

---

# On-Demand Dynamic Software Analysis

---

Joseph L. Greathouse

Ph.D. Candidate

Advanced Computer Architecture Laboratory

University of Michigan

November 29, 2011

---

# Software Errors Abound

- NIST: SW errors cost U.S. ~\$60 billion/year as of 2002

A problem has been detected and windows has been shut down to prevent damage to your computer.

The problem seems to be caused by the following file: SPCMDCON.SYS

PAGE\_FAULT\_IN\_NONPAGED\_AREA

If this is the first time you've seen this stop error screen, restart your computer. If this screen appears again, follow these steps:

Check to make sure any new hardware or software is properly installed. If this is a new installation, ask your hardware or software manufacturer for any windows updates you might need.

If problems continue, disable or remove any newly installed hardware or software. Disable BIOS memory options such as caching or shadowing. If you need to use Safe Mode to remove or disable components, restart your computer, press F8 to select Advanced Startup Options, and then select Safe Mode.

Technical information:

\*\*\* STOP: 0x00000050 (0xFD3094C2,0x00000001,0xFBFE7617,0x00000000)

\*\*\* SPCMDCON.SYS - Address FBFE7617 base at FBFE5000, Datestamp 3d6dd67c

# Software Errors Abound

- NIST: SW errors cost U.S. ~\$60 billion/year as of 2002

```
A problem has been detected and windows has been shut down to prevent damage
to your computer.

The problem seems to be caused by the following file: SPCMDCON.SYS

PAGE_FAULT_IN_NONPAGED_AREA

If this is the first time you've seen this stop error screen,
restart your computer. If this screen appears again, follow
these steps:

Check to make sure any new hardware or software is properly installed.
If this is a new installation, ask your hardware or software manufacturer
for any windows updates you might need.

If problems continue, disable or remove any newly installed hardware
or software. Disable BIOS memory options such as caching or shadowing.
If you need to use Safe Mode to remove or disable components, restart
your computer, press F8 to select Advanced Startup Options, and then
select Safe Mode.

Technical information:

*** STOP: 0x00000050 (0xFD3094C2,0x00000001,0xFBFE7617,0x00000000)

*** SPCMDCON.SYS - Address FBFE7617 base at FBFE5000, DateStamp 3d6dd67c
```

# Software Errors Abound

- NIST: SW errors cost U.S. ~\$60 billion/year as of 2002
- FBI CCS: Security Issues \$67 billion/year as of 2005
  - >1/3 from viruses, network intrusion, etc.

```
A problem has been detected and windows has been shut down to prevent damage
to your computer.

The problem seems to be caused by the following file: SPCMDCON.SYS

PAGE_FAULT_IN_NONPAGED_AREA

If this is the first time you've seen this stop error screen,
restart your computer. If this screen appears again, follow
these steps:

Check to make sure any new hardware or software is properly installed.
If this is a new installation, ask your hardware or software manufacturer
for any windows updates you might need.

If problems continue, disable or remove any newly installed hardware
or software. Disable BIOS memory options such as caching or shadowing.
If you need to use Safe Mode to remove or disable components, restart
your computer, press F8 to select Advanced Startup Options, and then
select Safe Mode.

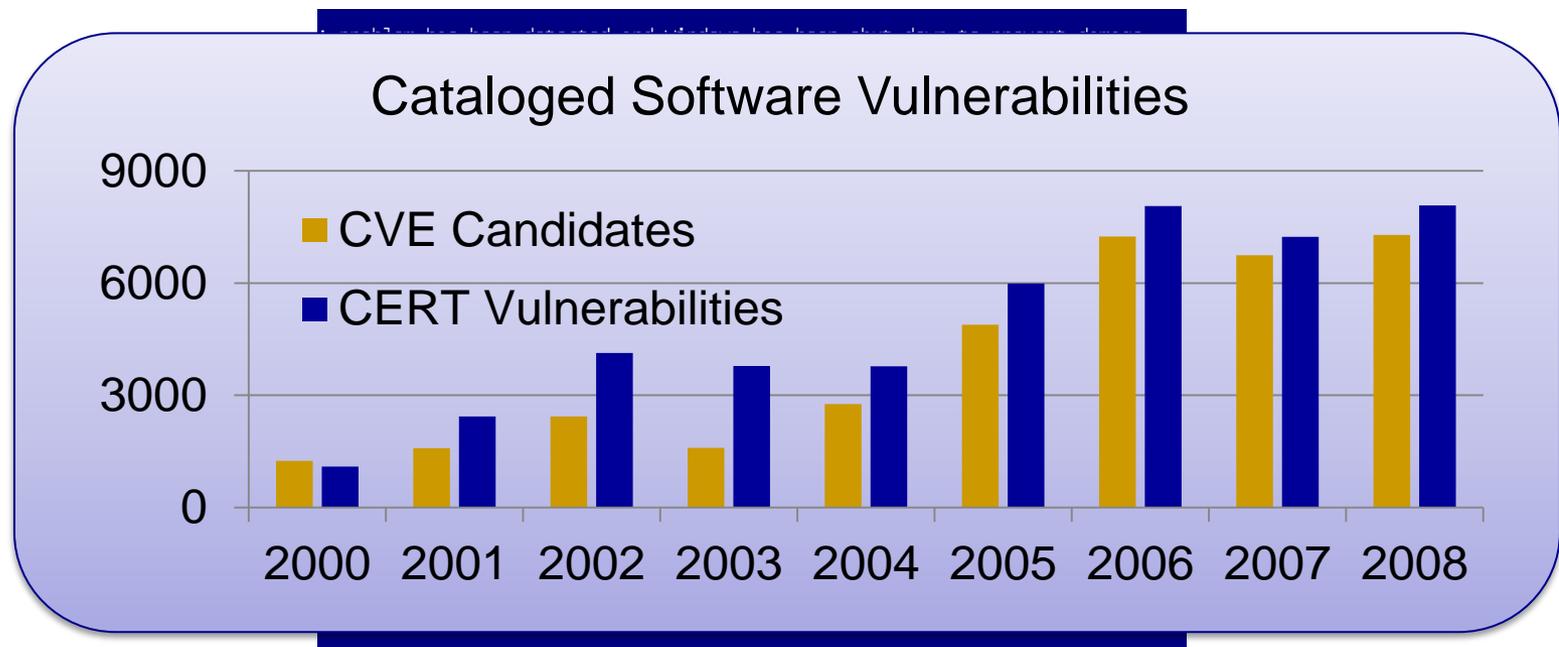
Technical information:

*** STOP: 0x00000050 (0xFD3094C2,0x00000001,0xFBFE7617,0x00000000)

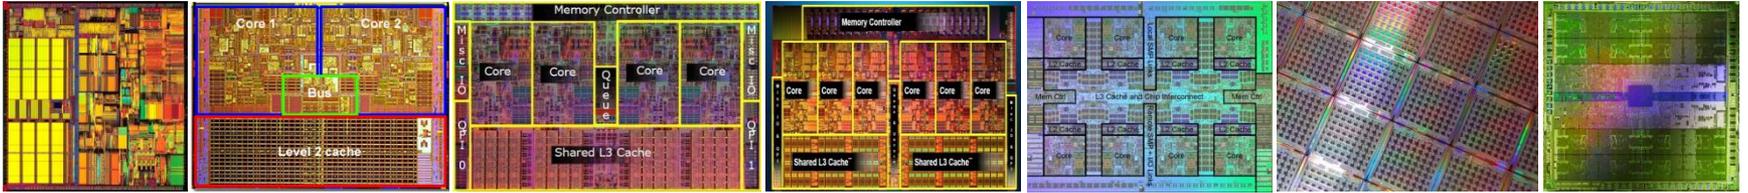
*** SPCMDCON.SYS - Address FBFE7617 base at FBFE5000, DateStamp 3d6dd67c
```

# Software Errors Abound

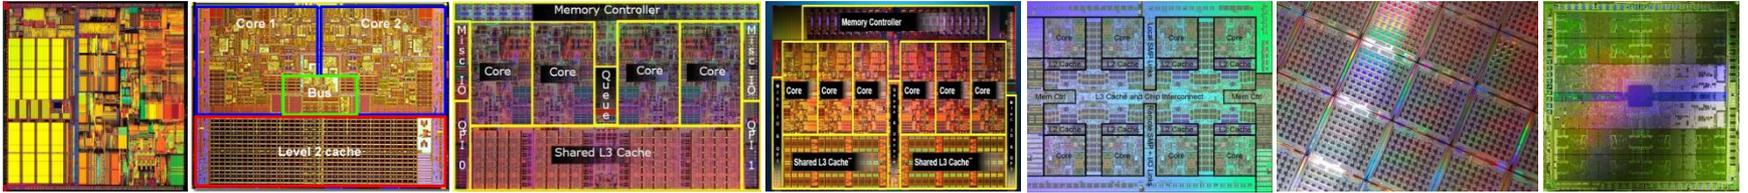
- NIST: SW errors cost U.S. ~\$60 billion/year as of 2002
- FBI CCS: Security Issues \$67 billion/year as of 2005
  - $>1/3$  from viruses, network intrusion, etc.



# Hardware Plays a Role



# Hardware Plays a Role



In spite of proposed solutions

Hardware Data  
Race Recording

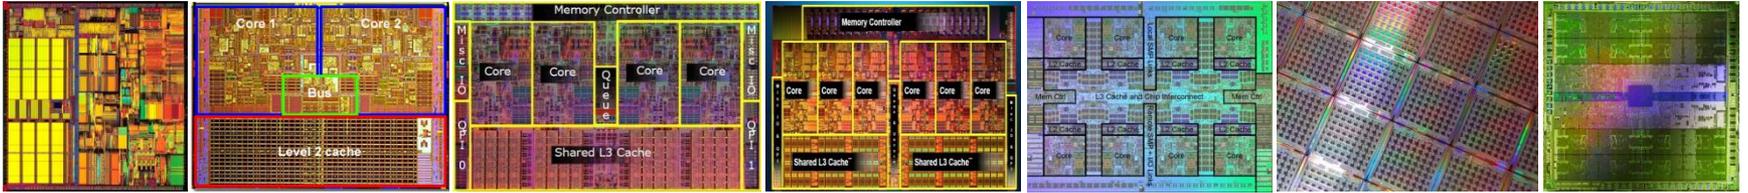
Bulk Memory  
Commits

Deterministic  
Execution/Replay

Bug-Free  
Memory Models

Atomicity Violation  
Detectors

# Hardware Plays a Role



In spite of proposed solutions

Hardware Data  
Race Recording

Bulk Memory  
Commits

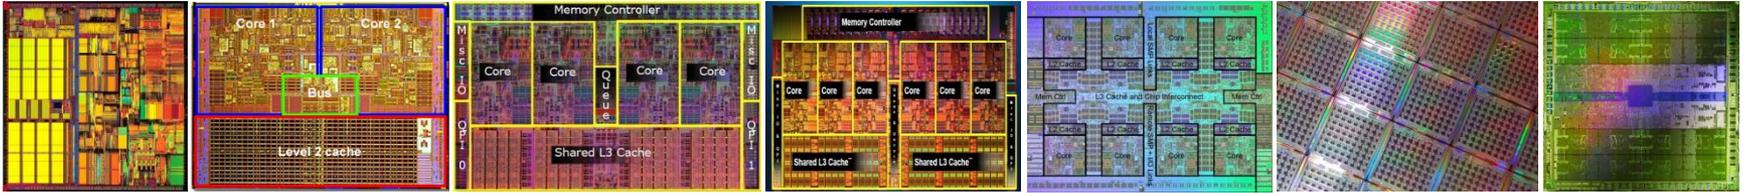
Deterministic  
Execution/Replay

Bulk  
Memory

**TRANSACTIONAL  
MEMORY**

Violation  
Detectors

# Hardware Plays a Role



In spite of proposed solutions

Hardware Data  
Race Recording

Bulk Memory  
Commits

Deterministic  
Execution/Replay

Bulk  
Memory



**TRANSACTIONAL  
MEMORY**

AMD  
ASF  
?

