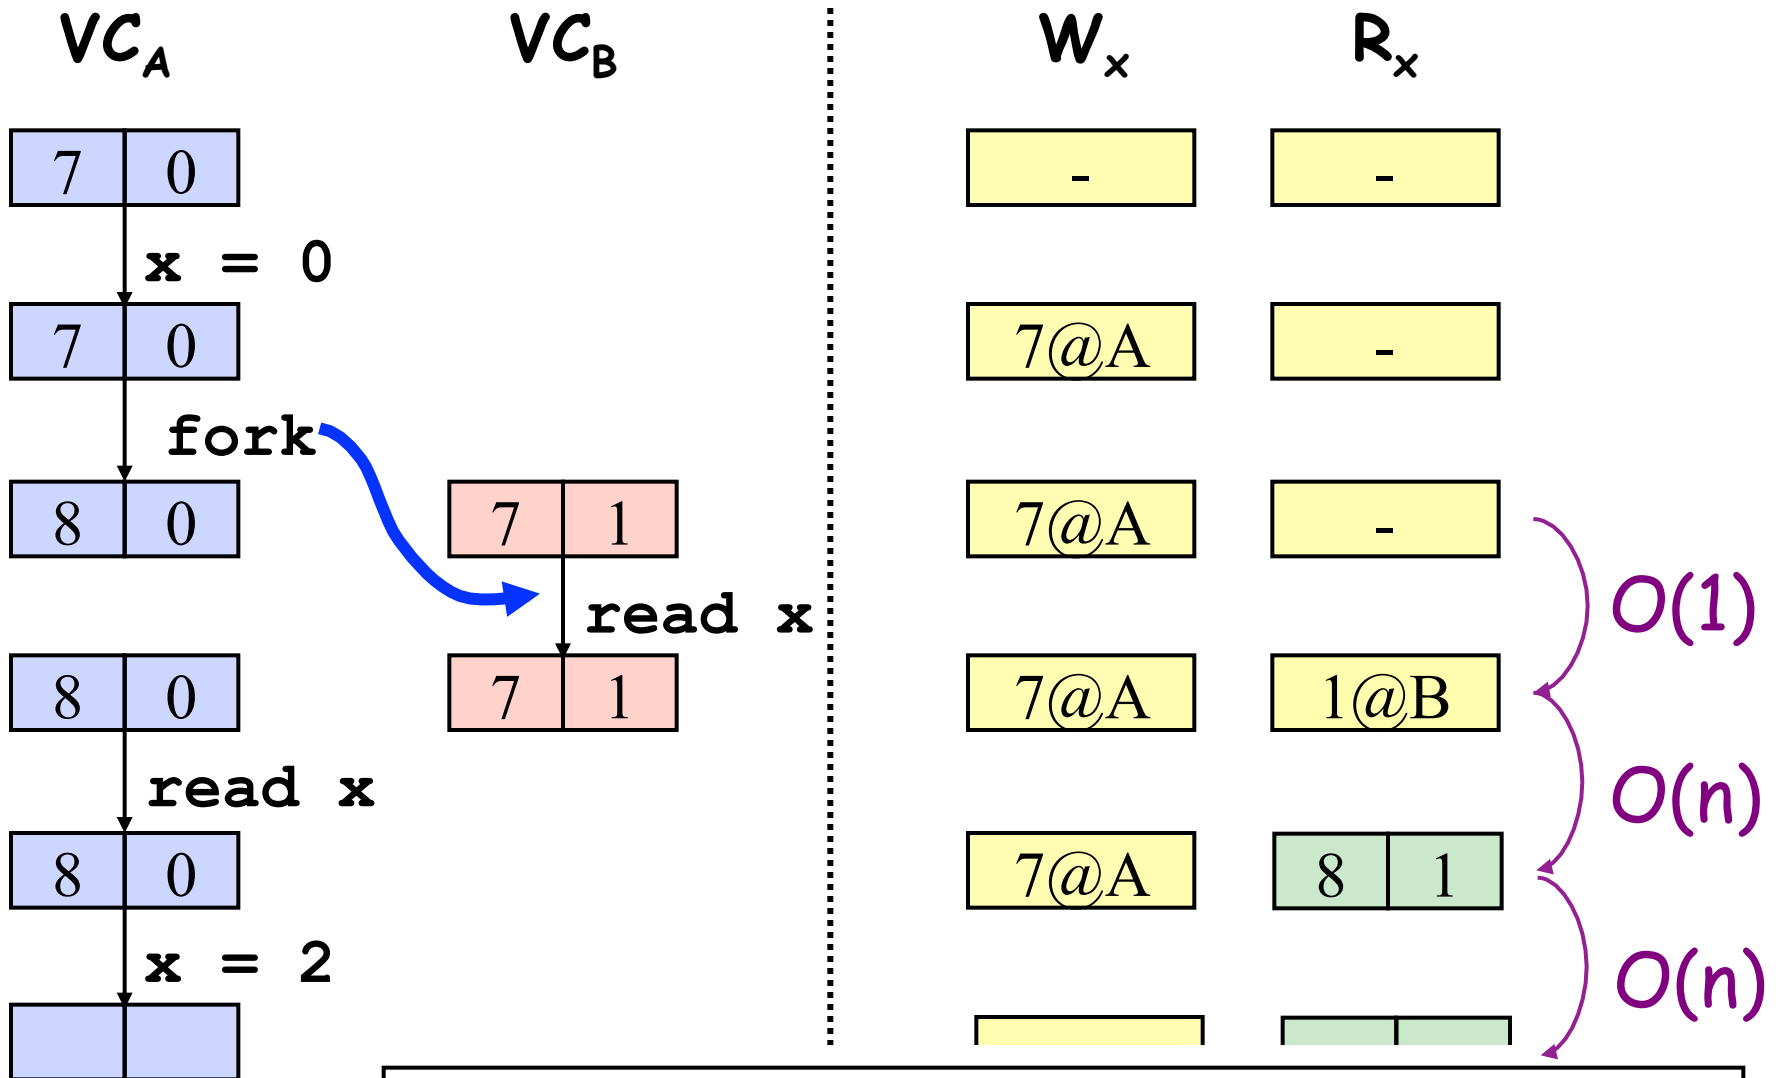
Thread D



Most common case: thread-local, lock-protected, ...

Read-Write Races -- Unordered Reads





Read-Write Check: $R_x \sqsubseteq VC_A$?