IBM  
BG/Q

Violation  
Detectors

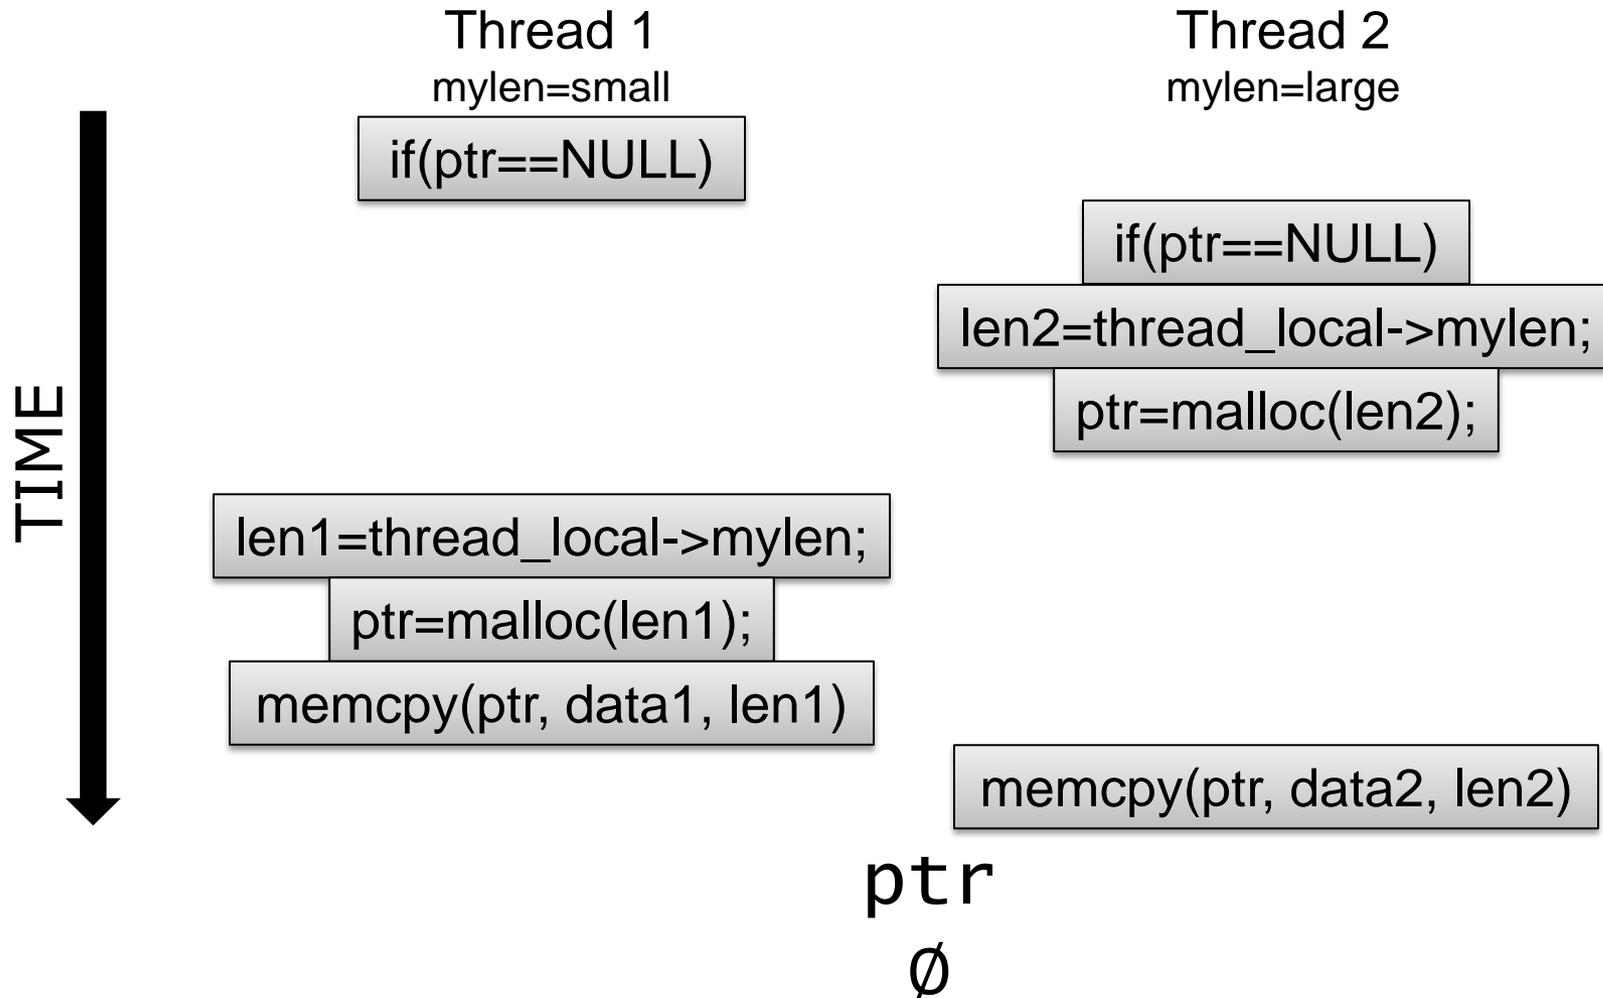


---

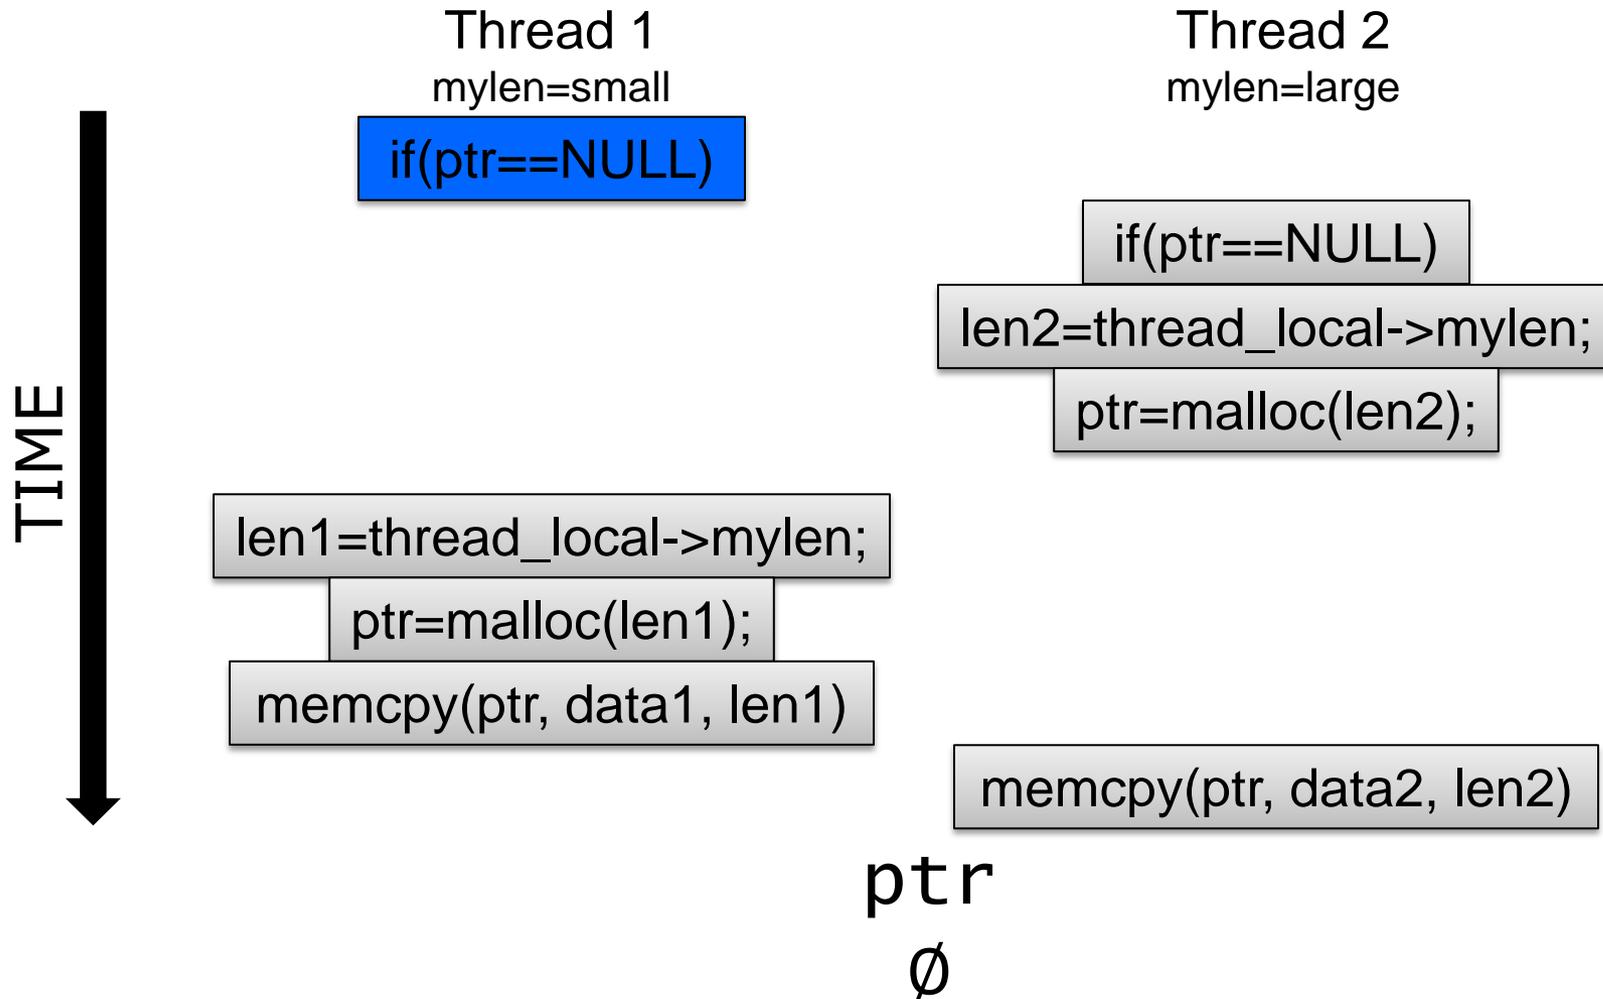
# Example of Modern Bug

Nov. 2010 OpenSSL Security Flaw

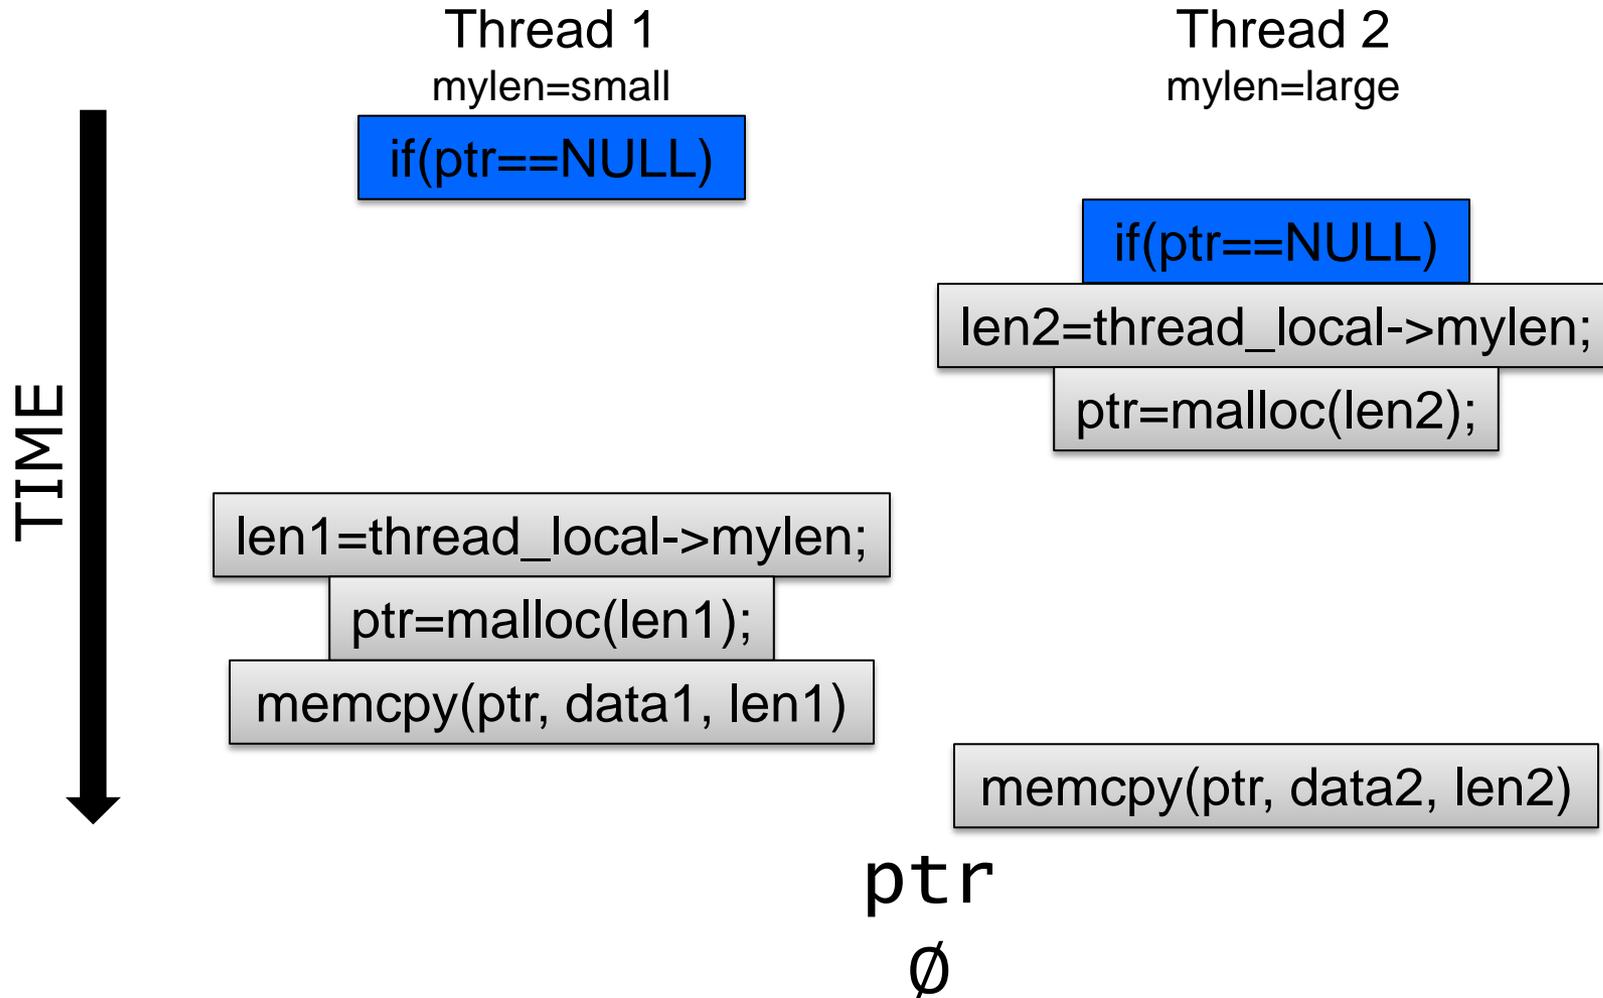
# Example of Modern Bug



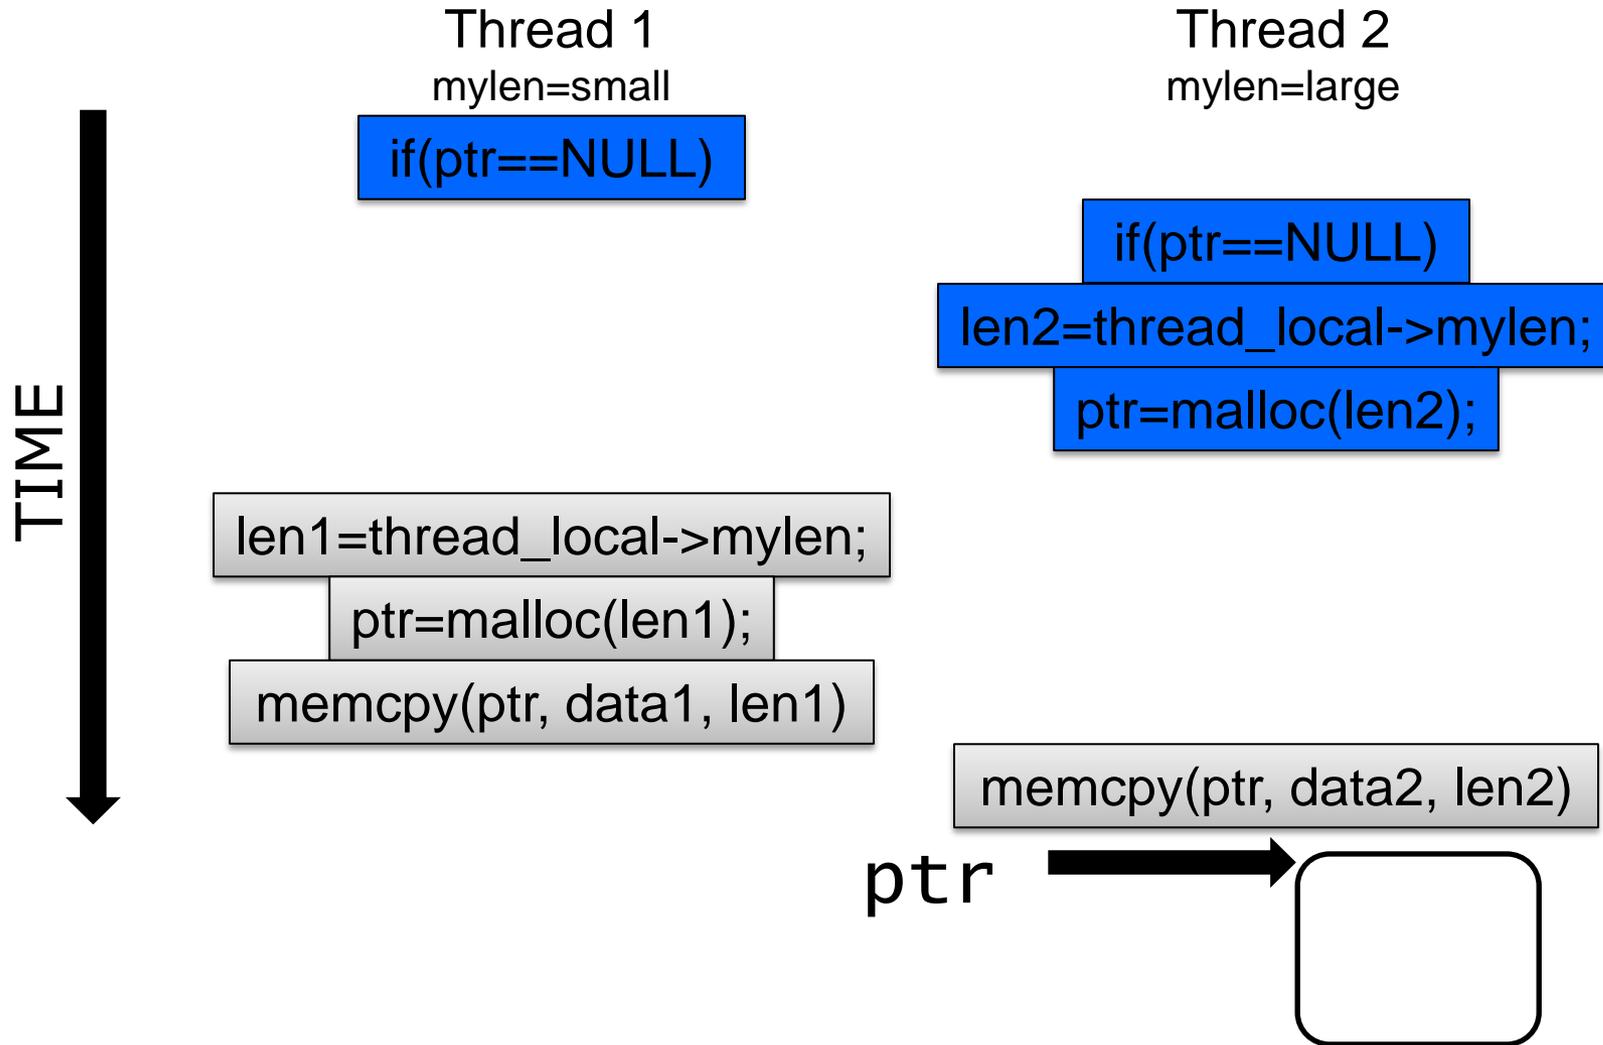
# Example of Modern Bug



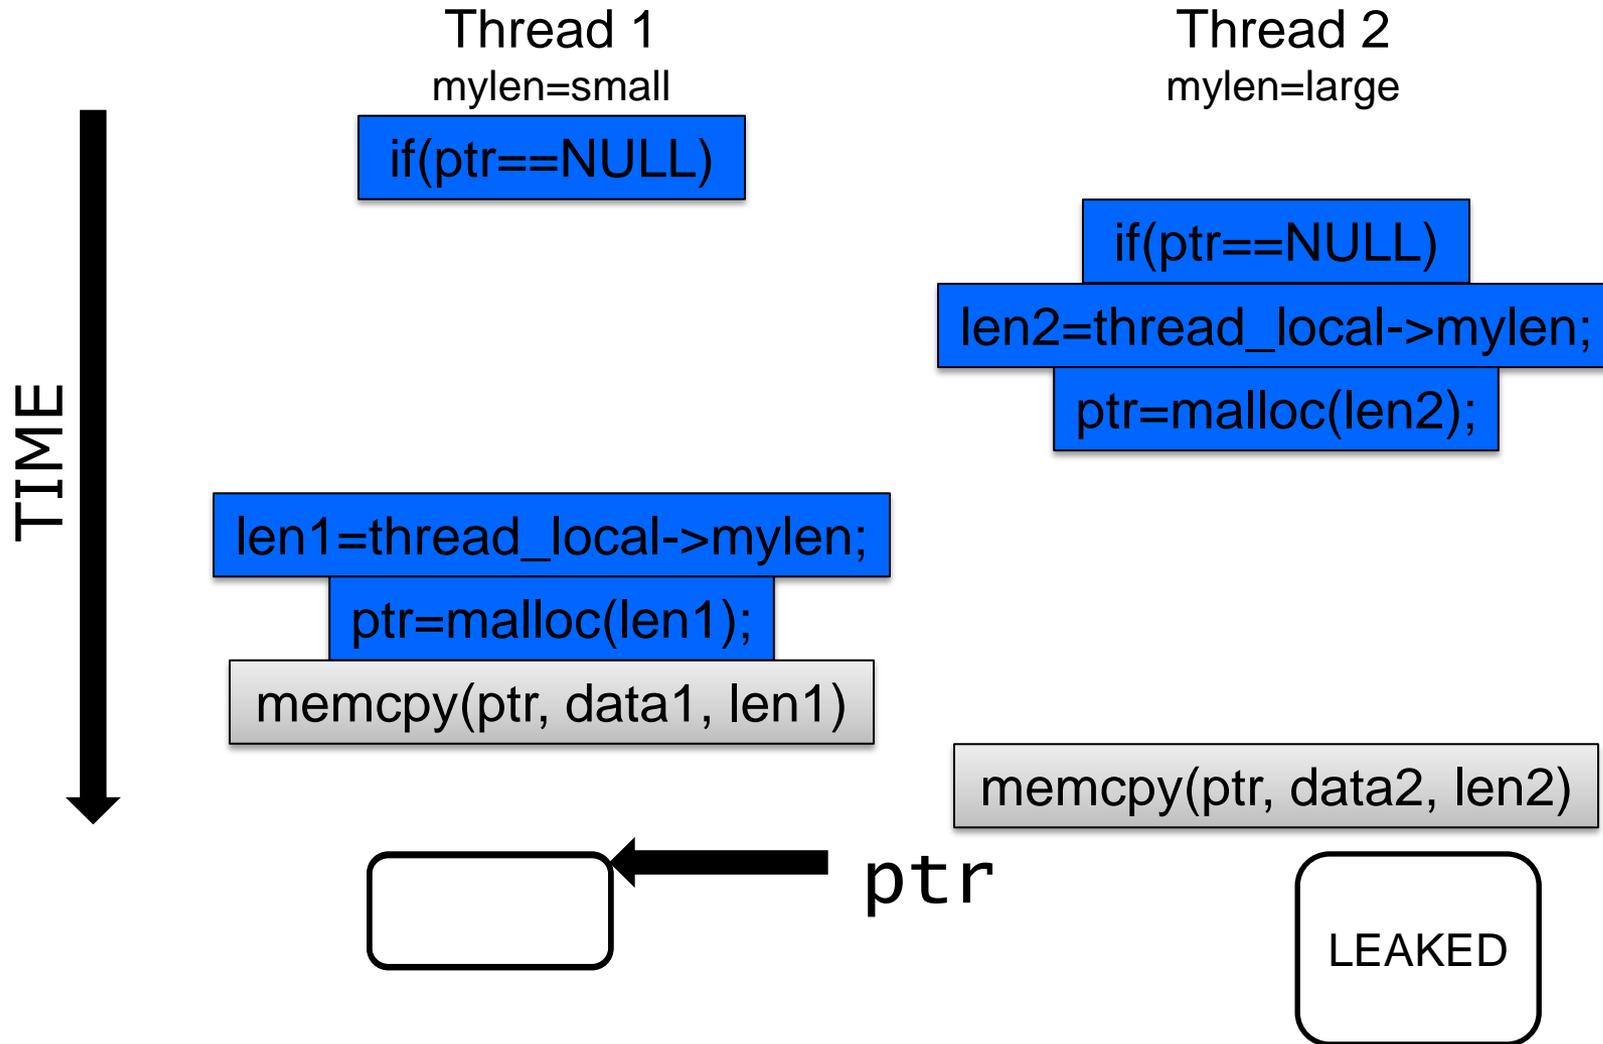
# Example of Modern Bug



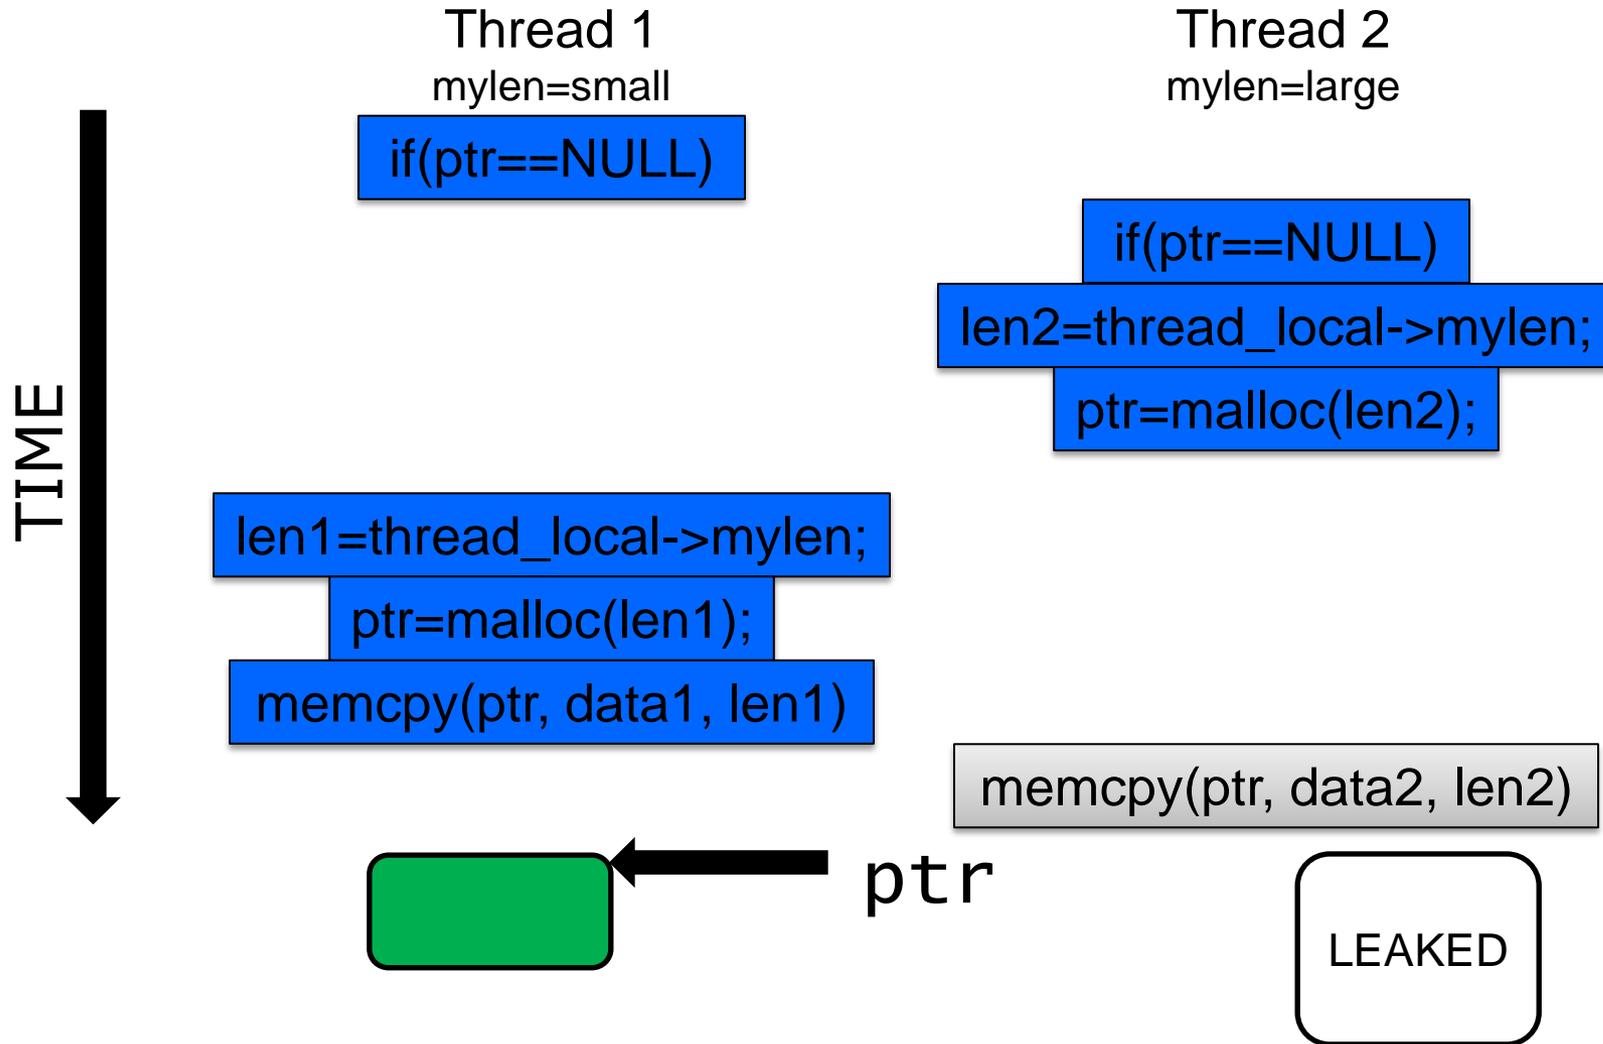
# Example of Modern Bug



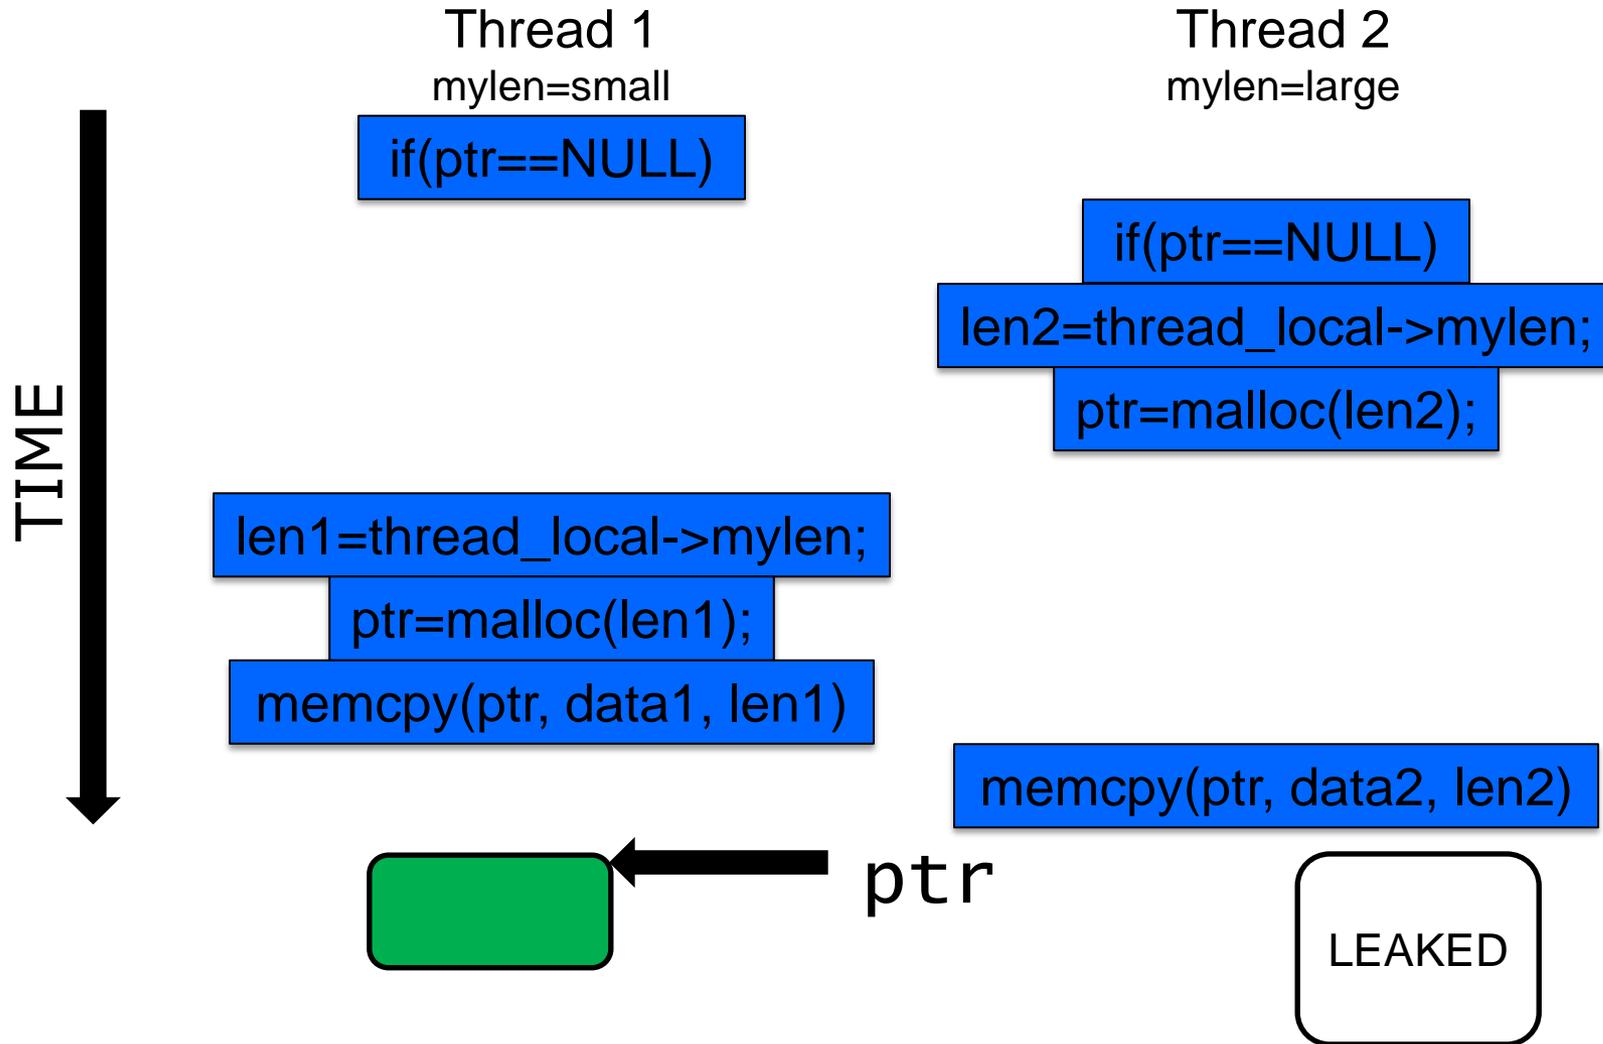
# Example of Modern Bug



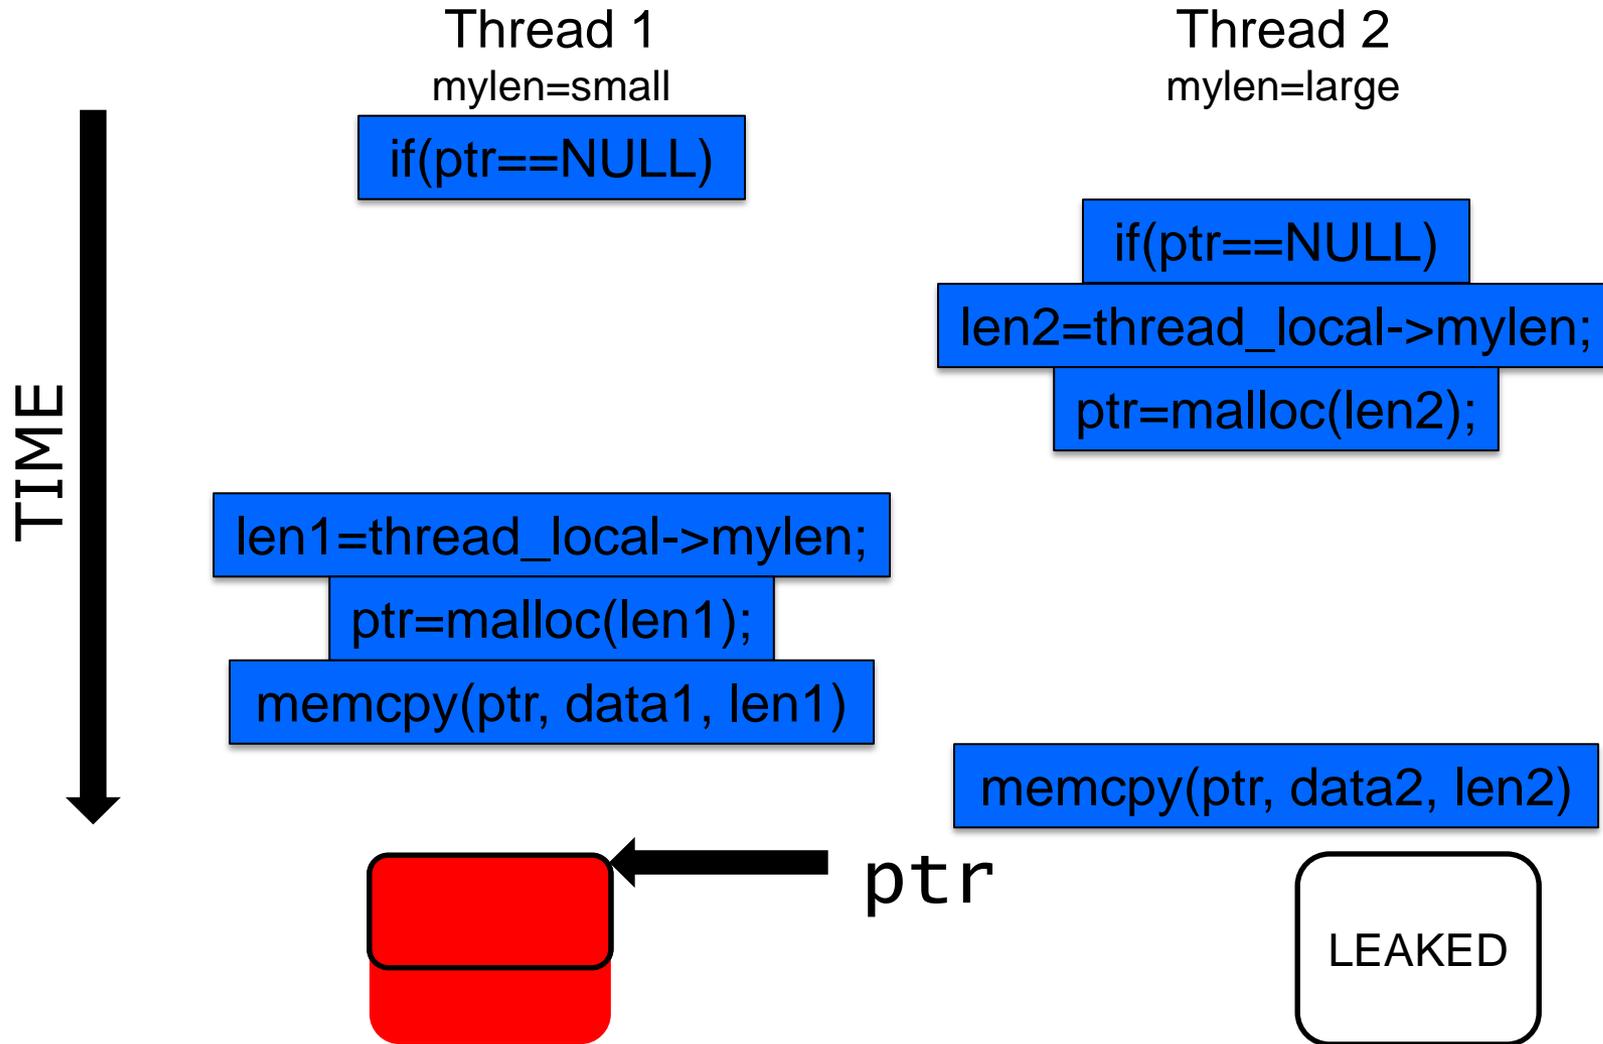
# Example of Modern Bug



# Example of Modern Bug



# Example of Modern Bug



# Dynamic Software Analysis

- Analyze the program as it runs
  - + System state, find errors on any executed path
  - LARGE runtime overheads, only test one path

## Analysis Instrumentation



**Developer**



**Program**



**In-House  
Test Server(s)**

# Dynamic Software Analysis

- Analyze the program as it runs
  - + System state, find errors on any executed path
  - LARGE runtime overheads, only test one path

## Analysis Instrumentation



# Dynamic Software Analysis

- Analyze the program as it runs
  - + System state, find errors on any executed path
  - LARGE runtime overheads, only test one path

## Analysis Instrumentation



**LONG run time**



**Developer**



**In-House  
Test Server(s)**

# Dynamic Software Analysis

- Analyze the program as it runs
  - + System state, find errors on any executed path
  - LARGE runtime overheads, only test one path

## Analysis Instrumentation



## In-House Test Server(s)

# Runtime Overheads: How Large?

- Data Race Detection  
(e.g. Inspector XE)

**2-300x**

- Taint Analysis  
(e.g. TaintCheck)

**2-200x**

- Memory Checking  
(e.g. MemCheck)

**5-50x**

- Dynamic Bounds Checking

**10-80x**

- Symbolic Execution

**10-200x**

---

# Could use Hardware

- Data Race Detection: HARD, CORD, etc.
  - Taint Analysis: Raksha, FlexiTaint, etc.
  - Bounds Checking: HardBound
- 
- None Currently Exist; Bugs Are Here Now
  - Single-Use Specialization
    - Won't be built due to HW, power, verification costs
    - Unchangeable algorithms locked in HW

---

# Goals of this Talk

- Accelerate SW Analyses Using Existing HW
- Run Tests **On Demand**: Only When Needed
- Explore Future **Generic HW Additions**

---

# Outline

- Problem Statement
- Background Information
  - Demand-Driven Dynamic Dataflow Analysis
- Proposed Solutions
  - Demand-Driven Data Race Detection
  - Unlimited Hardware Watchpoints

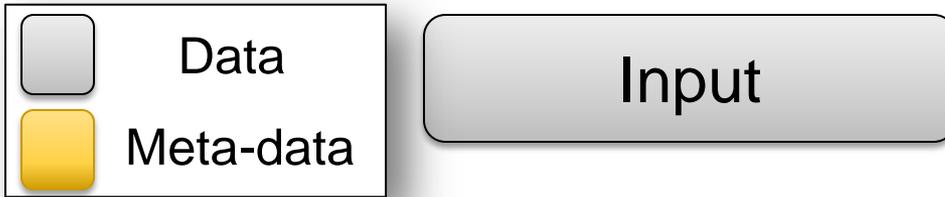
---

# Dynamic Dataflow Analysis

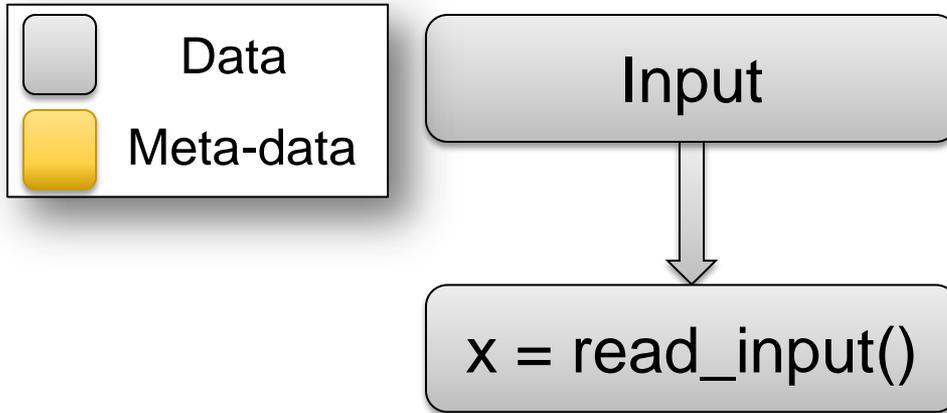
- **Associate** meta-data with program values
- **Propagate/Clear** meta-data while executing
- **Check** meta-data for safety & correctness
- Forms dataflows of meta/shadow information