$\begin{bmatrix} 8 & 1 \end{bmatrix} \sqsubseteq \begin{bmatrix} 8 & 0 \end{bmatrix} ? **No**$

Thread A

Thread B

Thread C

Thread D




read x

?



read x

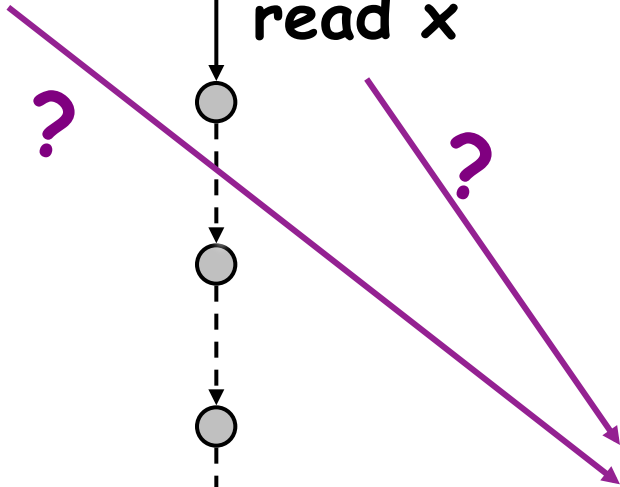
?