---

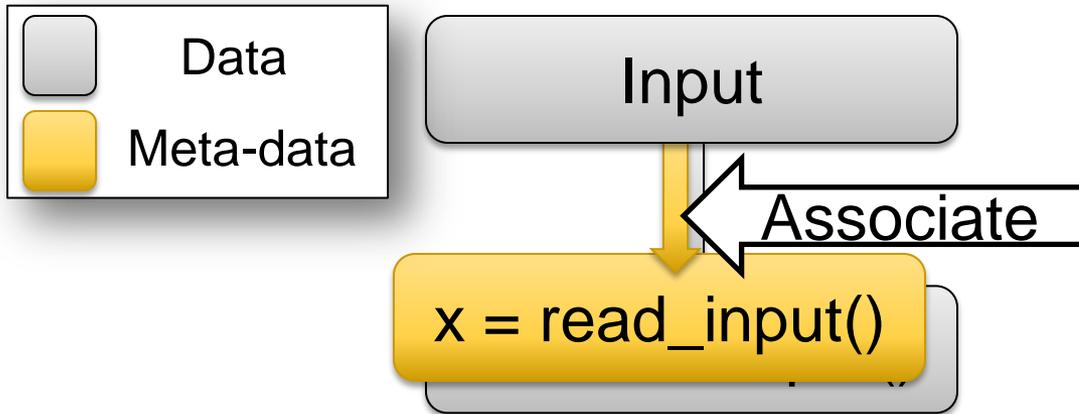
# Example: Taint Analysis



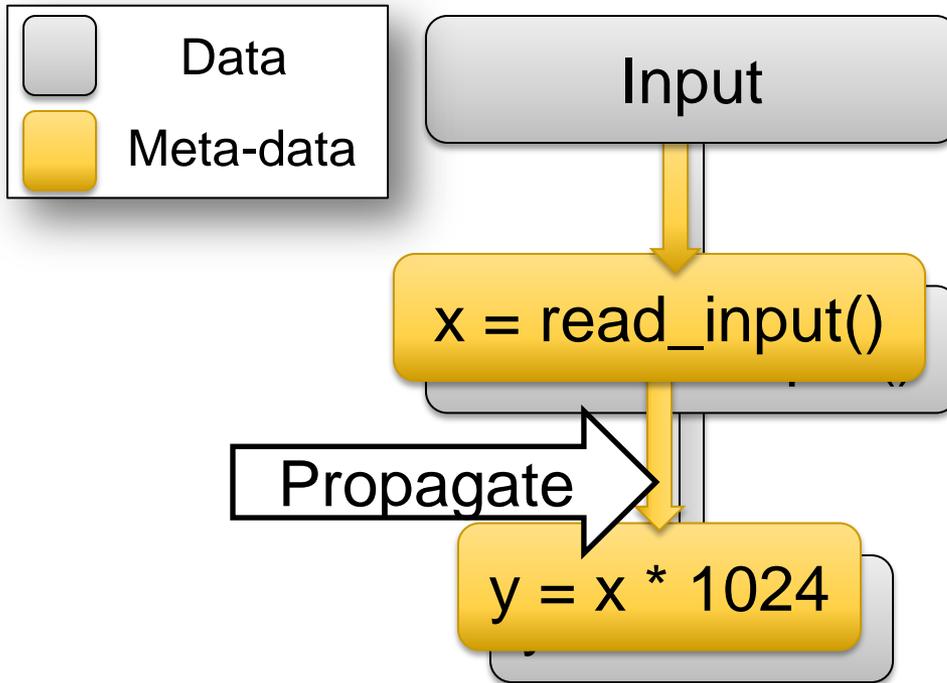
# Example: Taint Analysis



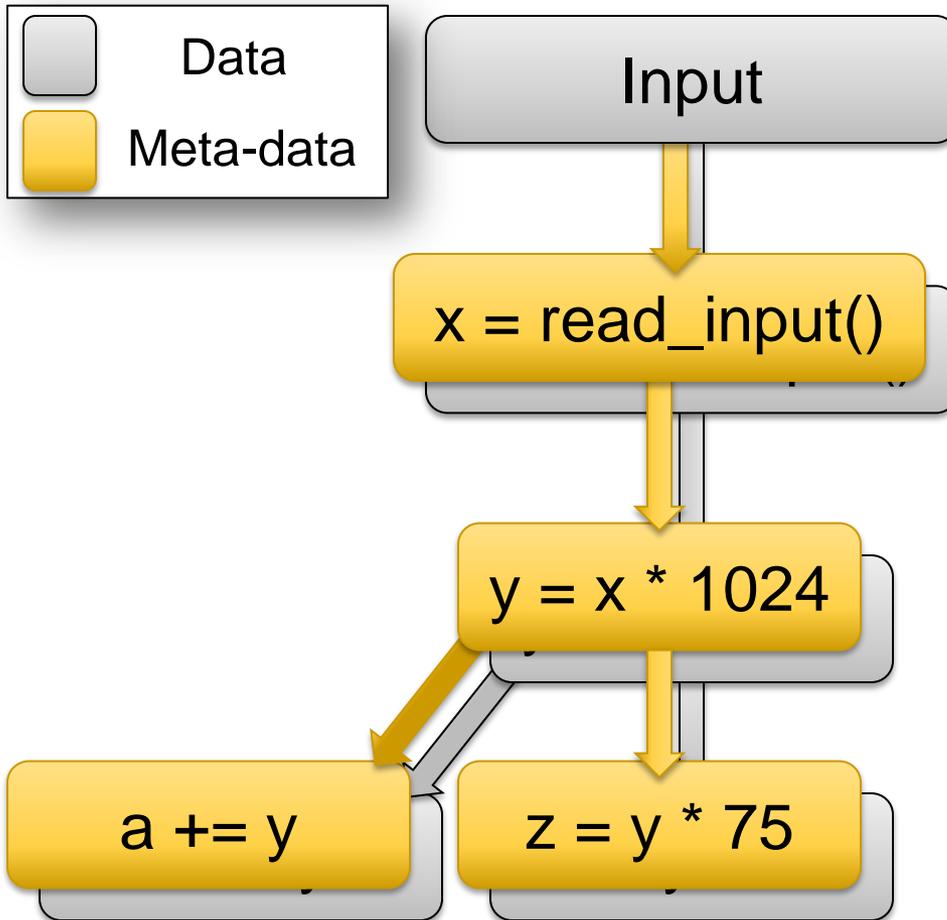
# Example: Taint Analysis



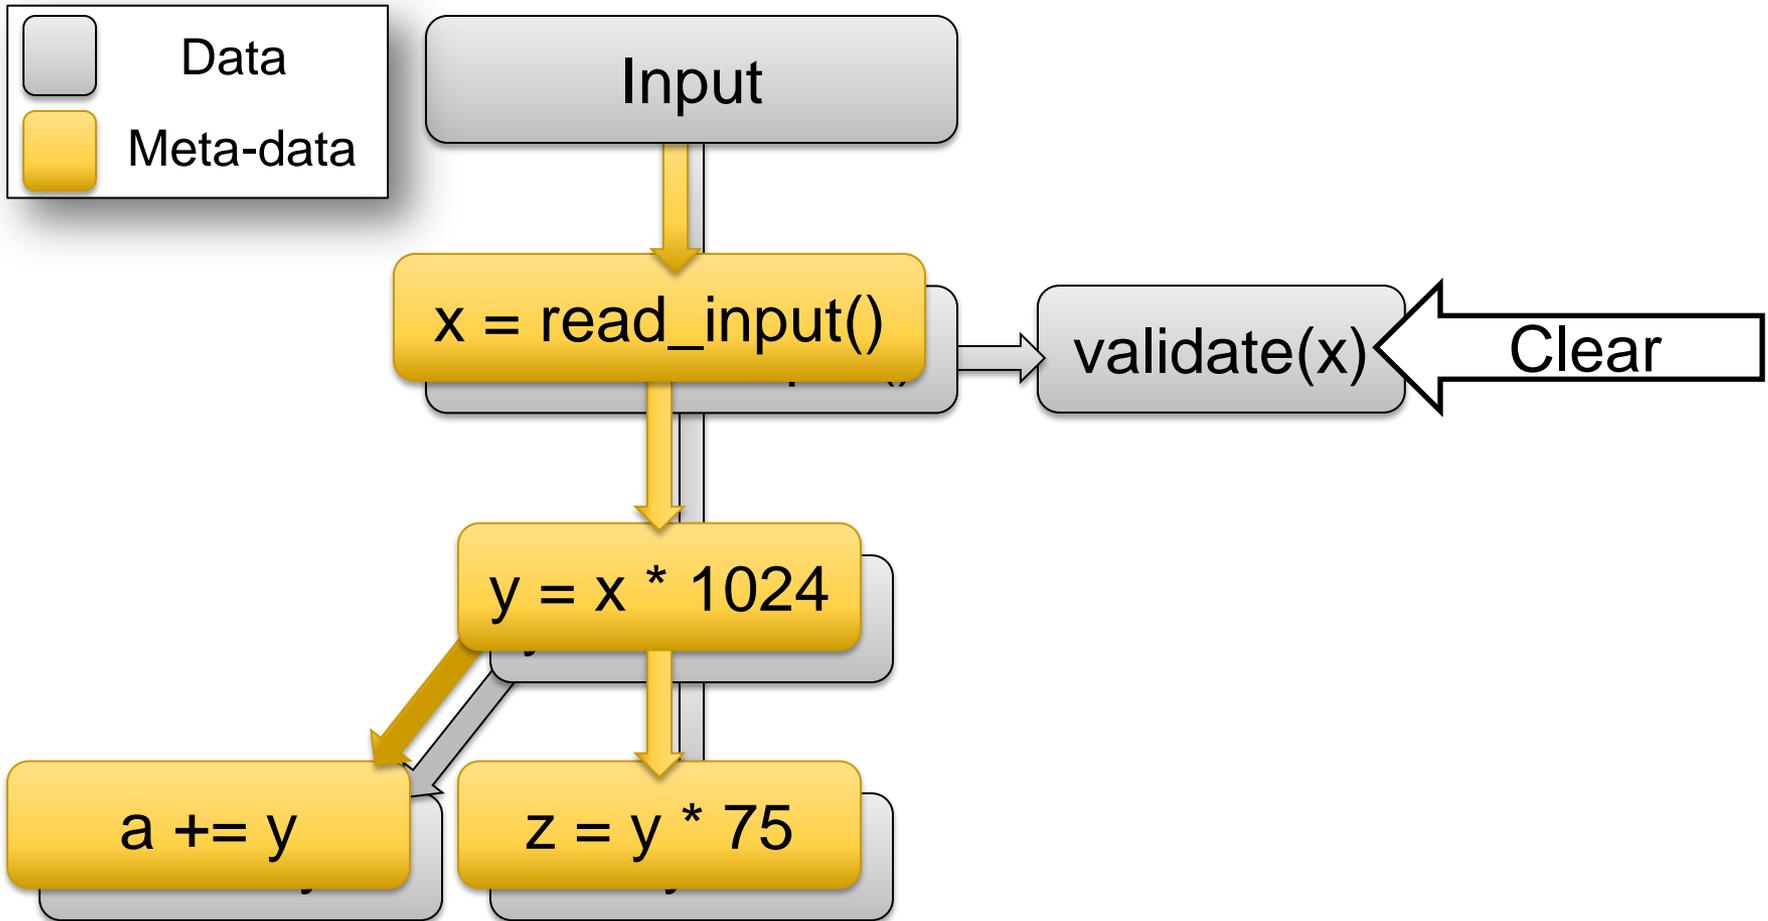
# Example: Taint Analysis



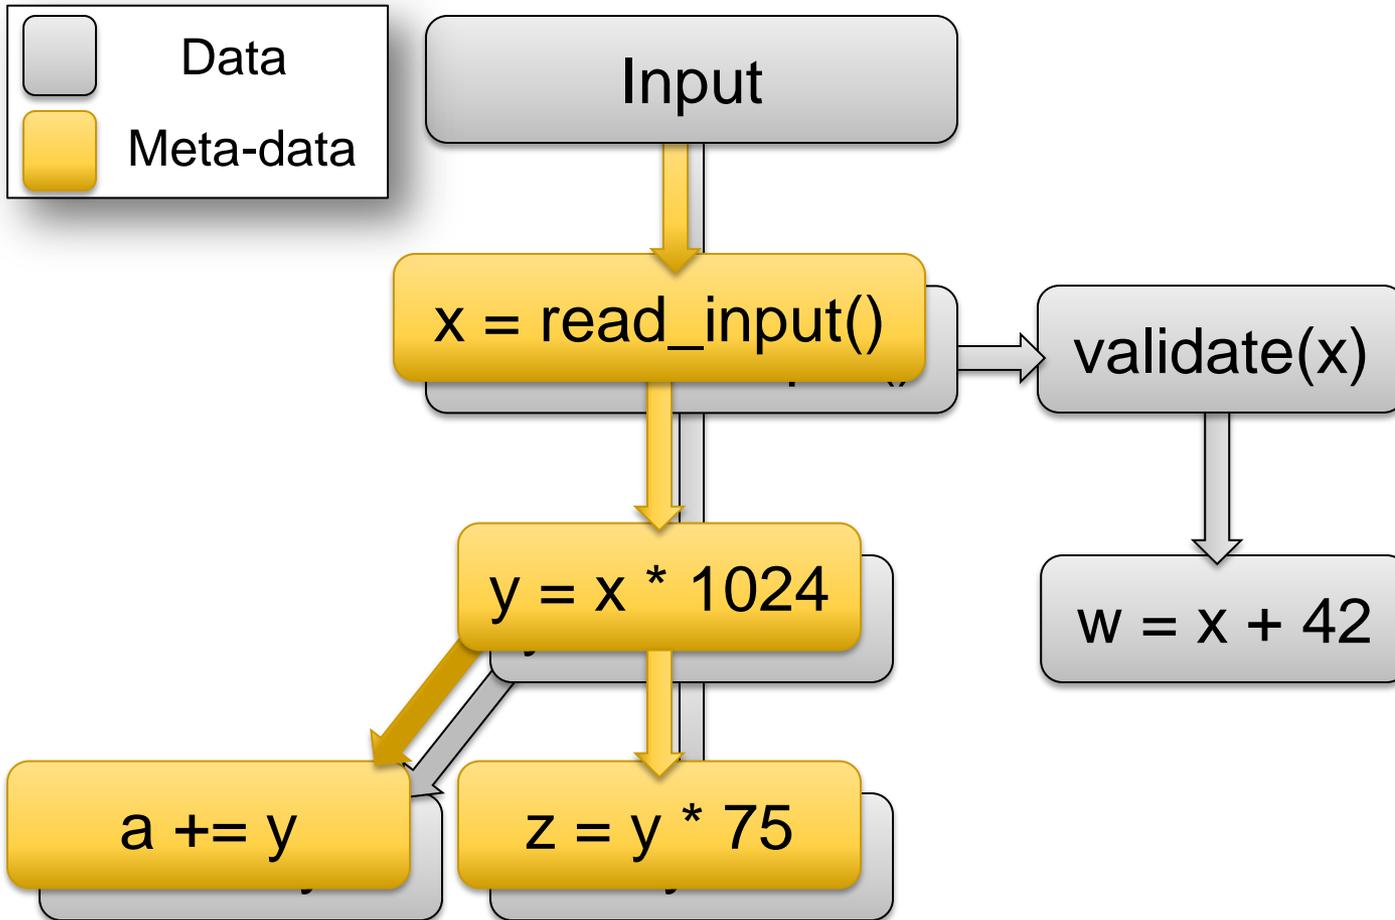
# Example: Taint Analysis



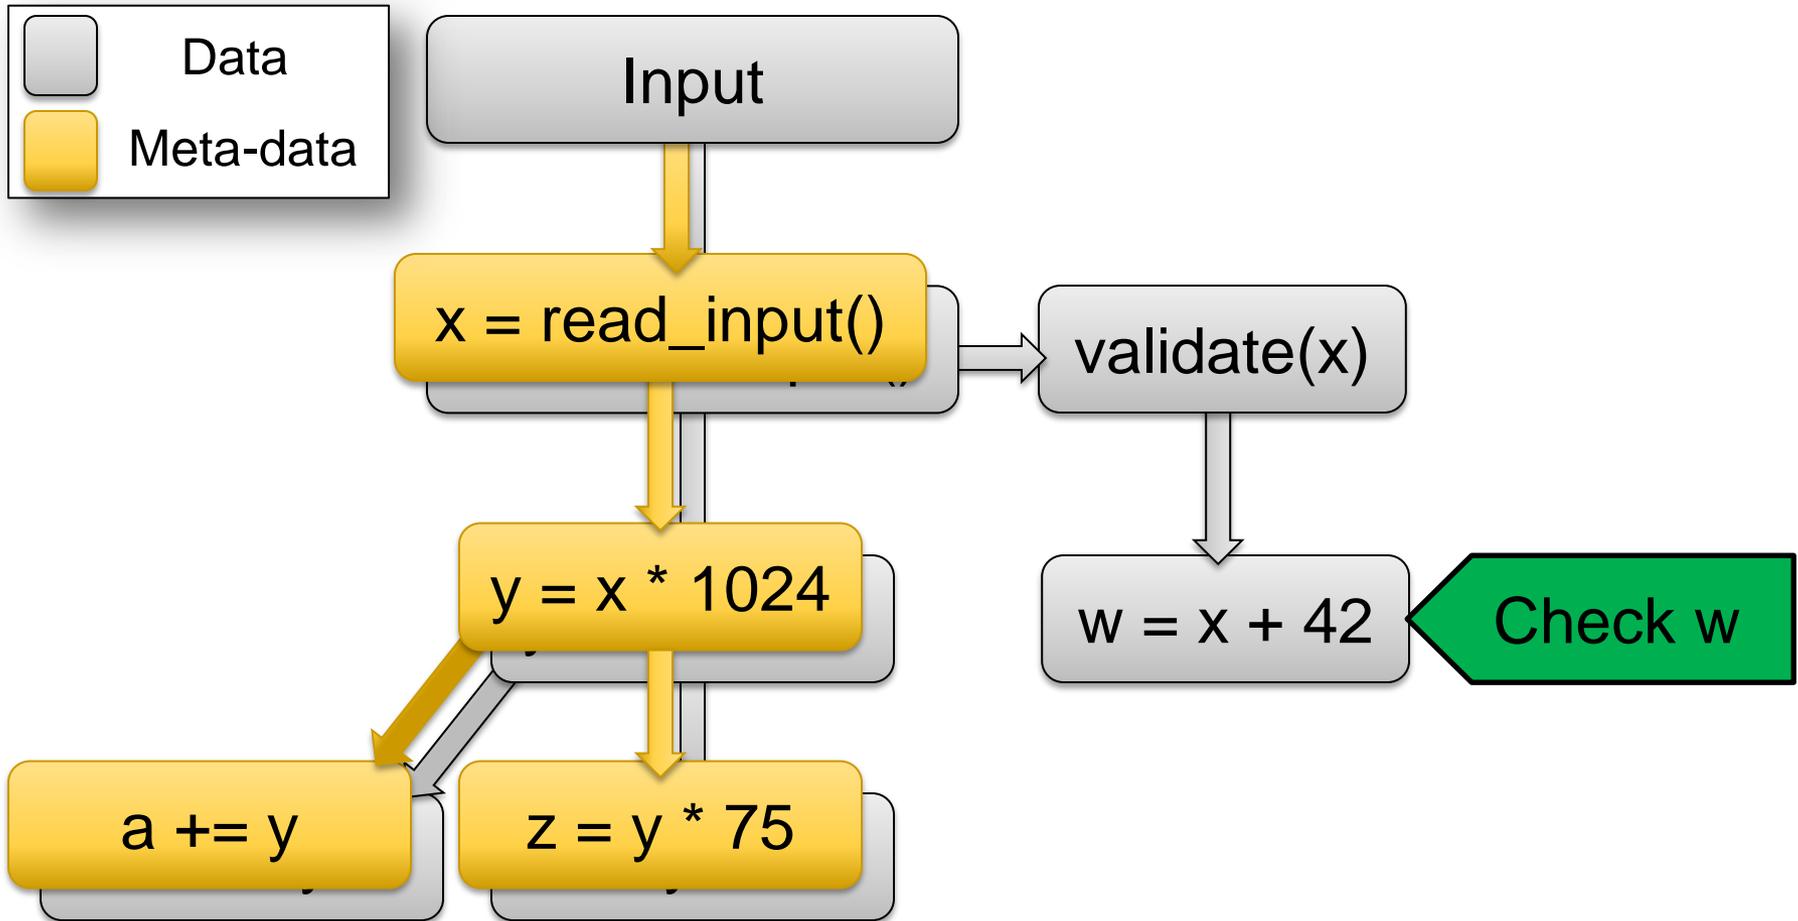
# Example: Taint Analysis



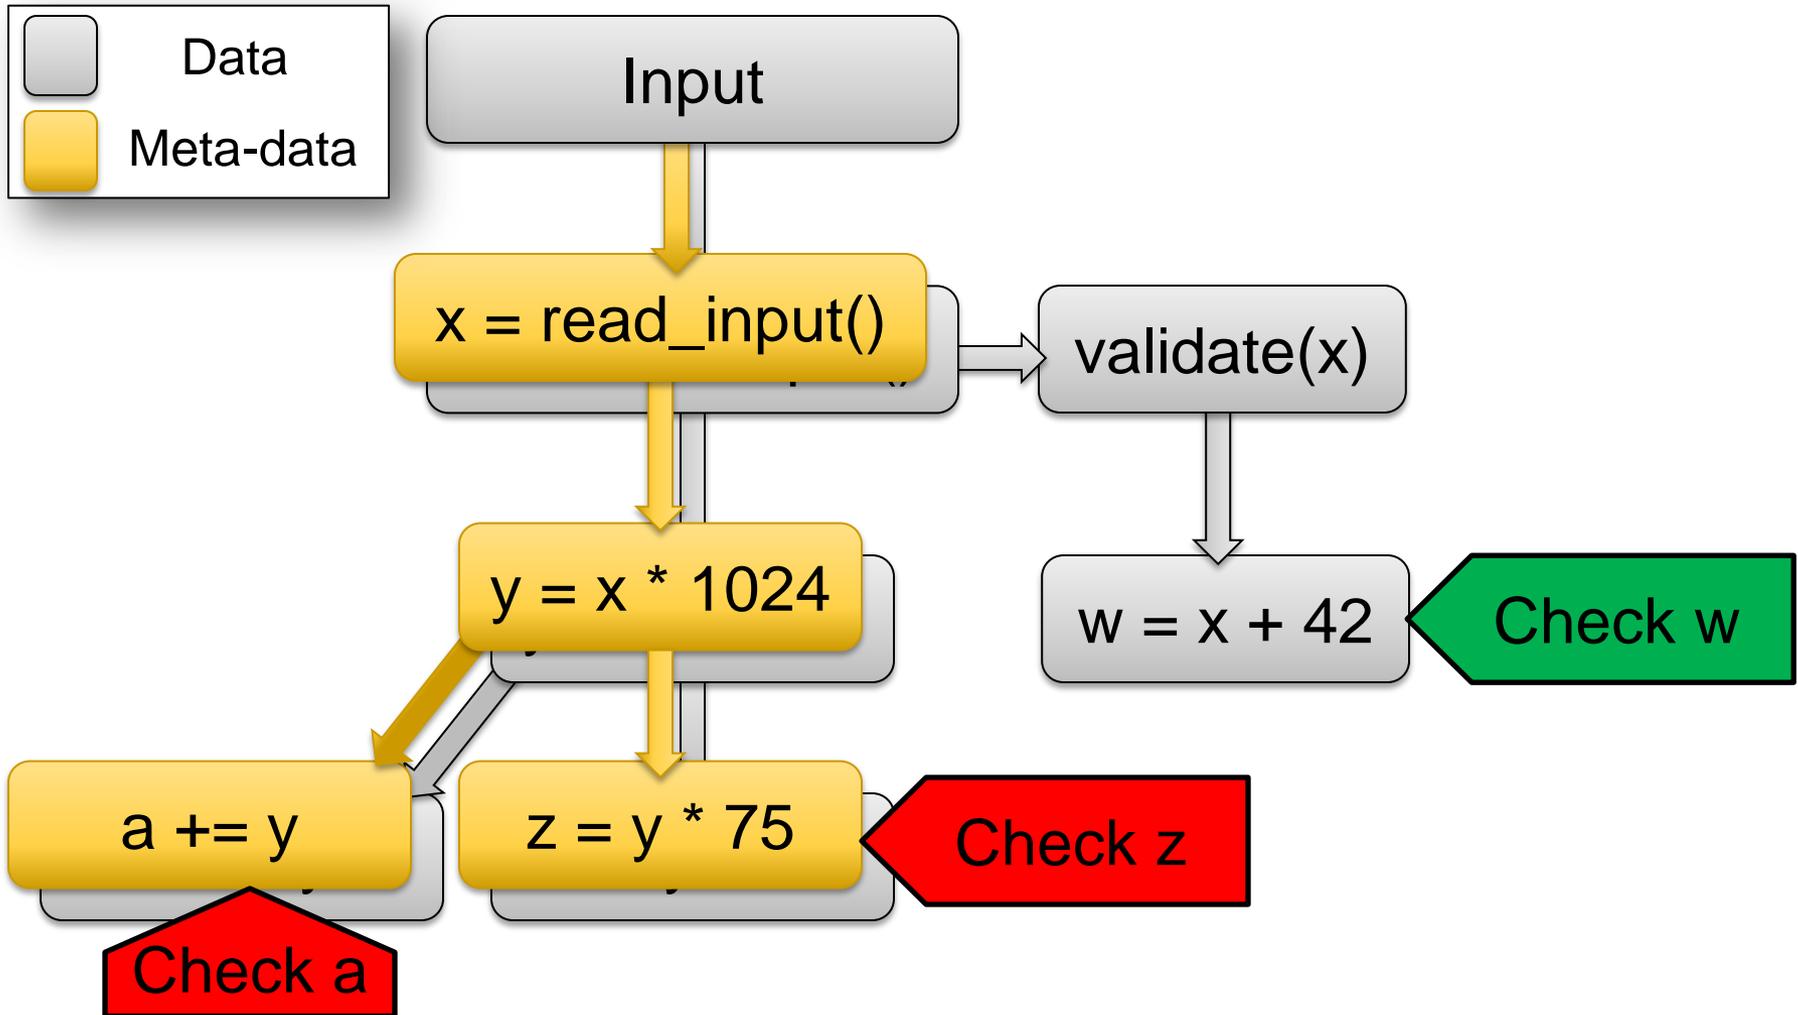
# Example: Taint Analysis



# Example: Taint Analysis

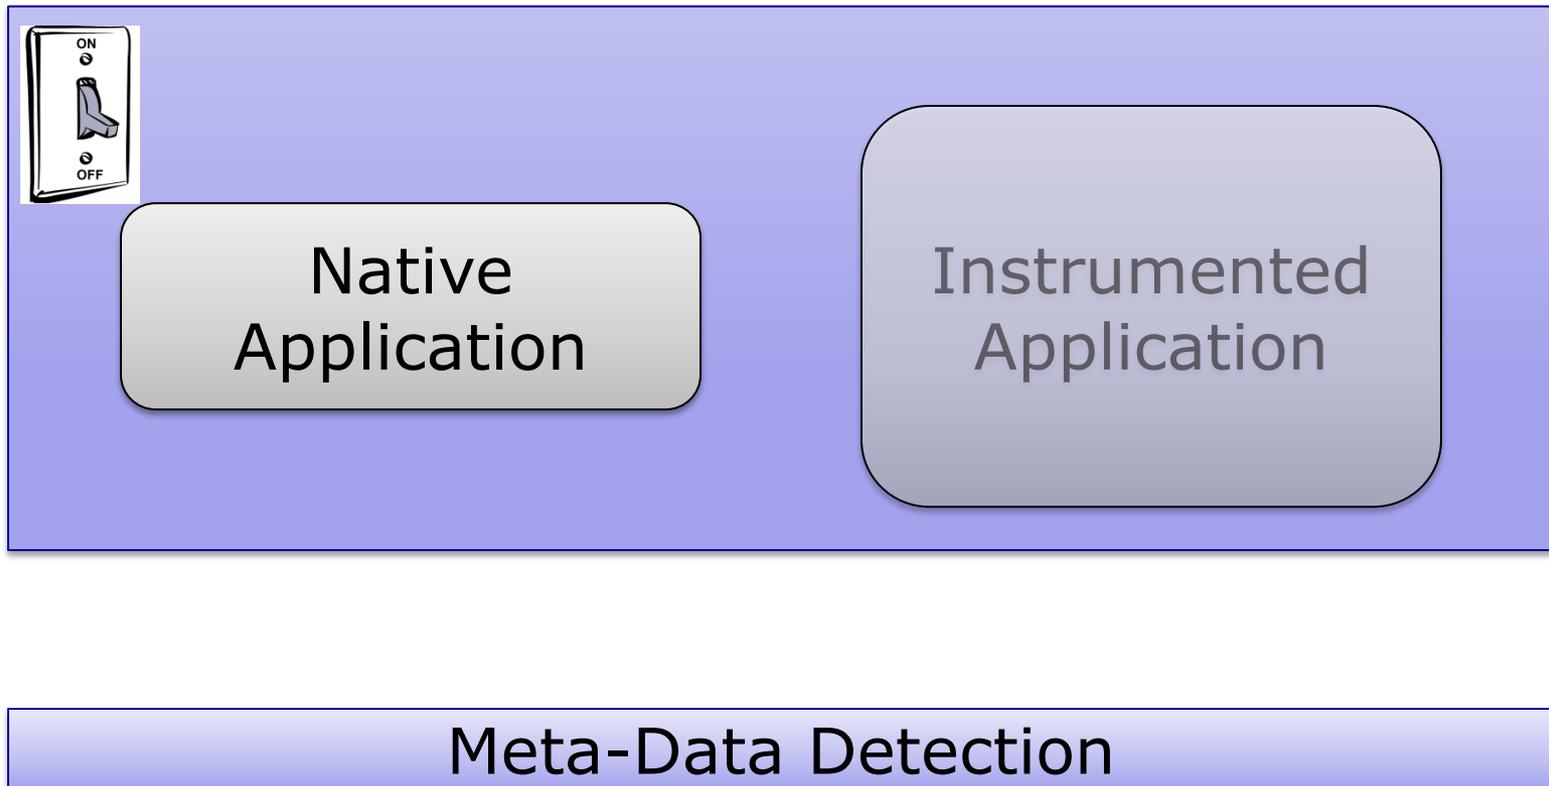


# Example: Taint Analysis



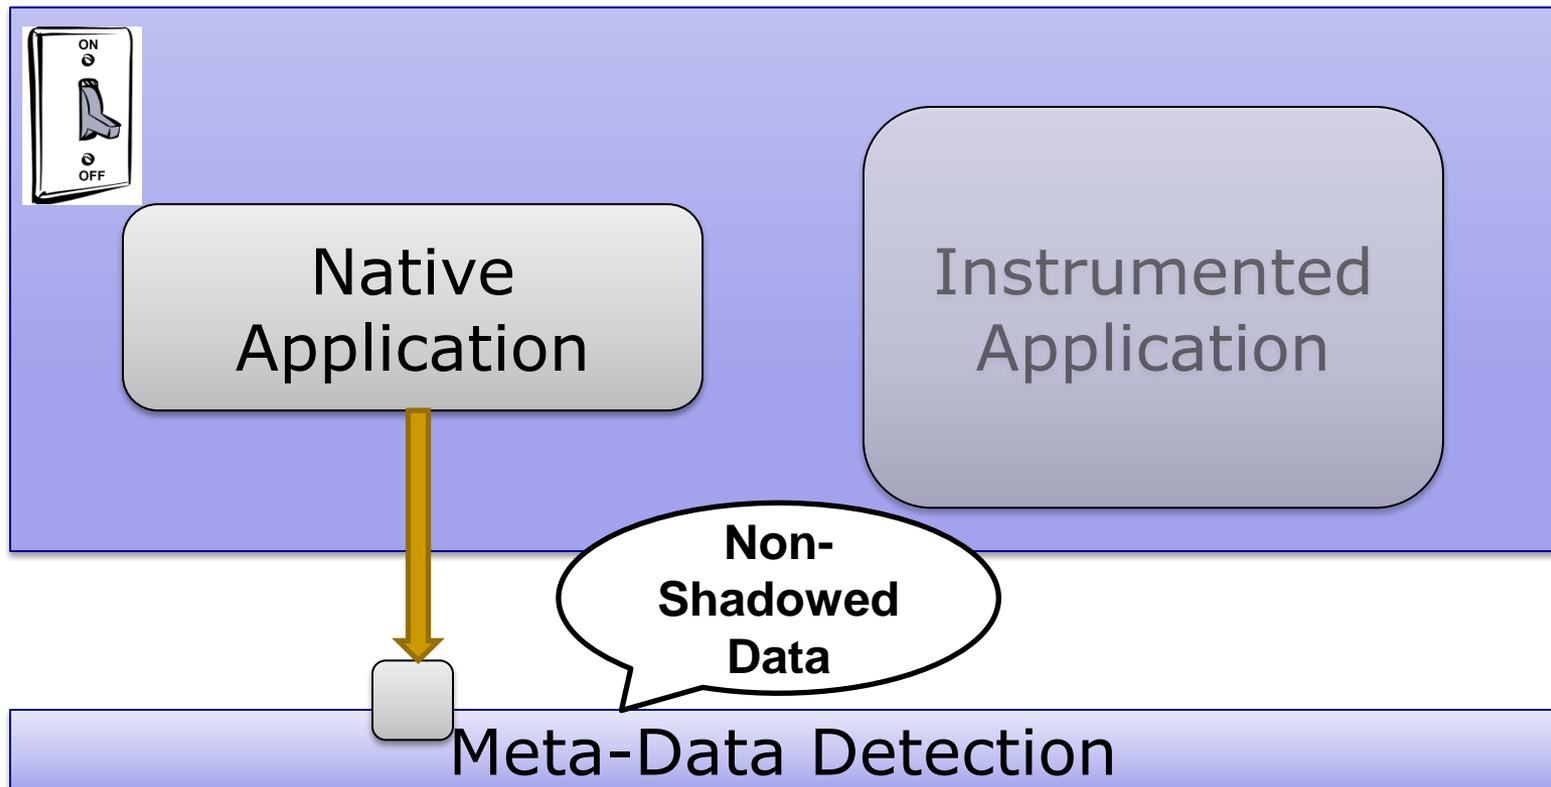
# Demand-Driven Dataflow Analysis

- Only Analyze Shadowed Data



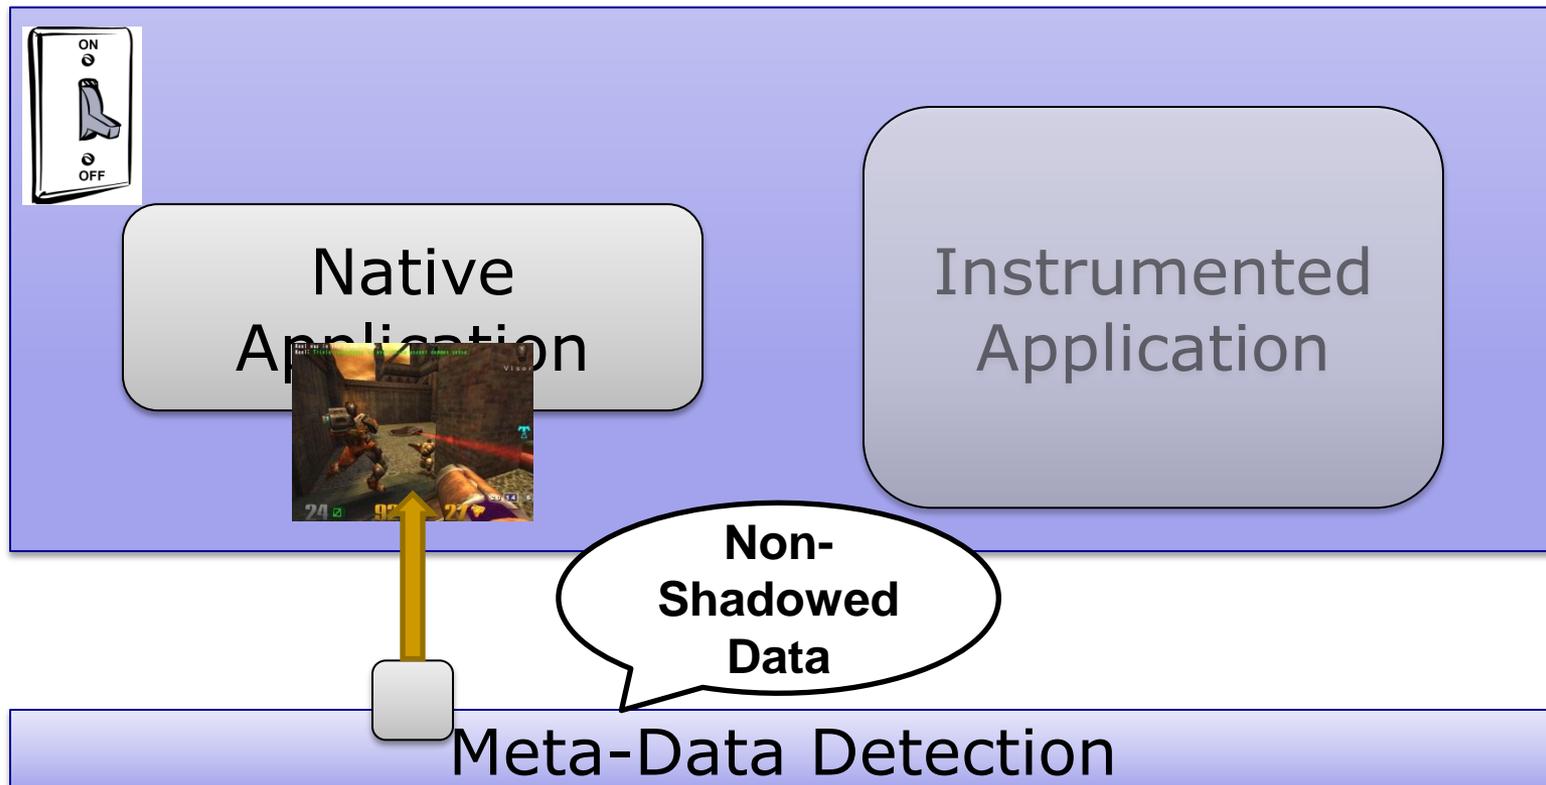
# Demand-Driven Dataflow Analysis

- Only Analyze Shadowed Data



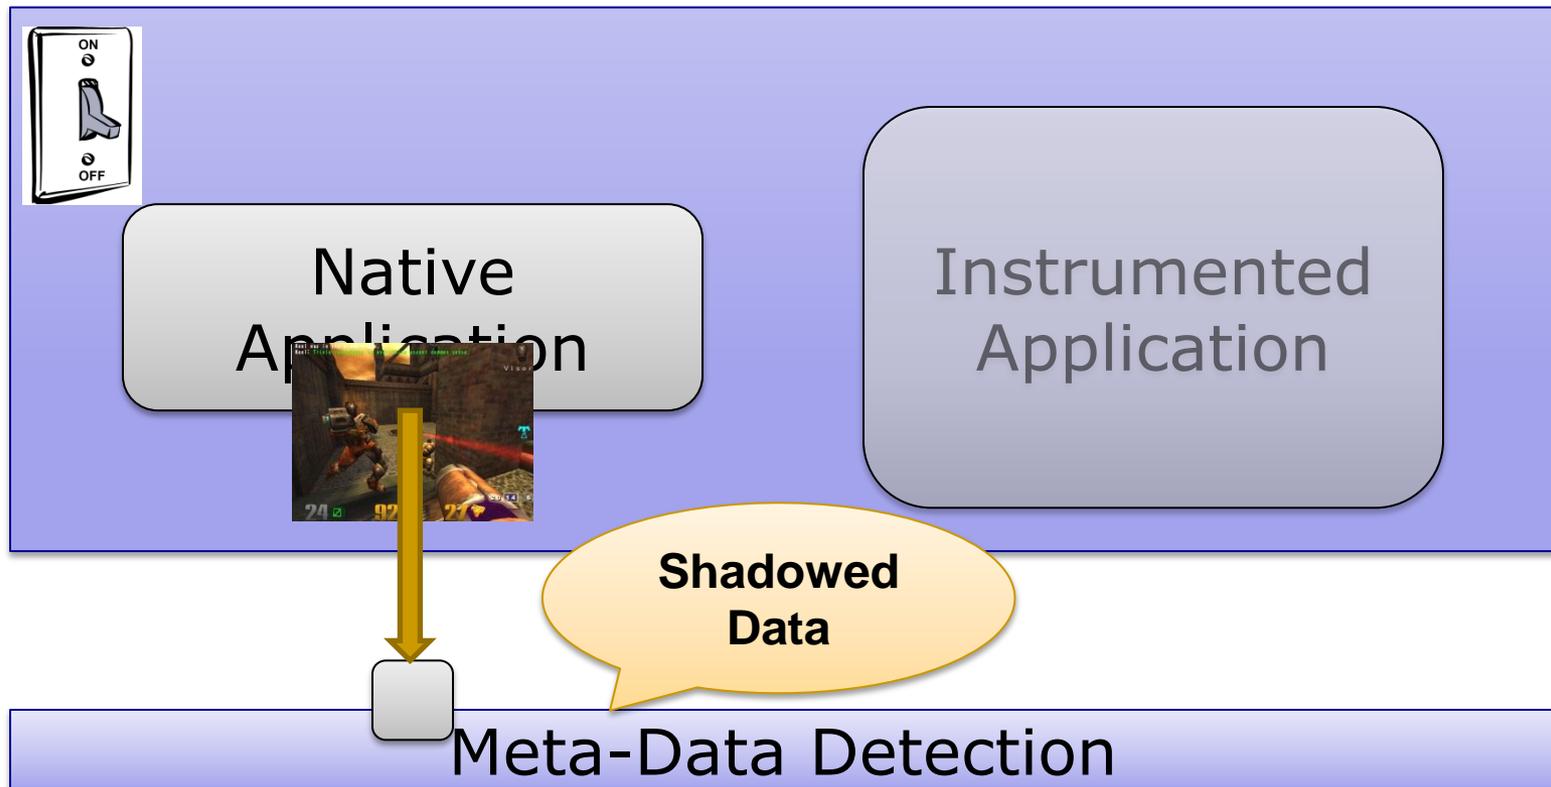
# Demand-Driven Dataflow Analysis

- Only Analyze Shadowed Data



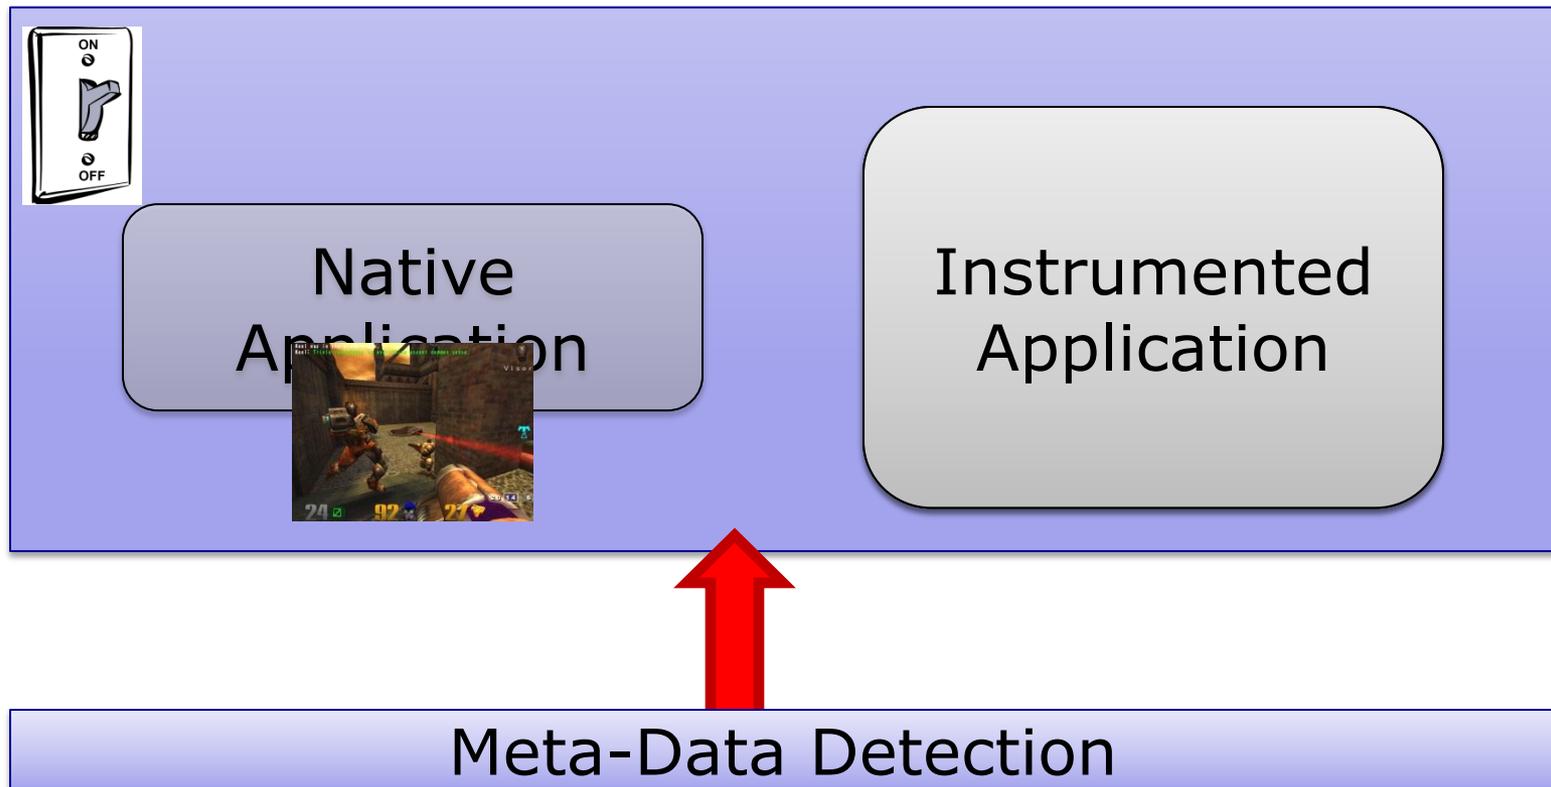
# Demand-Driven Dataflow Analysis

- Only Analyze Shadowed Data



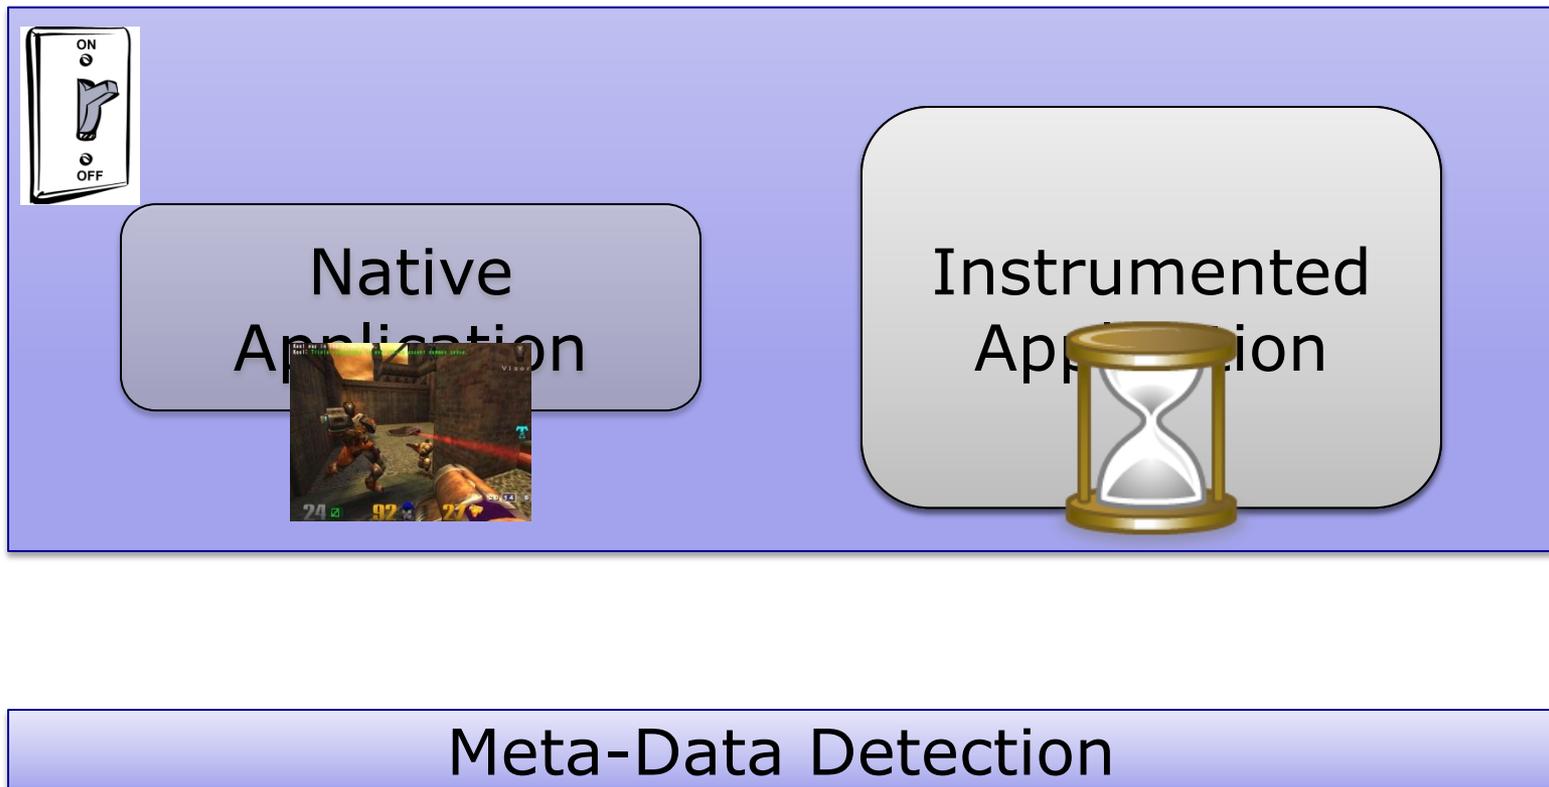
# Demand-Driven Dataflow Analysis

- Only Analyze Shadowed Data



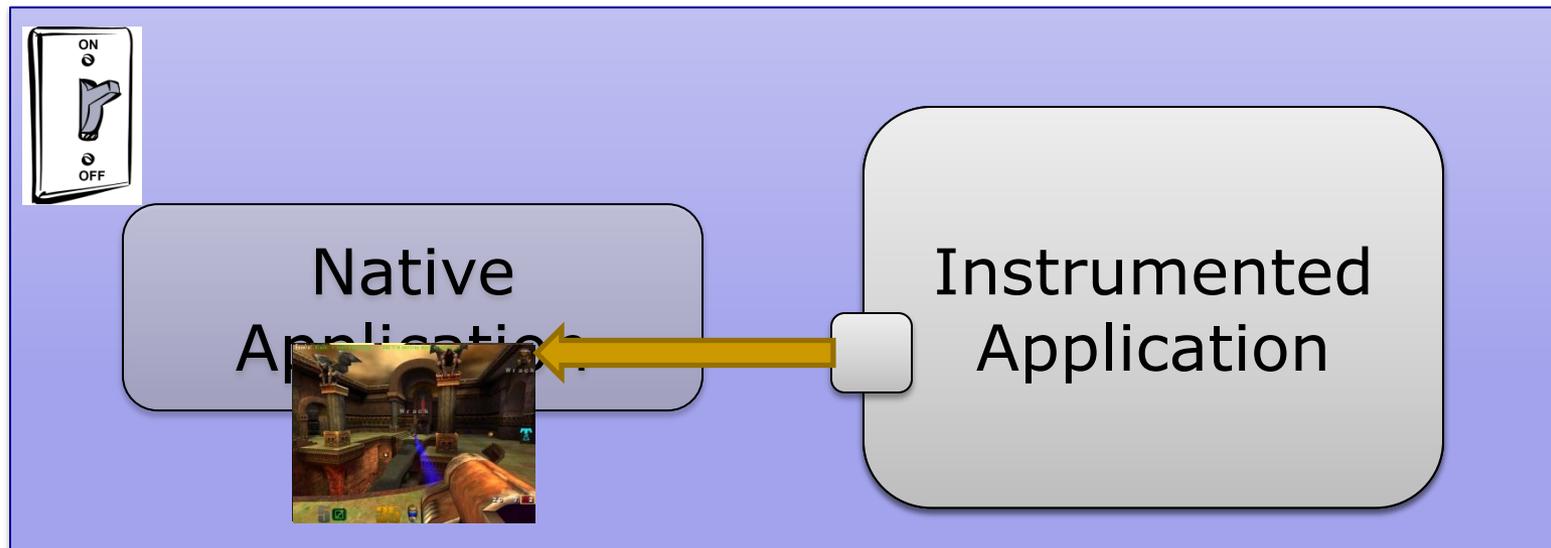
# Demand-Driven Dataflow Analysis

- Only Analyze Shadowed Data



# Demand-Driven Dataflow Analysis

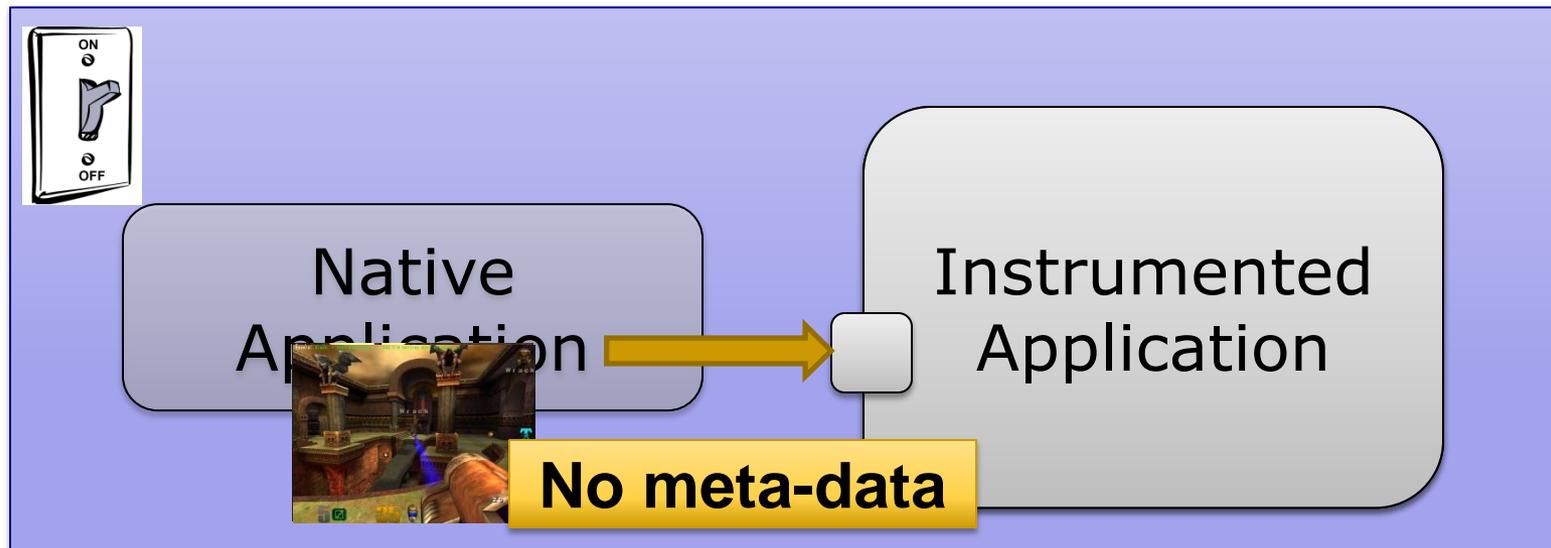
- Only Analyze Shadowed Data



Meta-Data Detection

# Demand-Driven Dataflow Analysis

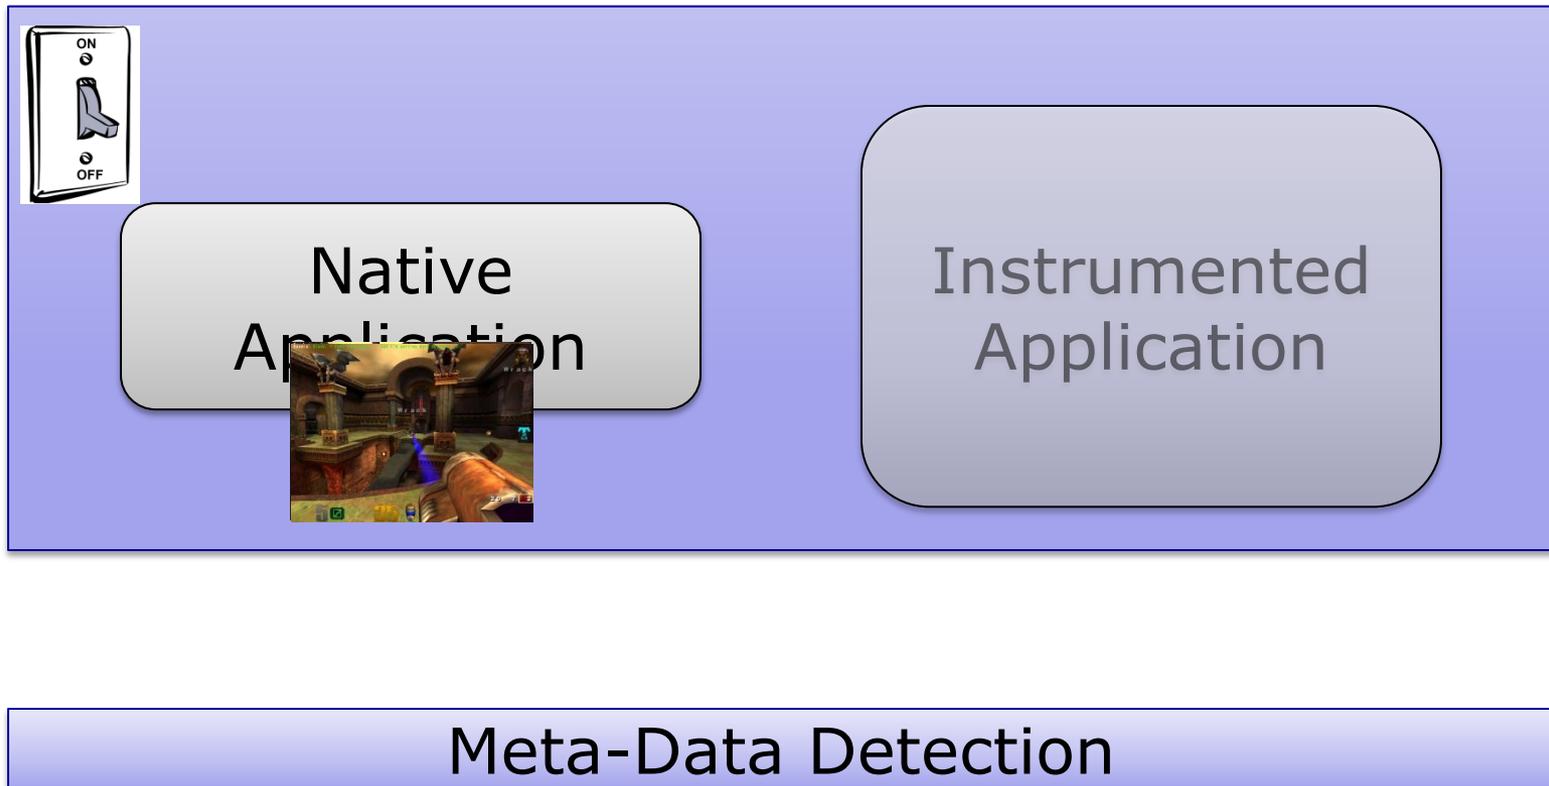
- Only Analyze Shadowed Data



Meta-Data Detection

# Demand-Driven Dataflow Analysis

- Only Analyze Shadowed Data



---

# Finding Meta-Data

- No additional overhead when no meta-data
  - Needs hardware support
- Take a fault when touching shadowed data

---

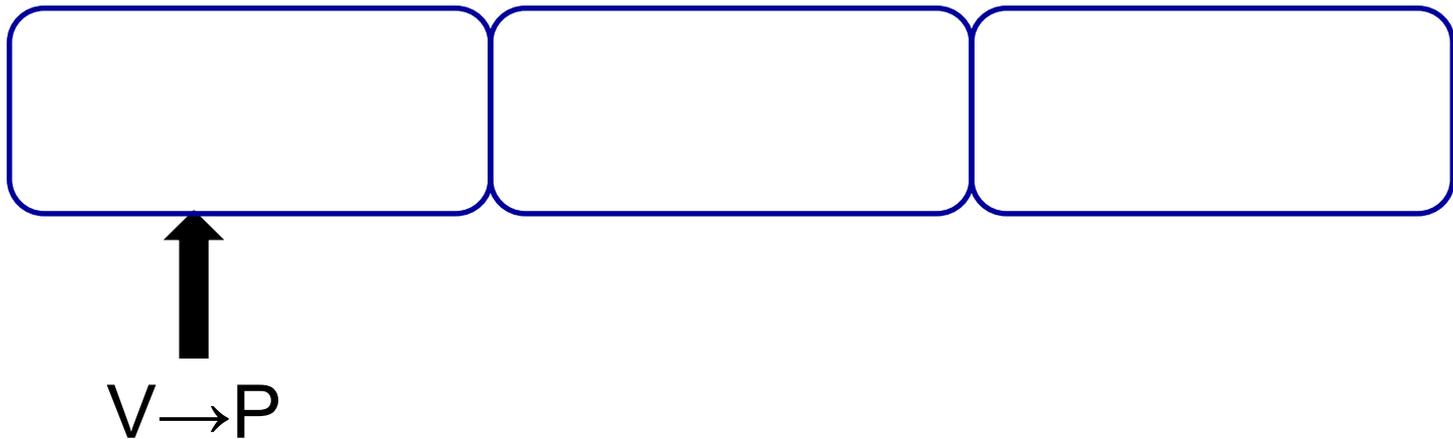
# Finding Meta-Data

- No additional overhead when no meta-data
  - Needs hardware support
- Take a fault when touching shadowed data
- Solution: Virtual Memory Watchpoints



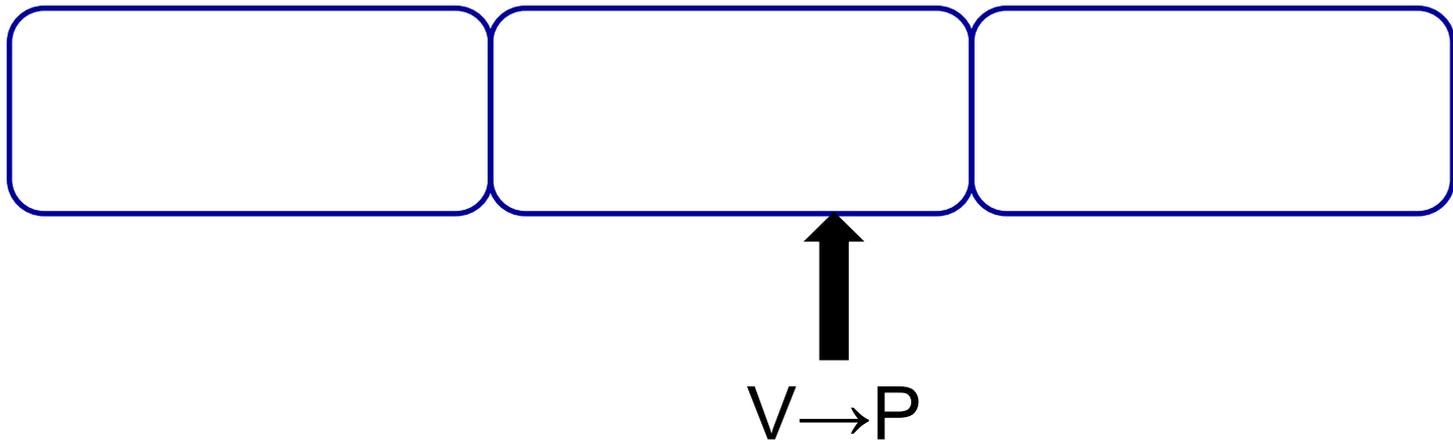
# Finding Meta-Data

- No additional overhead when no meta-data
  - Needs hardware support
- Take a fault when touching shadowed data
- Solution: Virtual Memory Watchpoints



# Finding Meta-Data

- No additional overhead when no meta-data
  - Needs hardware support
- Take a fault when touching shadowed data
- Solution: Virtual Memory Watchpoints



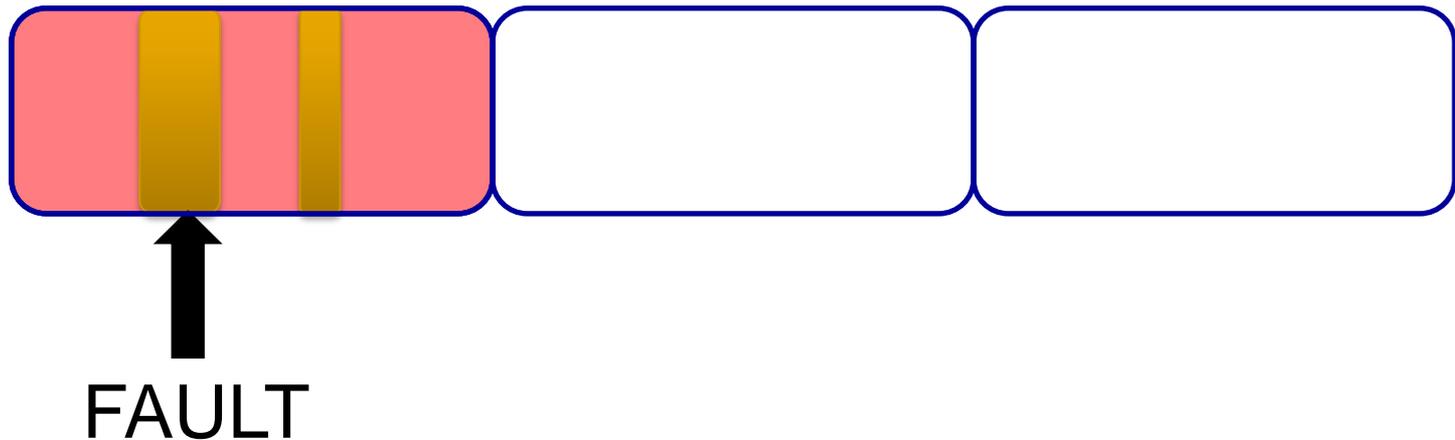
# Finding Meta-Data

- No additional overhead when no meta-data
  - Needs hardware support
- Take a fault when touching shadowed data
- Solution: Virtual Memory Watchpoints