x = 2



$O(n)$




Thread A

Thread B

Thread C


Thread D



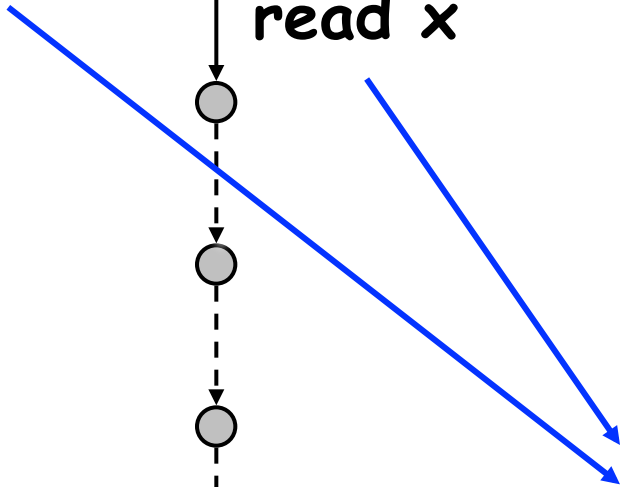
read x



read x



x = 2




Thread A


Thread B

Thread C


Thread D



read x




read x



x = 2



x = 3



$O(n)$

Thread A

Thread B

Thread C


Thread D



read x



read x





x = 2



x = 3

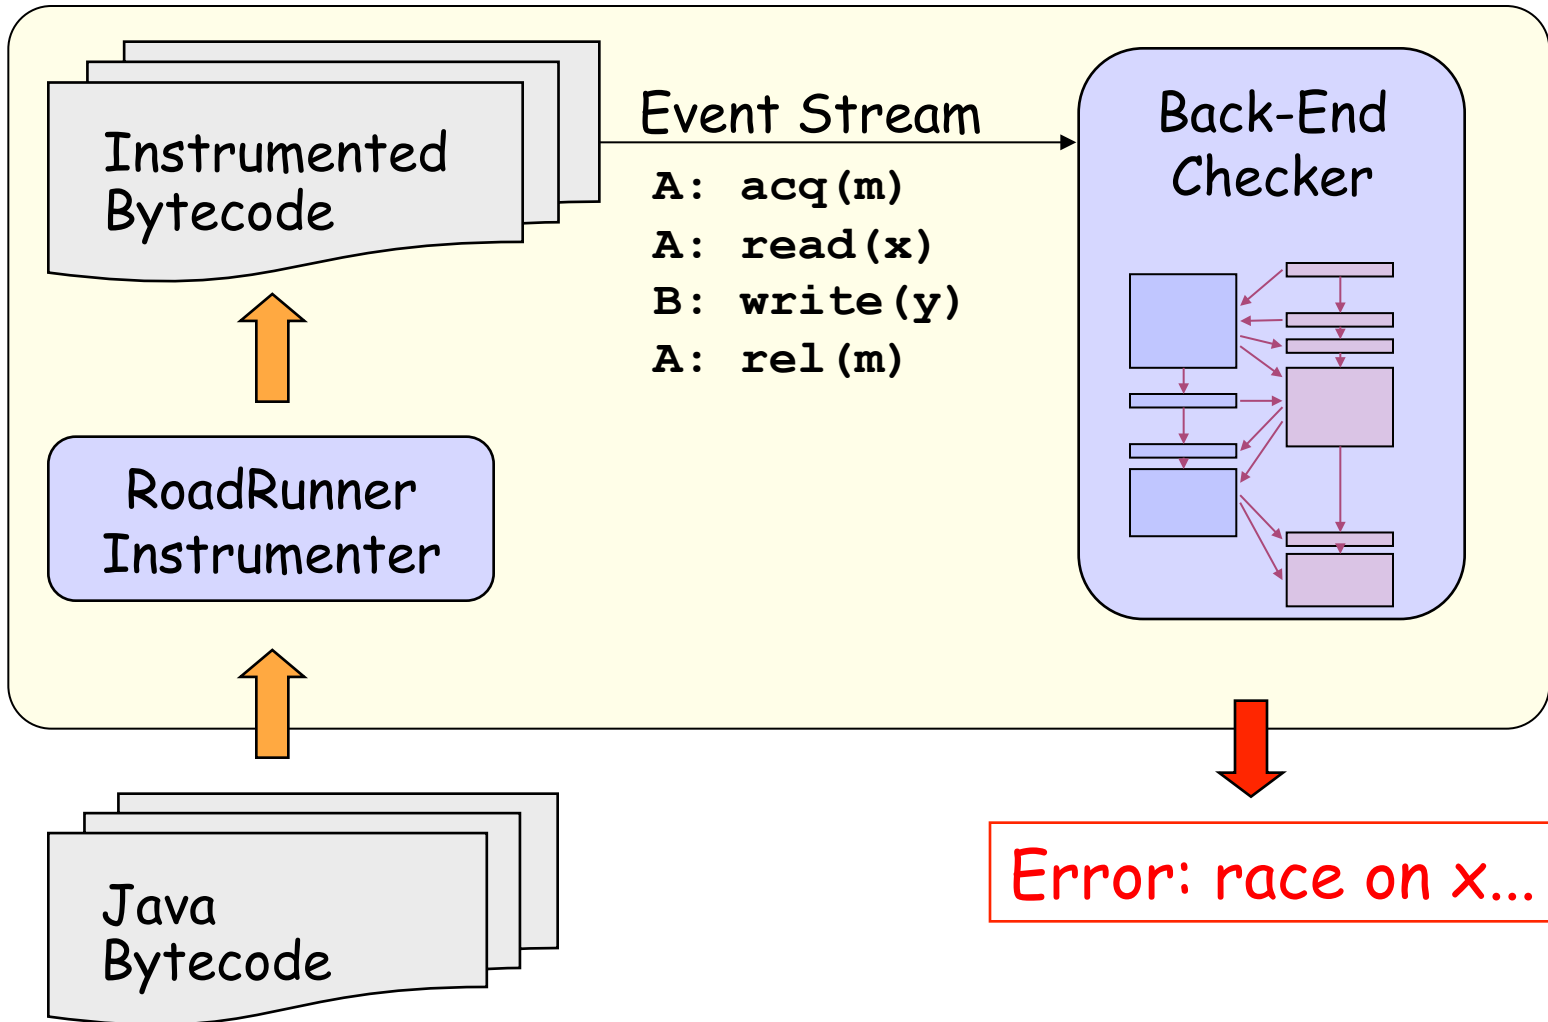
O(1)

Forget VC for R_x
and switch back
to "last read epoch"



RoadRunner Architecture

Standard JVM




Validation


- Six race condition checkers
 - all use RoadRunner
 - share common components (eg, VectorClock)
 - profiled and optimized
- Further optimization opportunities
 - unsound extensions, dynamic escape analysis, static analysis, implement inside JVM, hardware support, ...
- 15 Benchmarks
 - 250 KLOC
 - locks, wait/notify, fork/join, barriers, ...

Warnings


22 false positives
3 false negatives




Slowdown (x Base Time)



$O(n)$ Vector Clock Operations



$O(n)$ Vector Clock Operations



Memory Usage

- FastTrack allocated ~200x fewer VCs


Checker	Memory Overhead
Basic VC, DJIT+	7.9x
FastTrack	2.8x

(Note: VCs for dead objects can be garbage collected)

- Improvements
 - accordion clocks [CB 01]
 - analysis granularity [PS 03, YRC 05] (see paper)