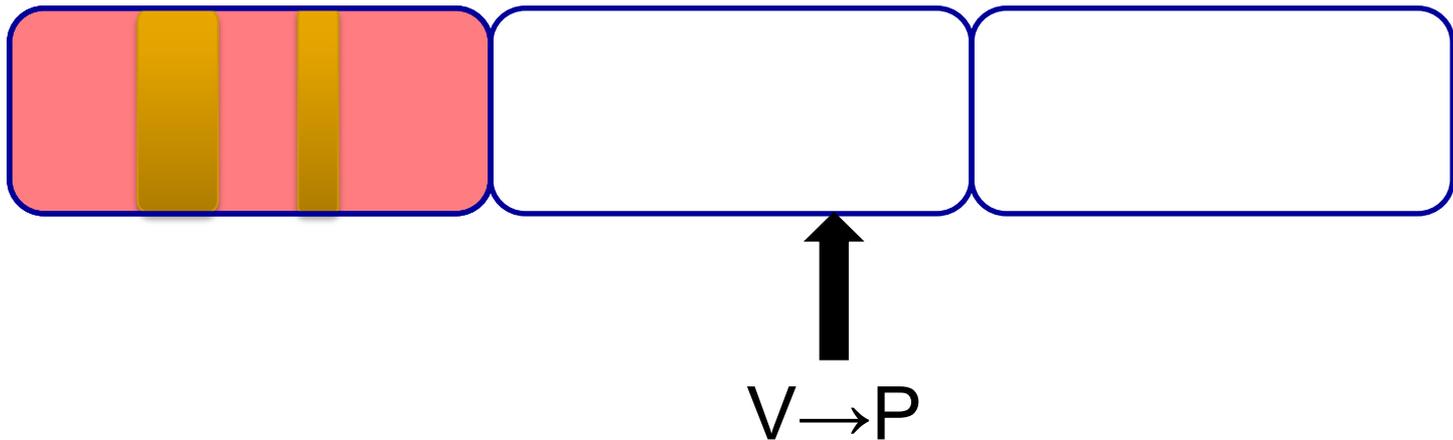
# Finding Meta-Data

- No additional overhead when no meta-data
  - Needs hardware support
- Take a fault when touching shadowed data
- Solution: Virtual Memory Watchpoints



# Finding Meta-Data

- No additional overhead when no meta-data
  - Needs hardware support
- Take a fault when touching shadowed data
- Solution: Virtual Memory Watchpoints

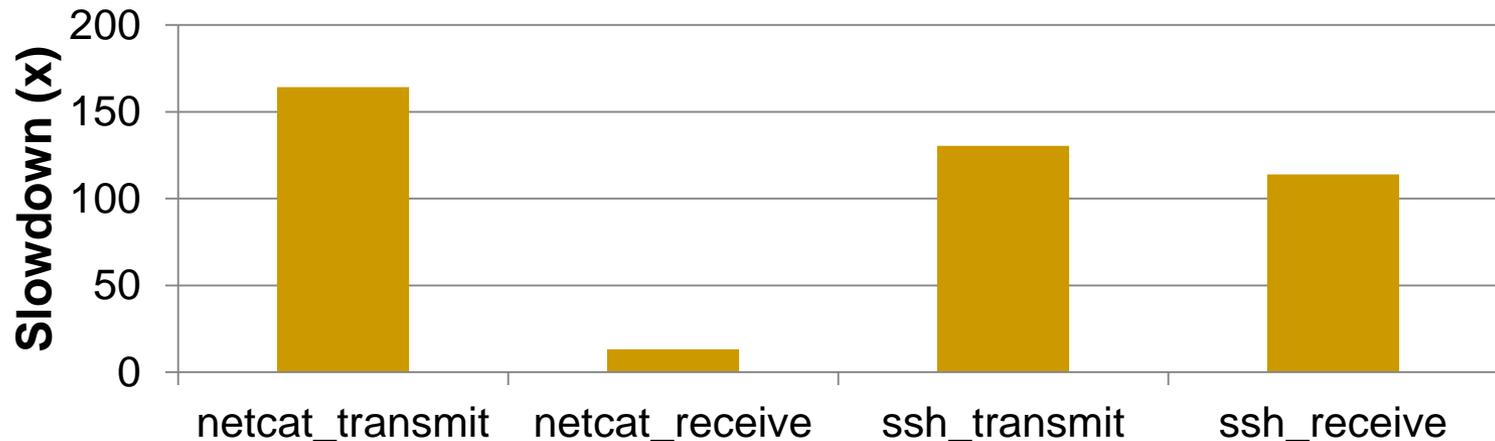


# Results by Ho et al.

## ■ Imbench Best Case Results:

System	Slowdown (normalized)
Taint Analysis	101.7x
On-Demand Taint Analysis	1.98x

## ■ Results when everything is tainted:



---

# Outline

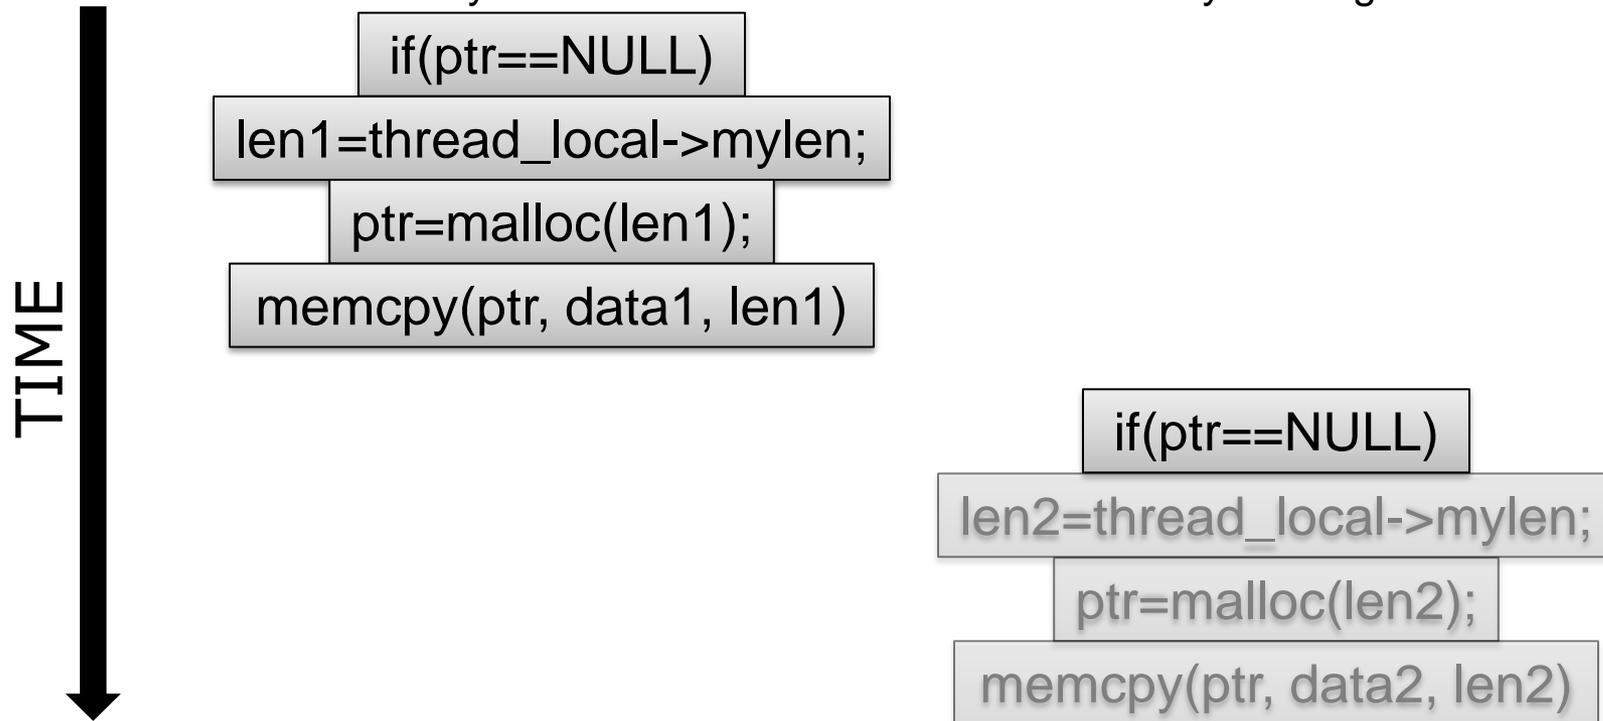
- Problem Statement
- Background Information
  - Demand-Driven Dynamic Dataflow Analysis
- Proposed Solutions
  - Demand-Driven Data Race Detection
  - Unlimited Hardware Watchpoints

---

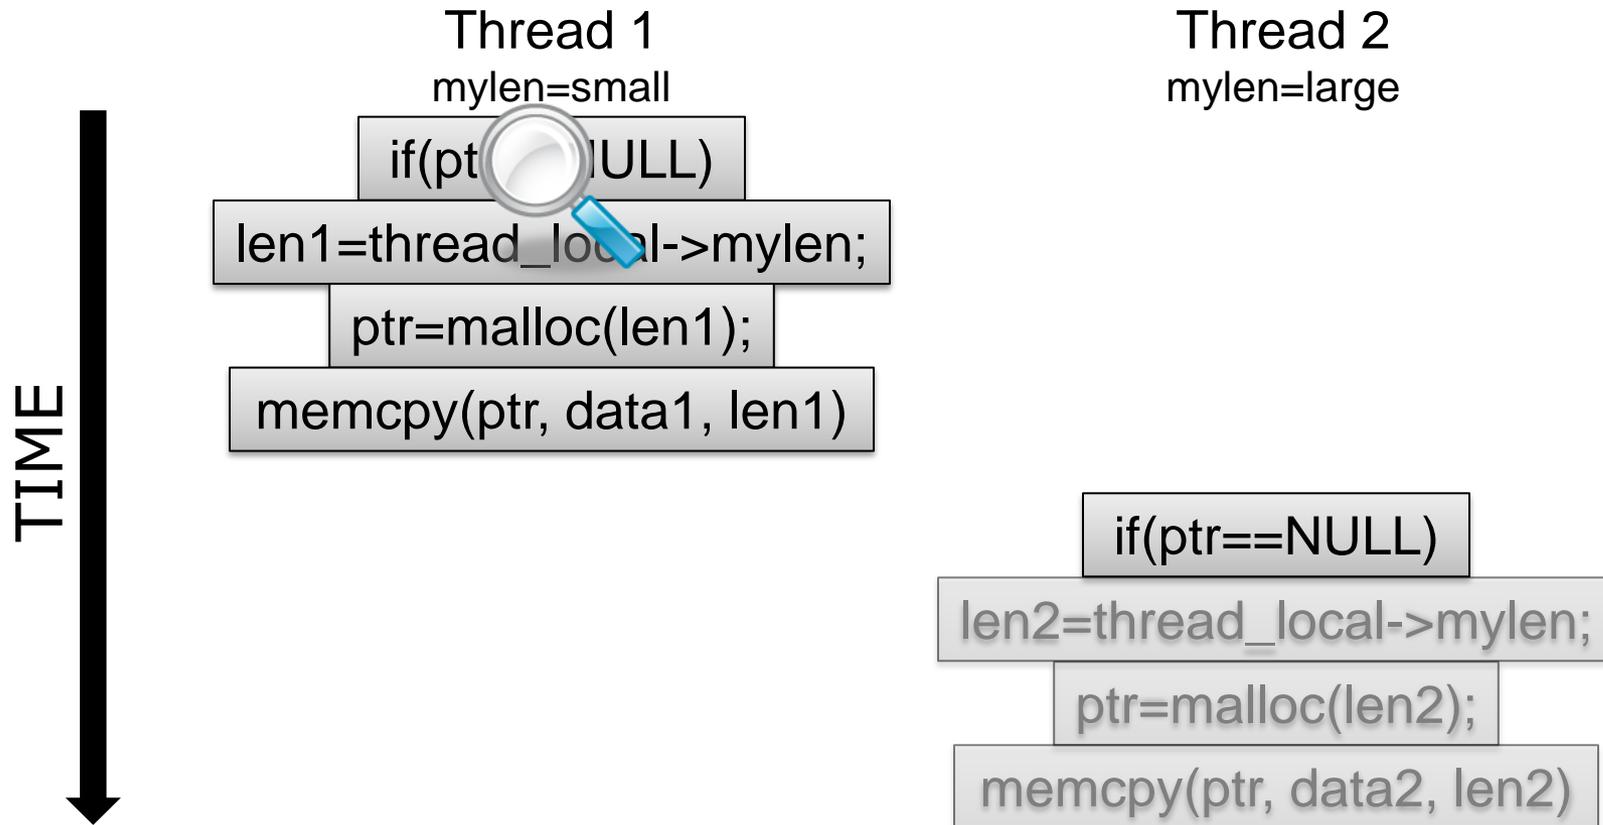
# Software Data Race Detection

- Add checks around every memory access
- Find inter-thread sharing events
- Synchronization between write-shared accesses?
  - No? Data race.

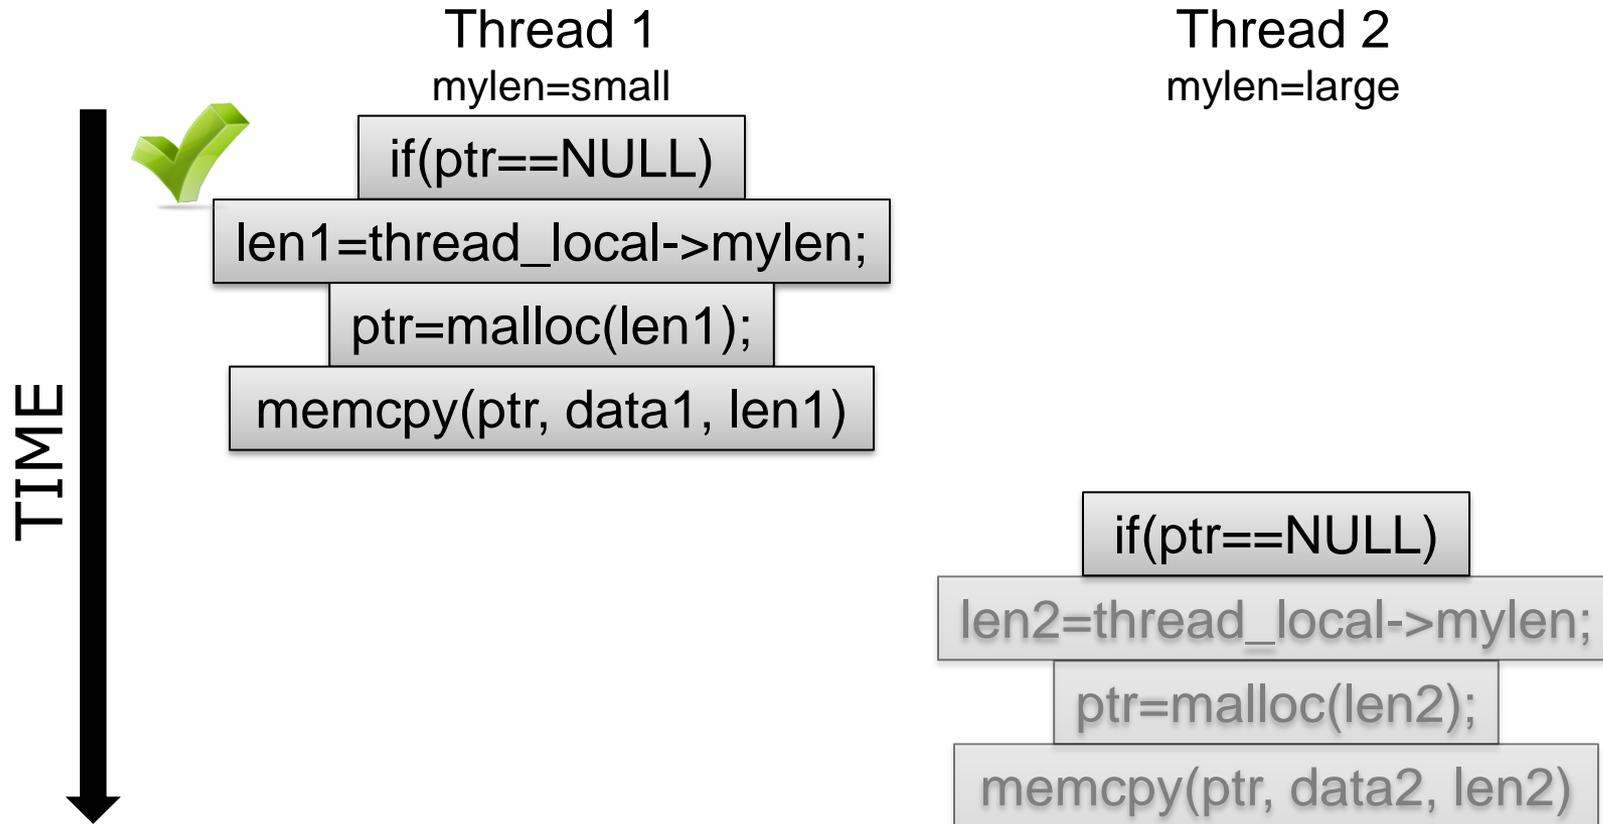
# Example of Data Race Detection



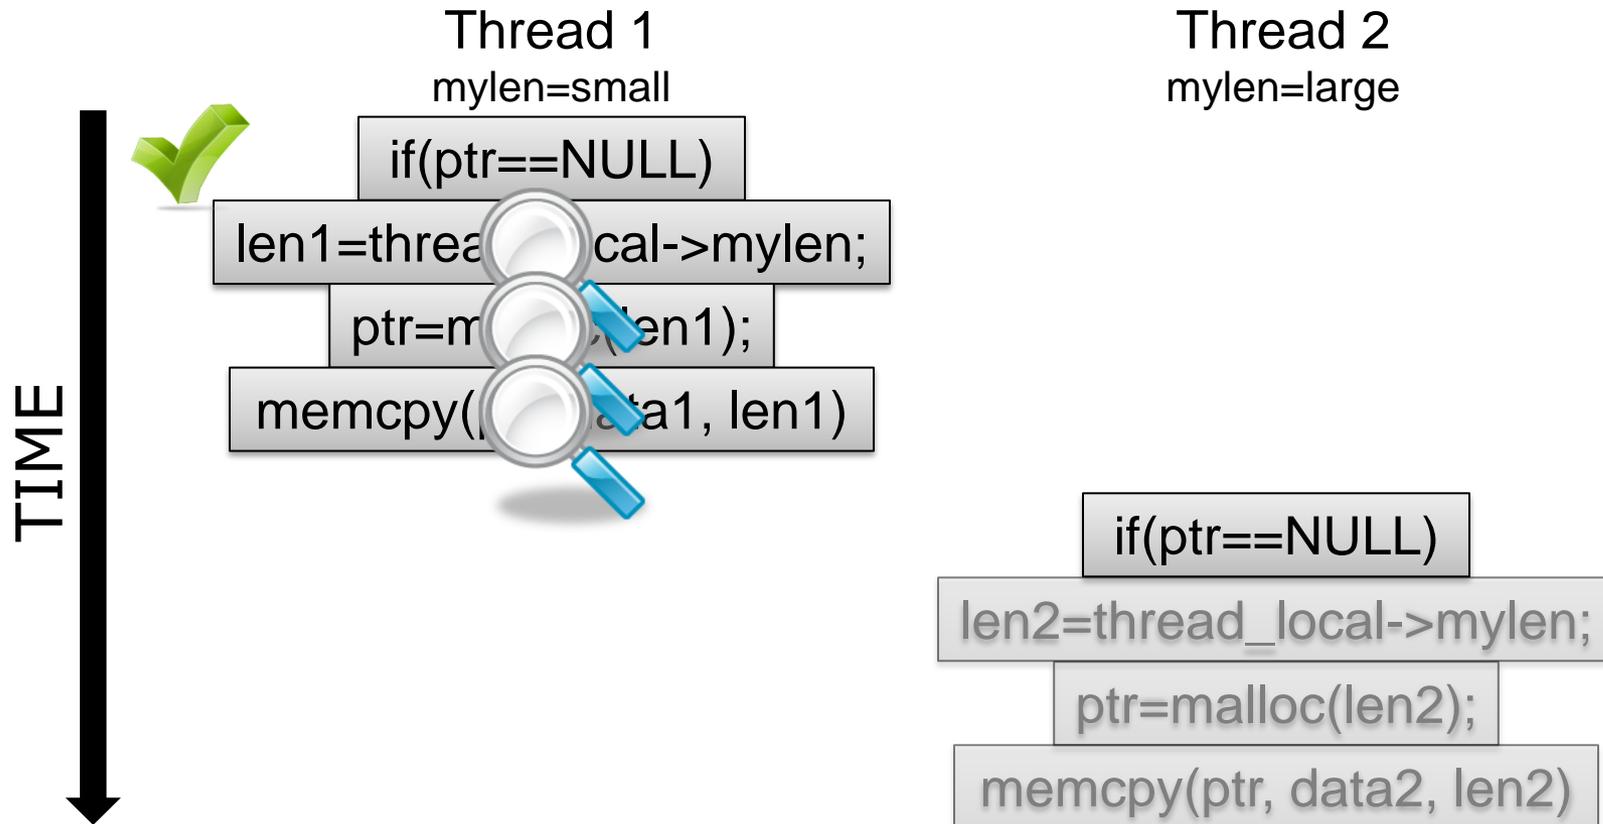
# Example of Data Race Detection



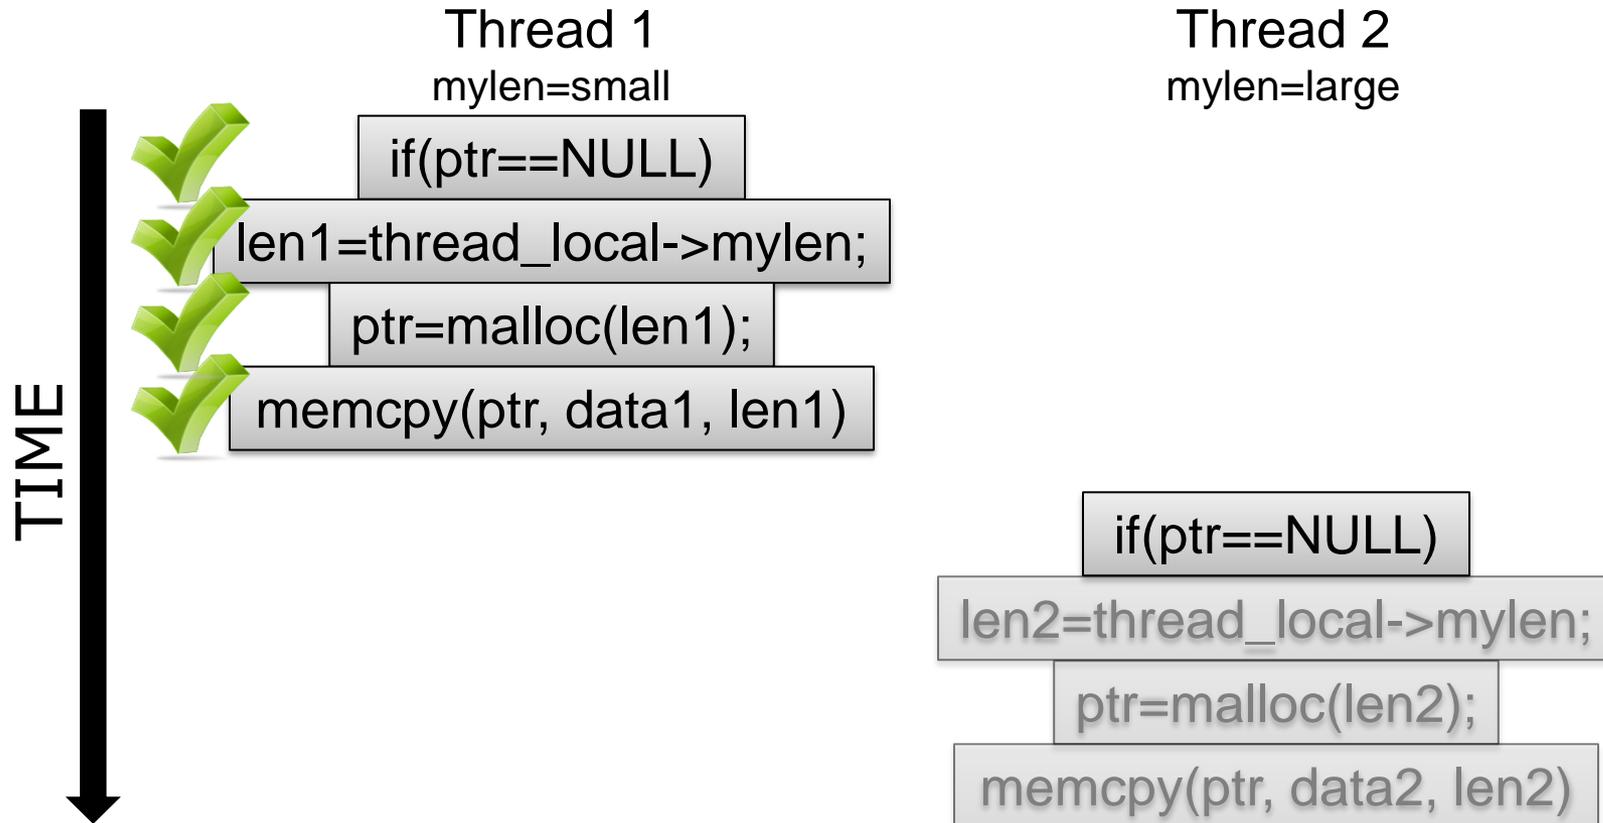
# Example of Data Race Detection



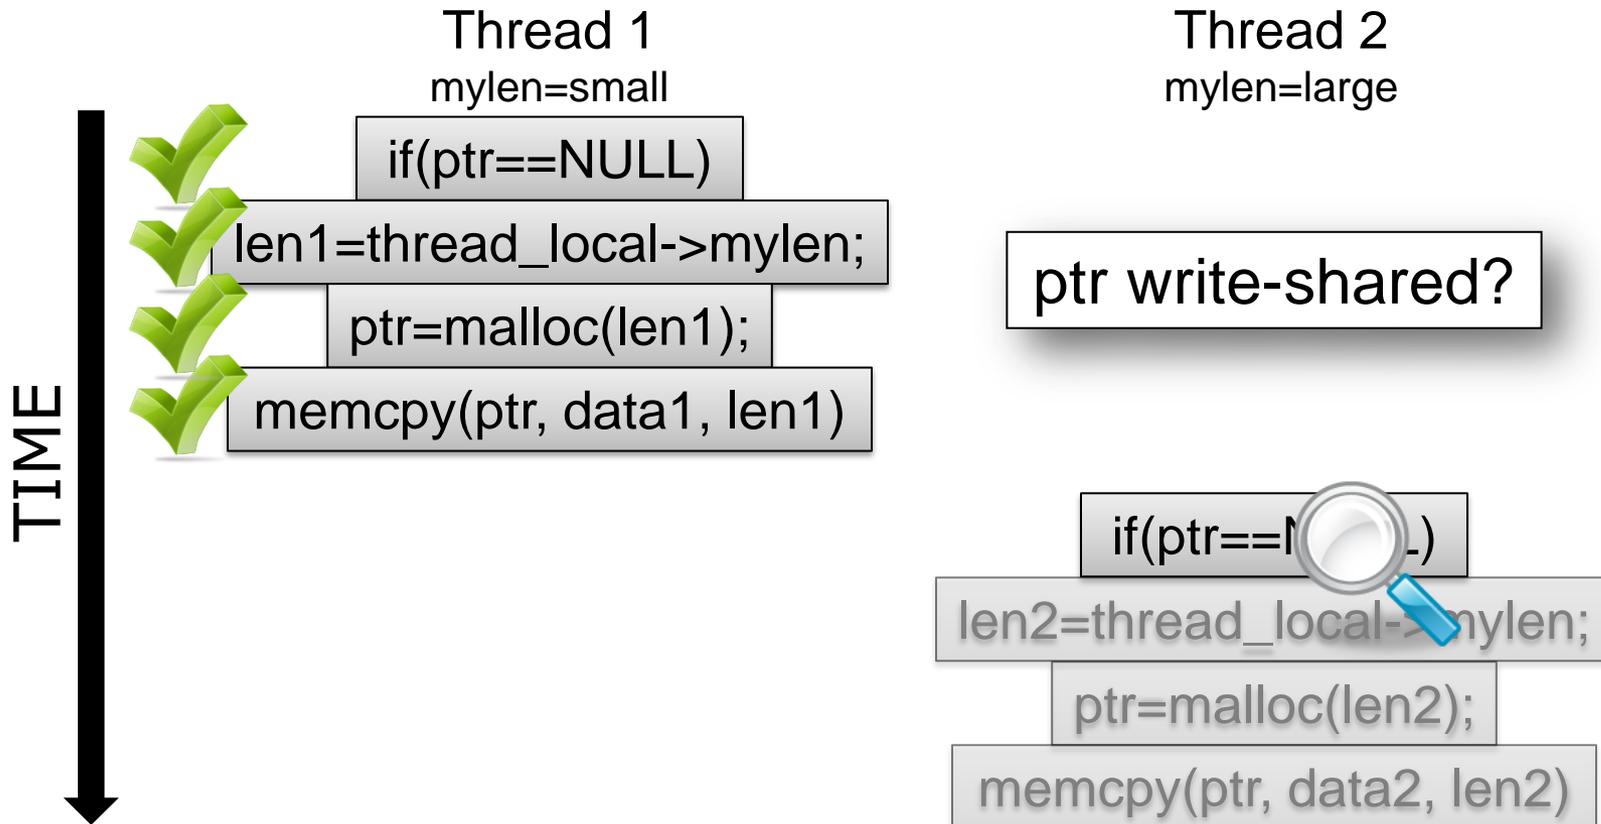
# Example of Data Race Detection



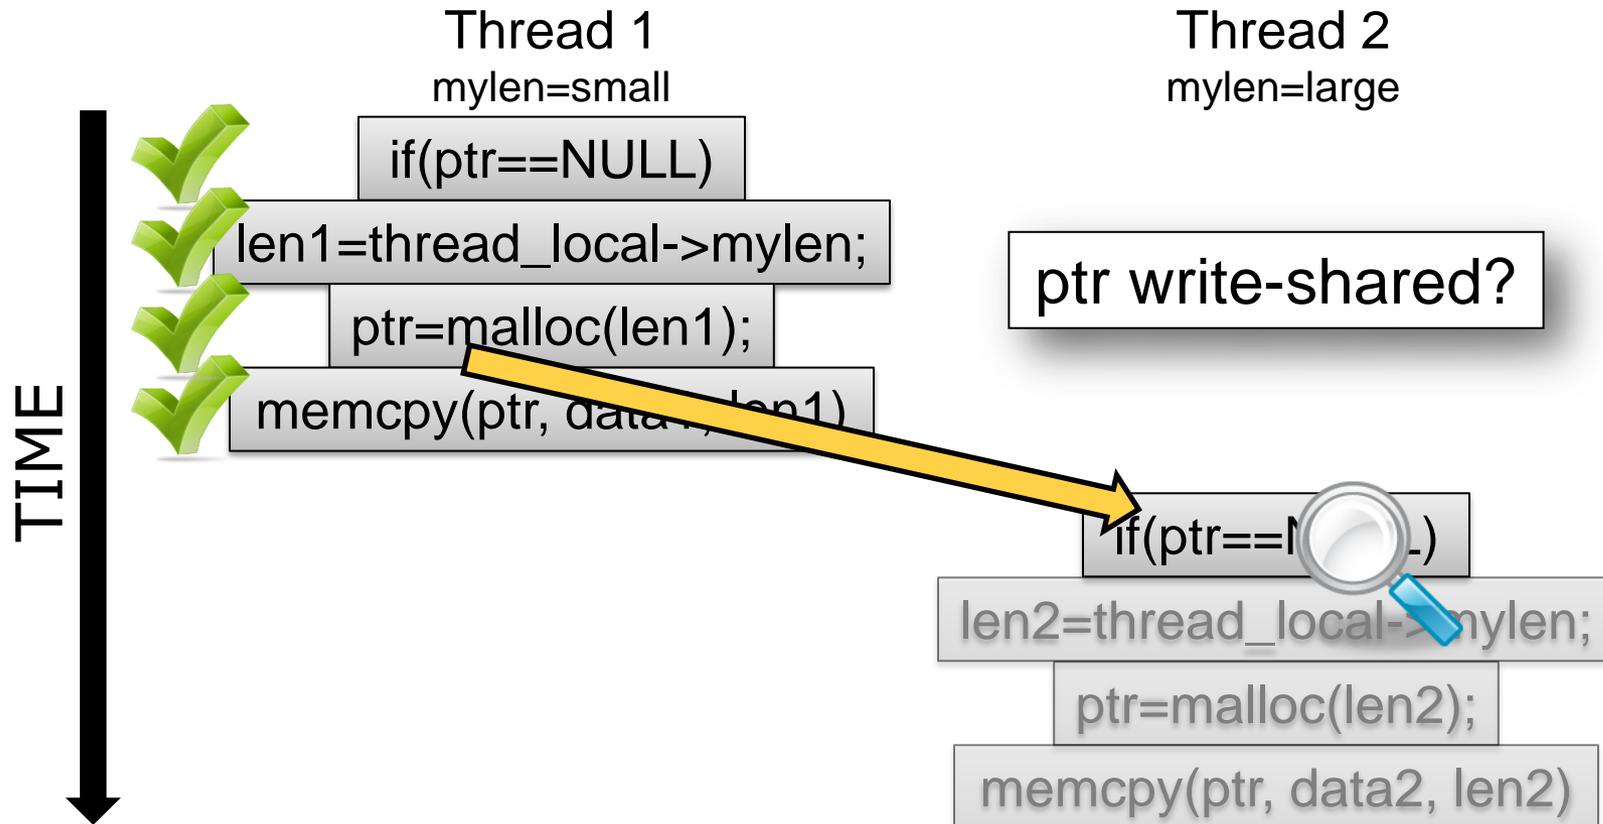
# Example of Data Race Detection



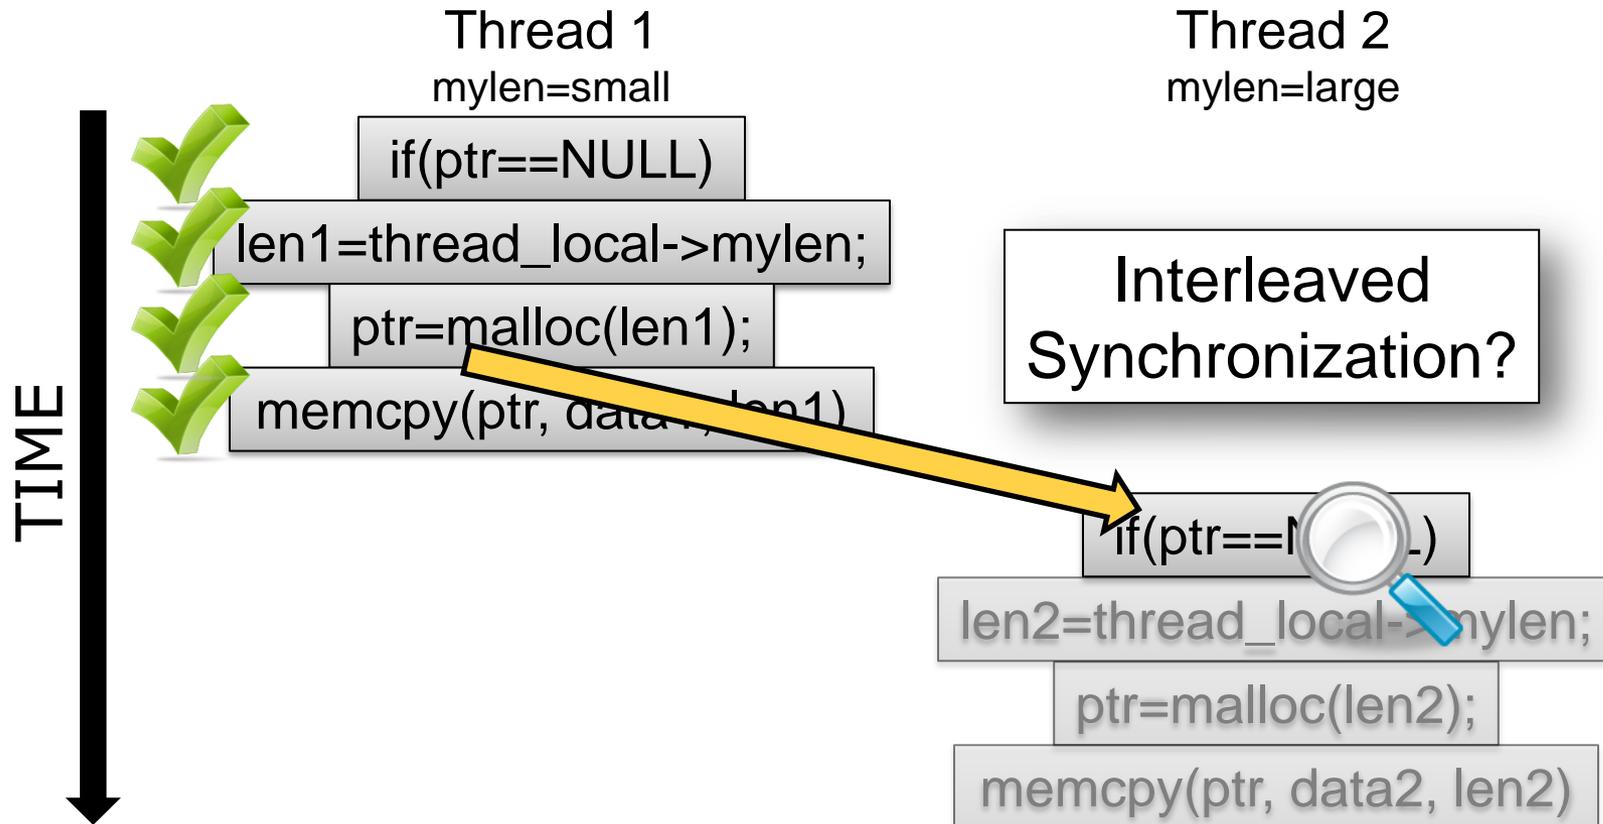
# Example of Data Race Detection



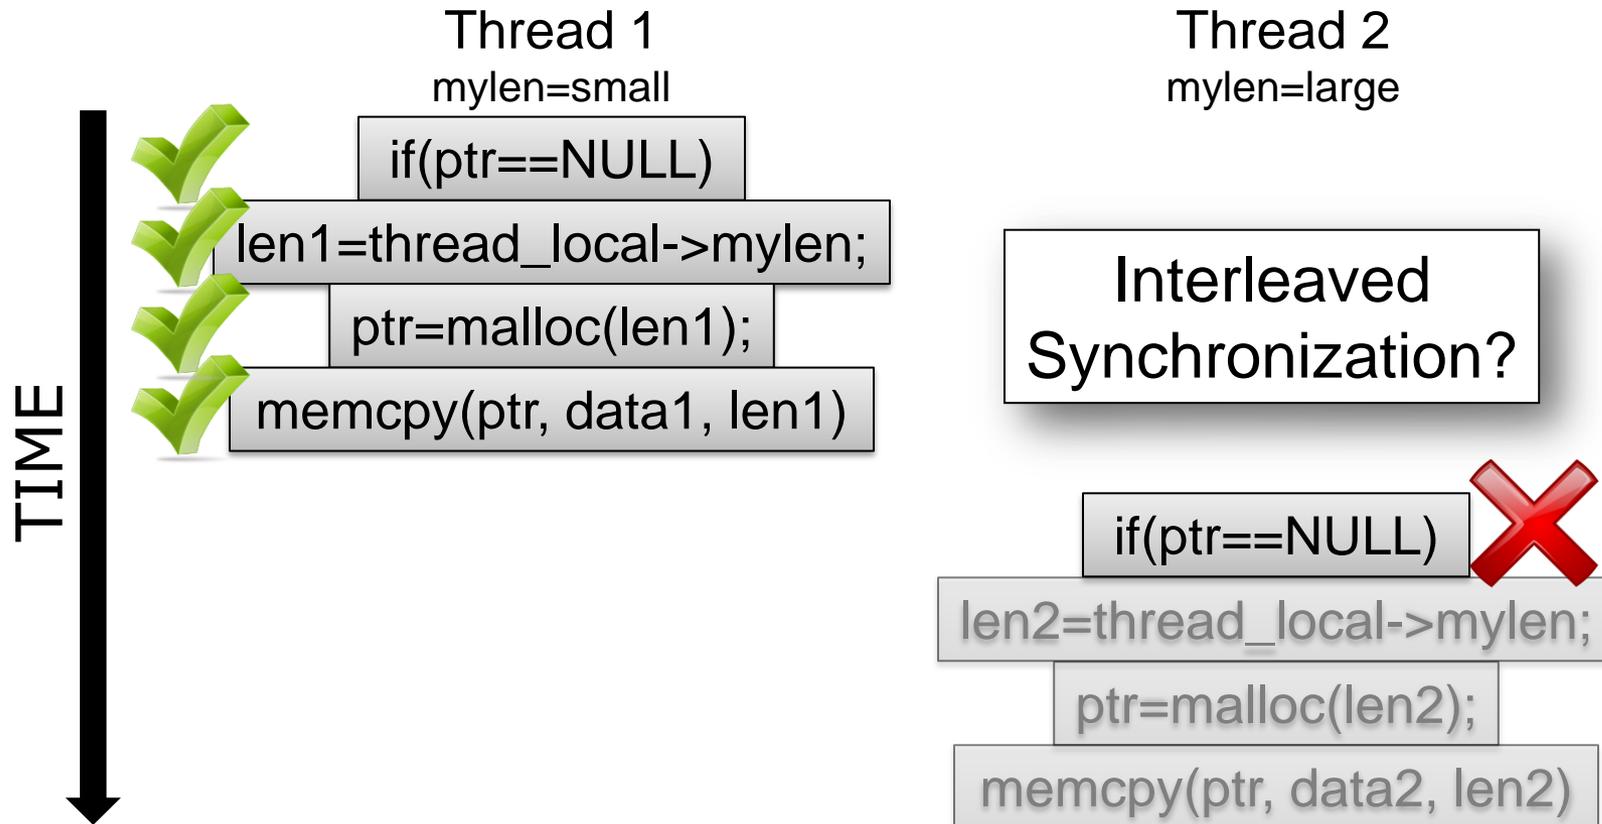
# Example of Data Race Detection



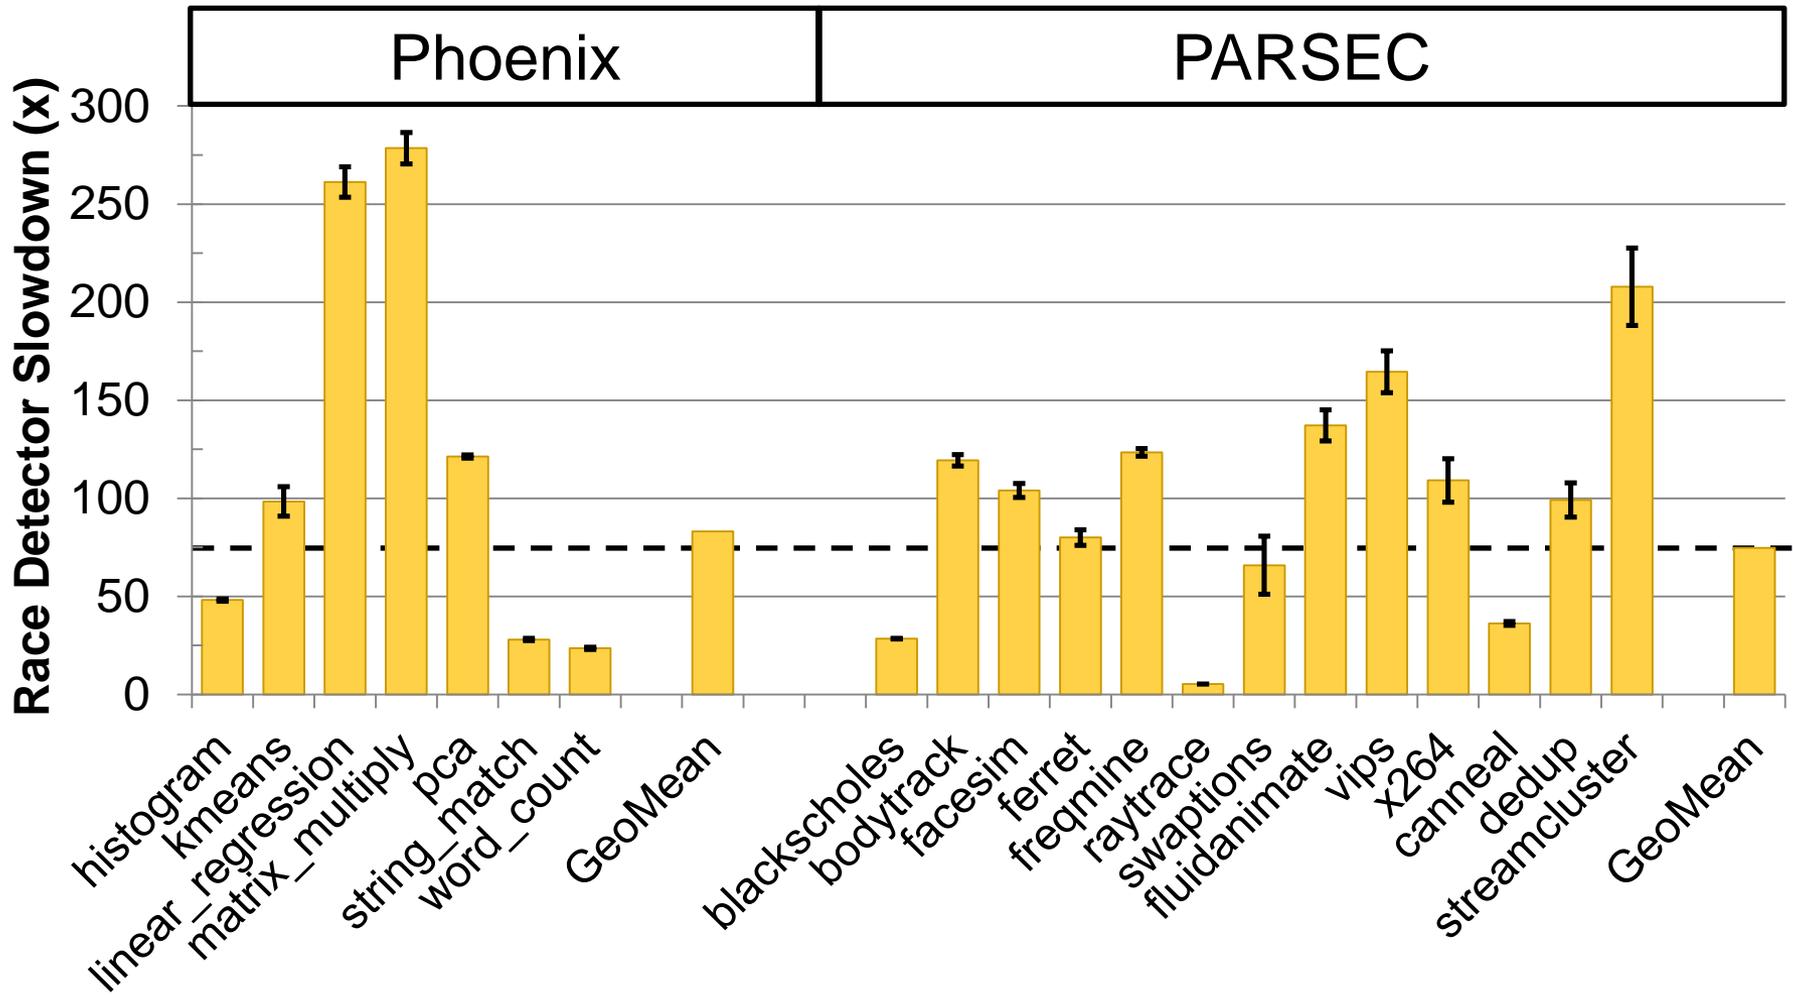
# Example of Data Race Detection



# Example of Data Race Detection

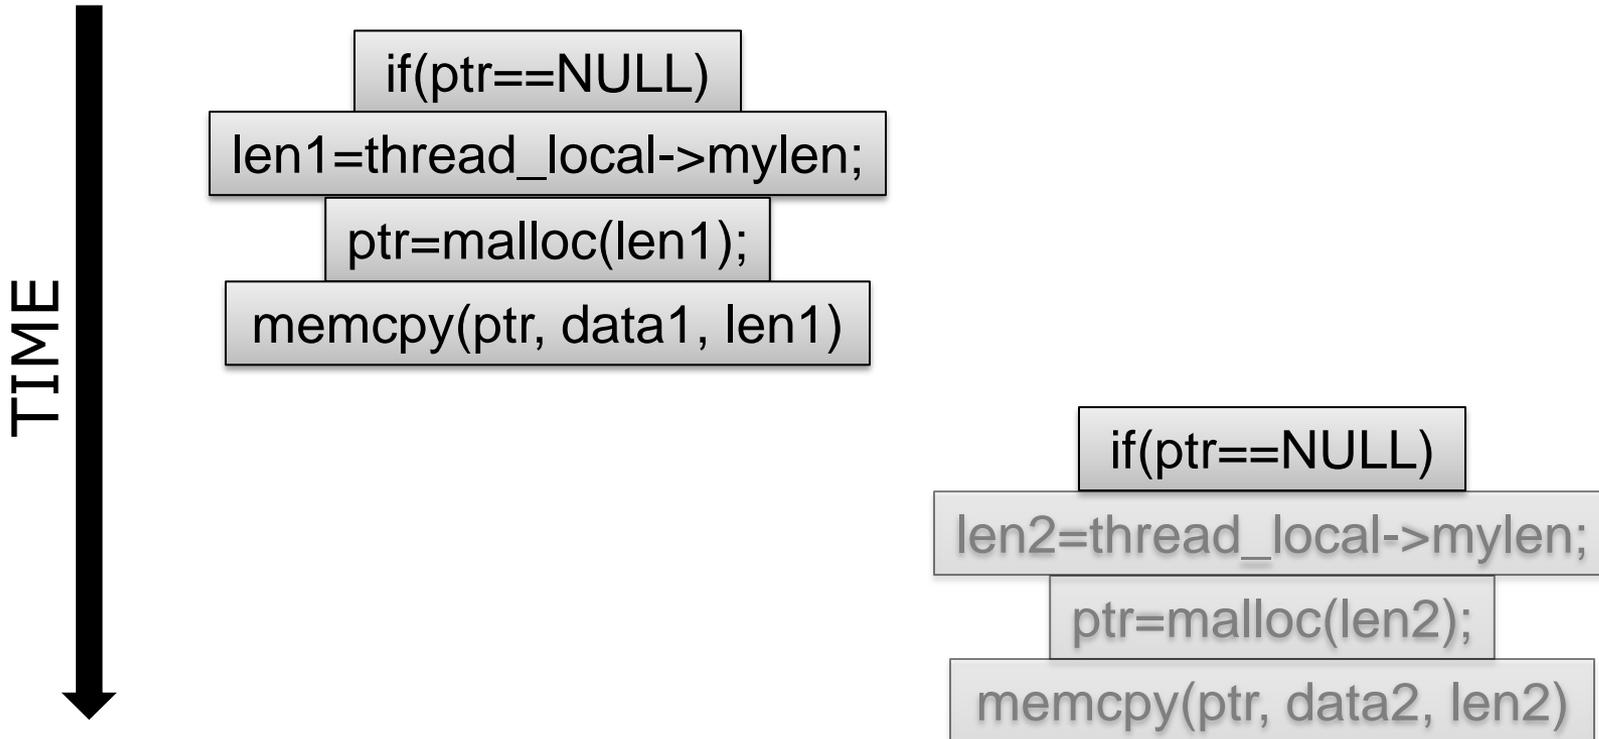


# SW Race Detection is Slow



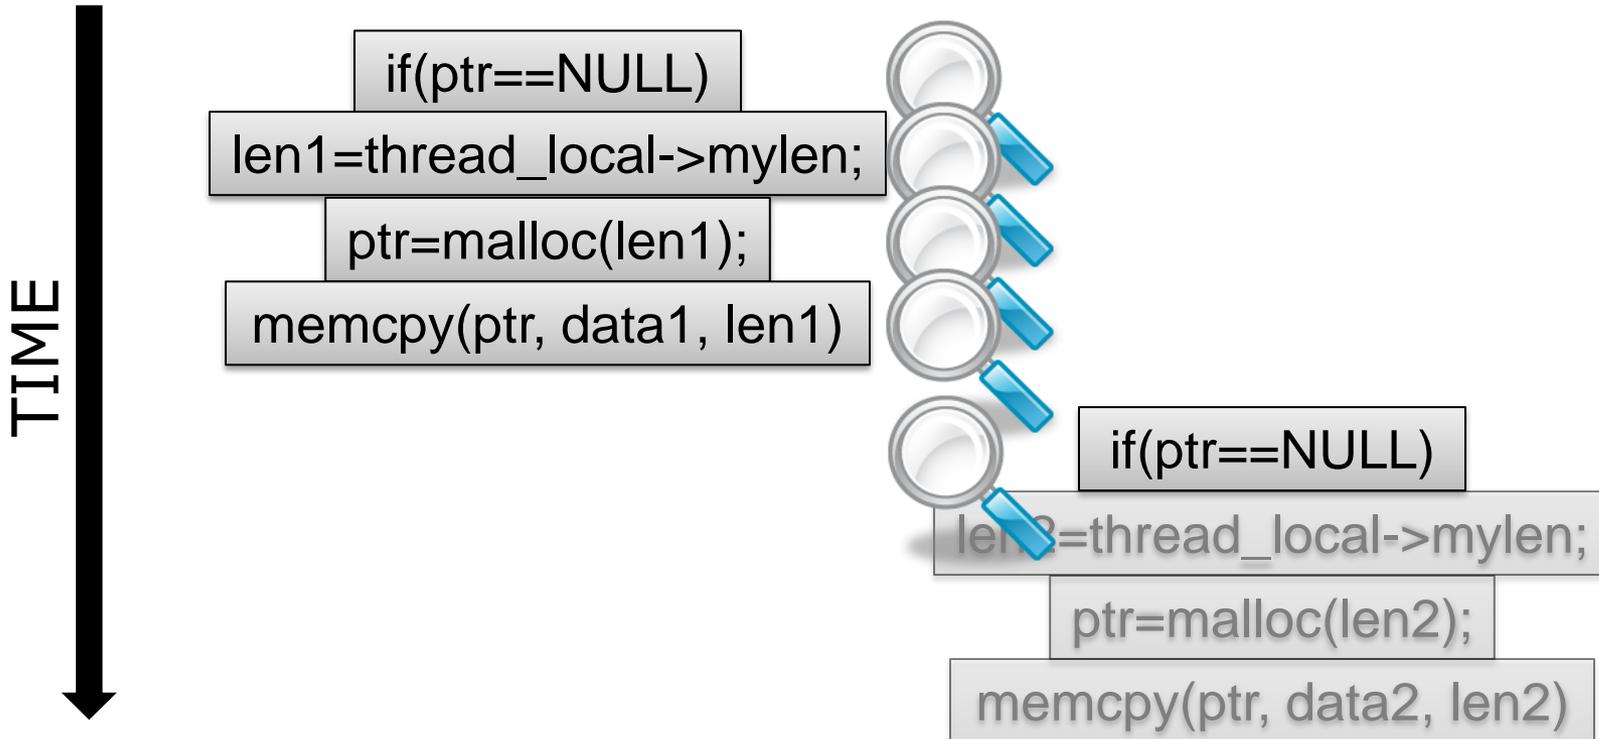
# Inter-thread Sharing is What's Important

“Data races ... are failures in programs that **access and update shared data** in critical sections” – Netzer & Miller, 1992



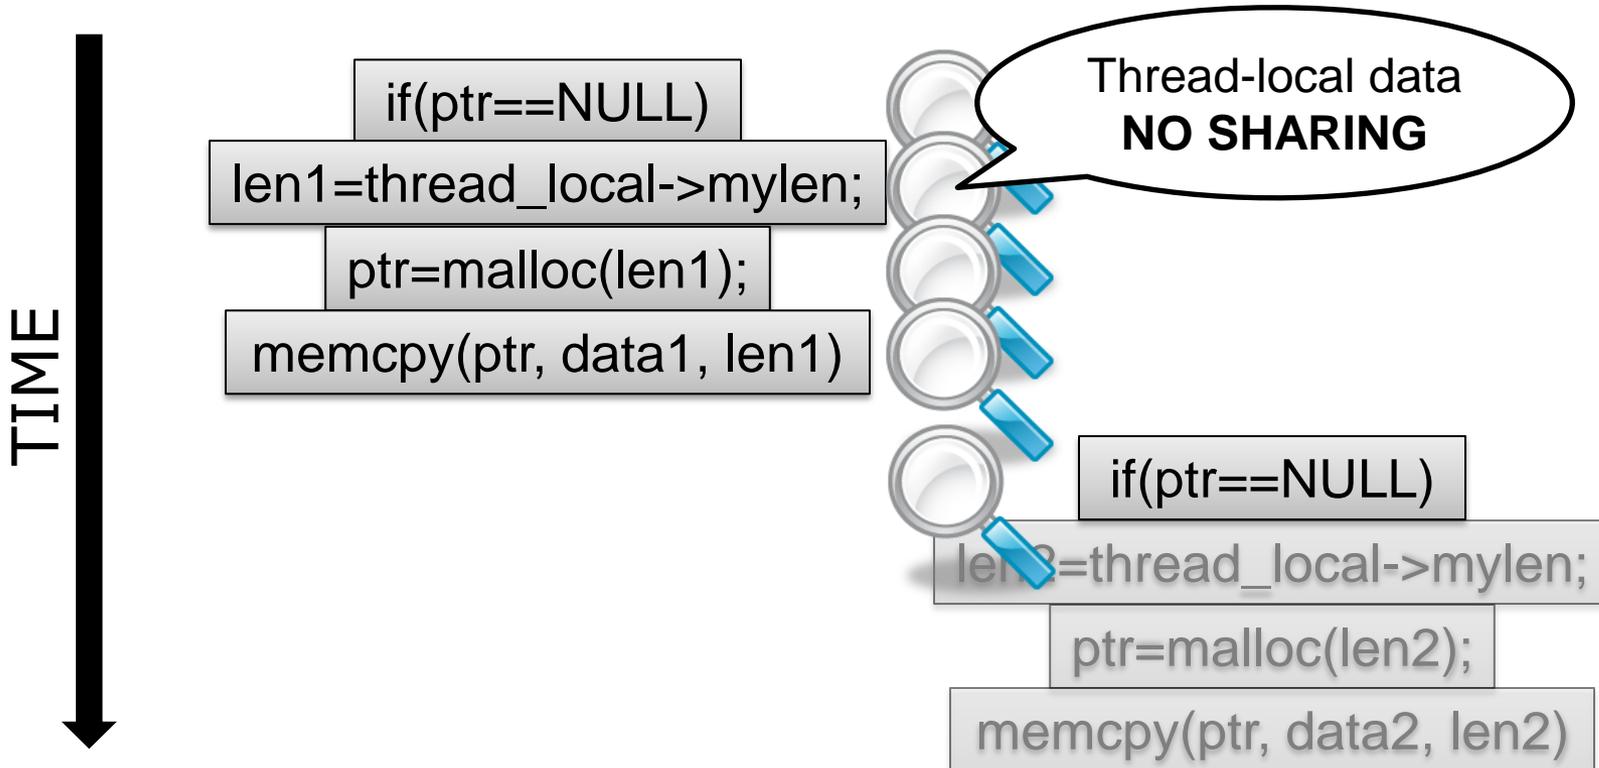