Eclipse 3.4

- Scale
 - > 6,000 classes
 - 24 threads
 - custom sync. idioms
- Precision (tested 5 common tasks)
 - Eraser: ~1000 warnings
 - FastTrack: ~30 warnings
- Performance on compute-bound tasks
 - > 2x speed of other precise checkers
 - same as Eraser



Beyond Detecting Race Conditions

- FastTrack finds real race conditions
 - races correlated with defects
 - cause unintuitive behavior on relaxed memory
- Which race conditions are real bugs?
 - that cause erroneous behaviors (crashes, etc)
 - and are not "benign race conditions"

```
class Point {
    double x, y;
    static Point p;

    Point() { x = 1.0; y = 1.0; }

    static Point get() {
        Point t = p;
        if (t != null) return t;
        synchronized (Point.class) {
            if (p==null) p = new Point();
            return p;
        }
    }

    static double slope() {
        return get().x / get().y;
    }

    public static void main(String[] args) {
        fork { System.out.println( slope() ); }
        fork { System.out.println( slope() ); }
    }
}
```

Thread 0

```
p = null  
px = 0  
py = 0  
fork 1,2
```

Thread 1

```
read p // non-null  
read px // ?
```

Thread 2

```
read p // null  
acquire  
read p // null  
p = new Point  
px = 1  
py = 1  
release  
read px // get 1  
read py // get 1
```


Thread 0





```
p = null  
px = 0  
py = 0  
fork 1,2
```

Thread 1

```
read p // non-null  
read px // ?
```

Thread 2

```
read p // null  
acquire  
read p // null  
p = new Point  
px = 1  
py = 1  
release  
read px // get 1  
read py // get 1
```



Thread 0

```
p = null  
px = 0  
py = 0  
fork 1,2
```

Thread 1

```
read p // non-null  
read px // ?
```

Thread 2

```
read p // null  
acquire  
read p // null  
p = new Point  
px = 1  
py = 1  
release  
read px // get 1  
read py // get 1
```

- Race: can return either write (mm non-determinism)
- Typical JVM: mostly sequentially consistent
- Adversarial memory
 - use heuristics to return older stale values

Adversarial Memory

- Record history of all writes (plus VCs) to racy variables
- At read
 - determine all visible writes legal under JMM
 - heuristically pick one likely to crash target program
- Six heuristics:
 - Sequentially consistent: return last write
 - Oldest: return "most stale" value
 - Oldest-but-different: never return same val twice
 - if (p != null) p.draw()
 - Random, Random-but-different

Experimental Results

Program	Field	Erroneous Behavior Observation Rate (%)						Destructive Race?
		No Jumble	JUMBLE configurations				Random but Different	
			Sequentially Consistent	Oldest	Oldest but Different	Random		
Figure 1	x	0	0	0	0	28	57	Yes
Figure 2	p	0	0	0	0	0	0	No
	p.x	0	0	60	52	32	30	Yes
	p.y	0	0	48	53	27	30	Yes
hedc	Task.thread	0	0	0	96	24	43	Yes
	MetaSearchResult.results	0	0	100	100	100	100	Yes
	MetaSearchResult.completed	0	0	33	36	25	26	Yes
	MetaSearchResult.request	0	0	0	0	0	0	No
	Task.valid	0	0	0	0	0	0	No
jbb	Company.elapsed_time	0	0	100	0	15	5	Yes
	Company.mode	0	0	100	100	95	98	Yes
montecarlo	Universal.UNIVERSAL_DEBUG	0	0	0	0	0	0	No
mtrt	RayTracer.threadCount	0	0	0	0	0	0	No
raytracer	JGFRayTracerBench.checksum1	0	0	100	100	100	100	Yes
tsp	TspSolver.MinTourLen	0	0	100	100	100	100	QoS
sor	array index [0] and [1]	0	0	100	100	100	100	Yes
lufact	array index [0] and [1]	0	0	100	100	100	100	Yes
moldyn	array index [0] and [1]	0	0	100	100	100	100	Yes