# Inter-thread Sharing is What's Important

“Data races ... are failures in programs that **access and update shared data** in critical sections” – Netzer & Miller, 1992



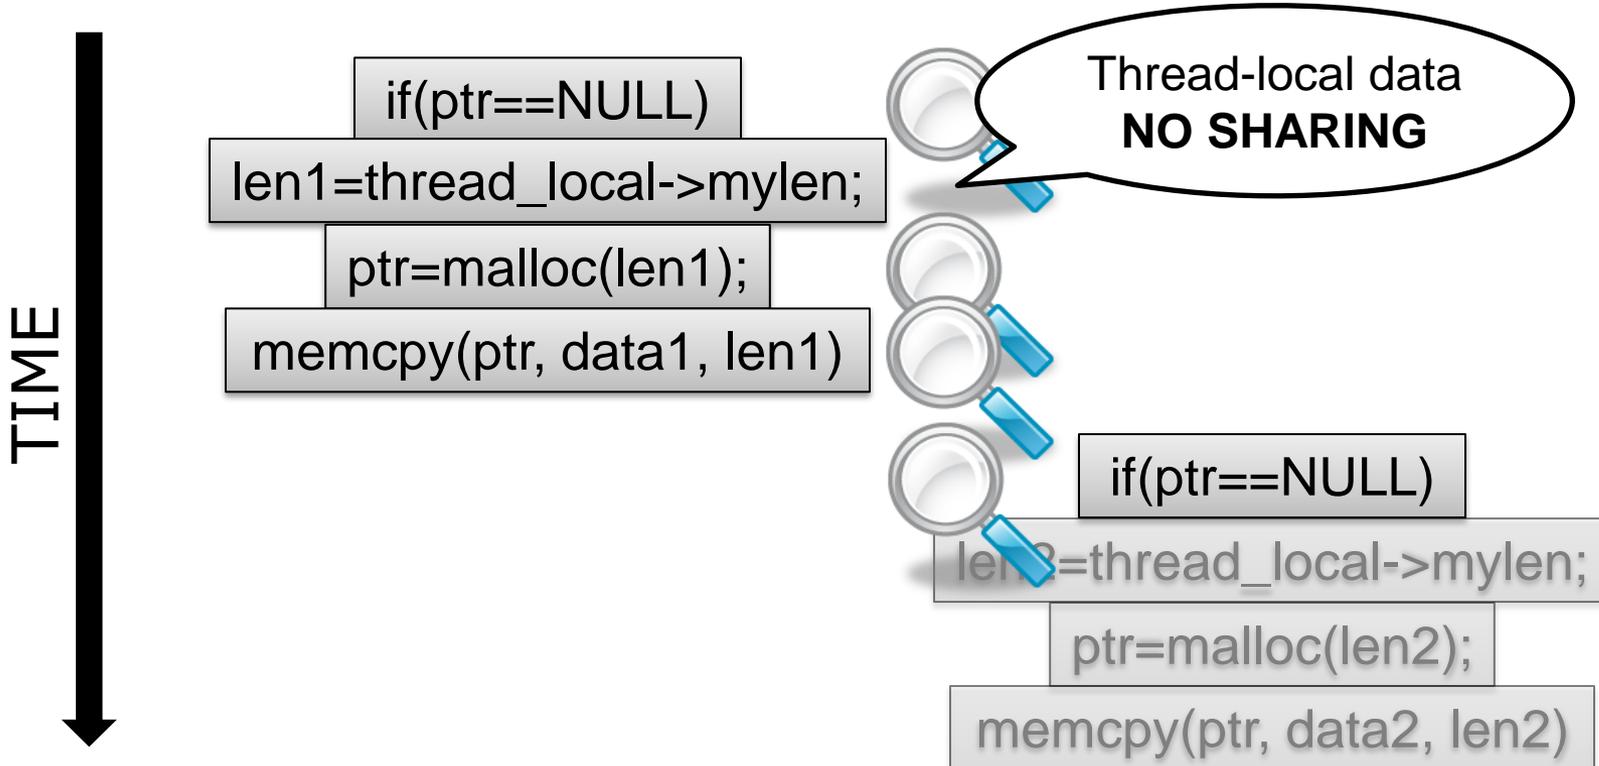
# Inter-thread Sharing is What's Important

“Data races ... are failures in programs that **access and update shared data** in critical sections” – Netzer & Miller, 1992



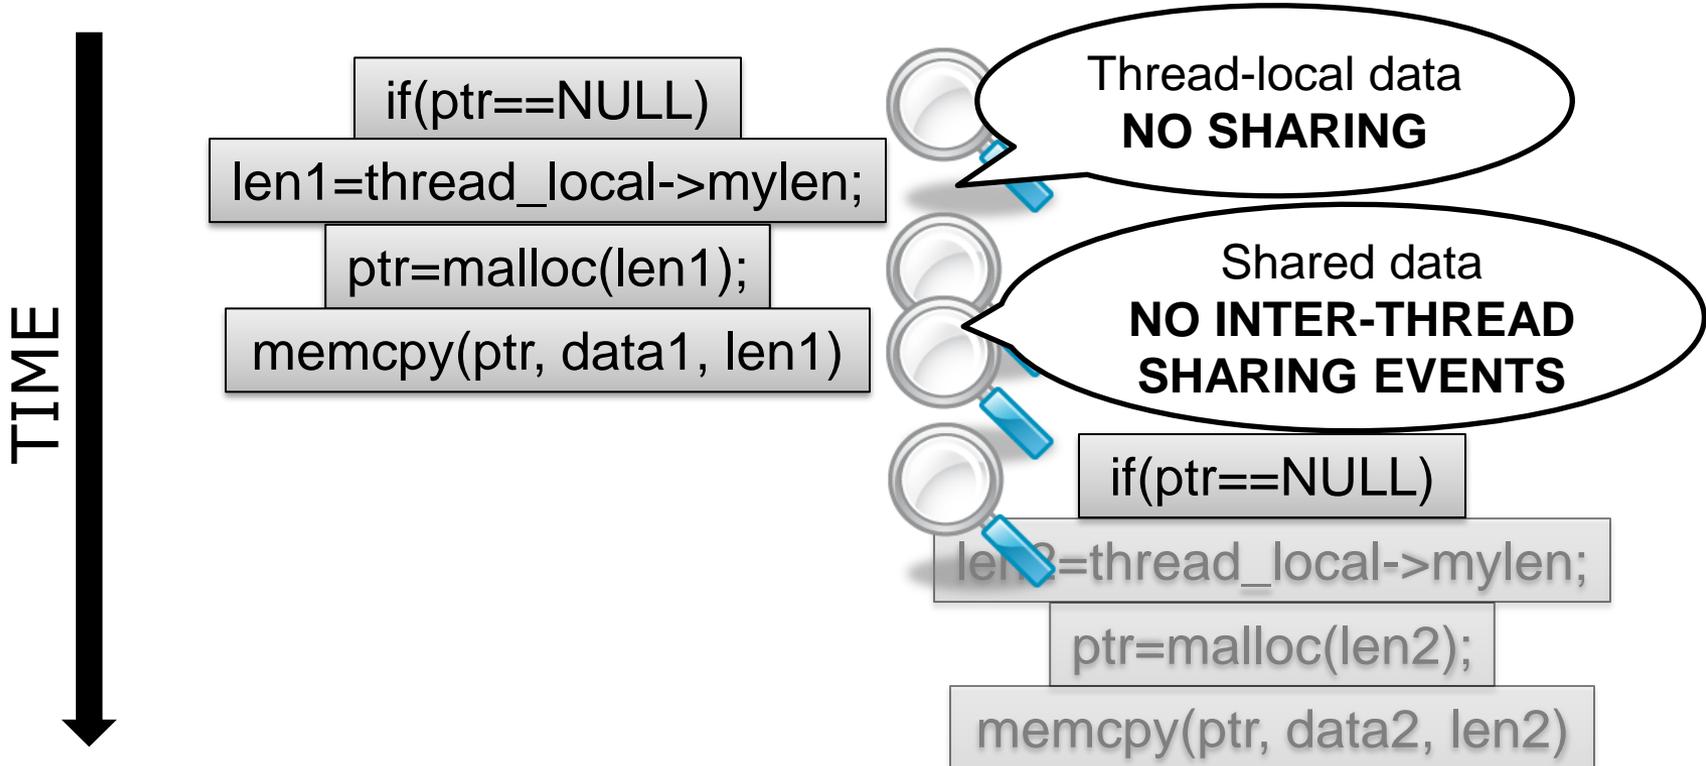
# Inter-thread Sharing is What's Important

“Data races ... are failures in programs that **access and update shared data** in critical sections” – Netzer & Miller, 1992



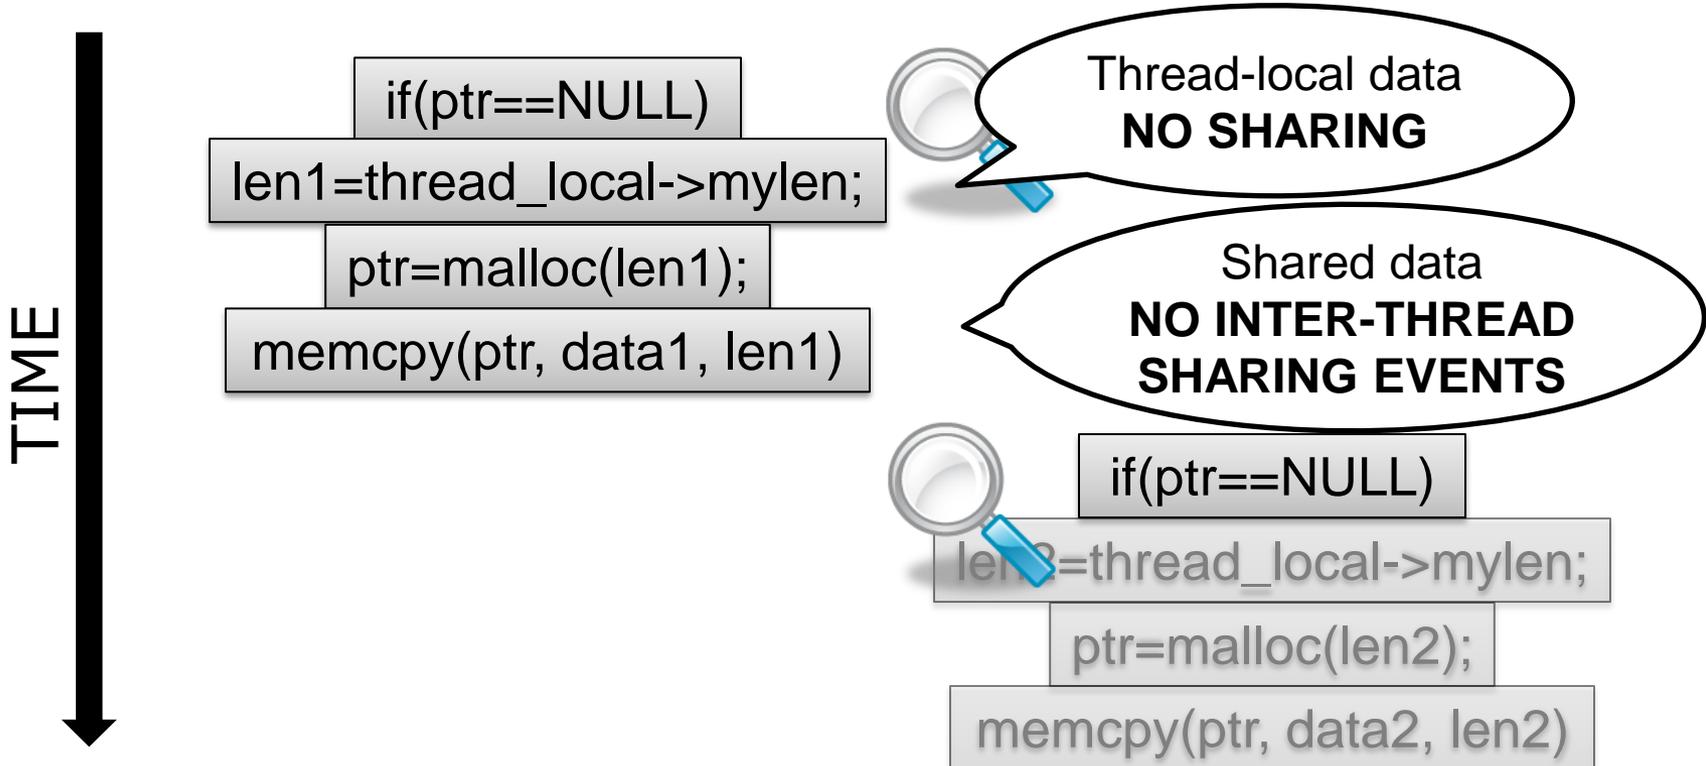
# Inter-thread Sharing is What's Important

“Data races ... are failures in programs that **access and update shared data** in critical sections” – Netzer & Miller, 1992



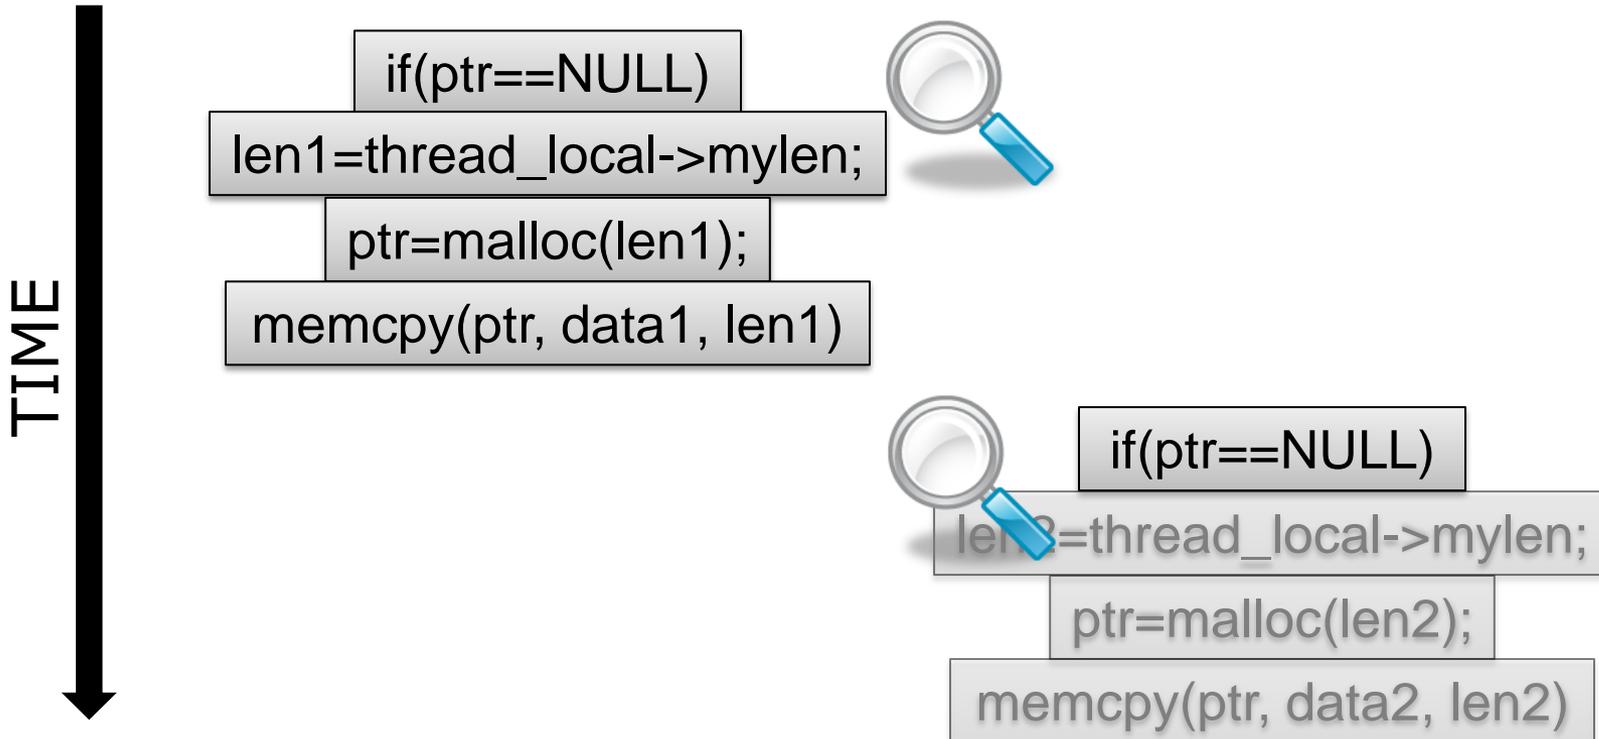
# Inter-thread Sharing is What's Important

“Data races ... are failures in programs that **access and update shared data** in critical sections” – Netzer & Miller, 1992



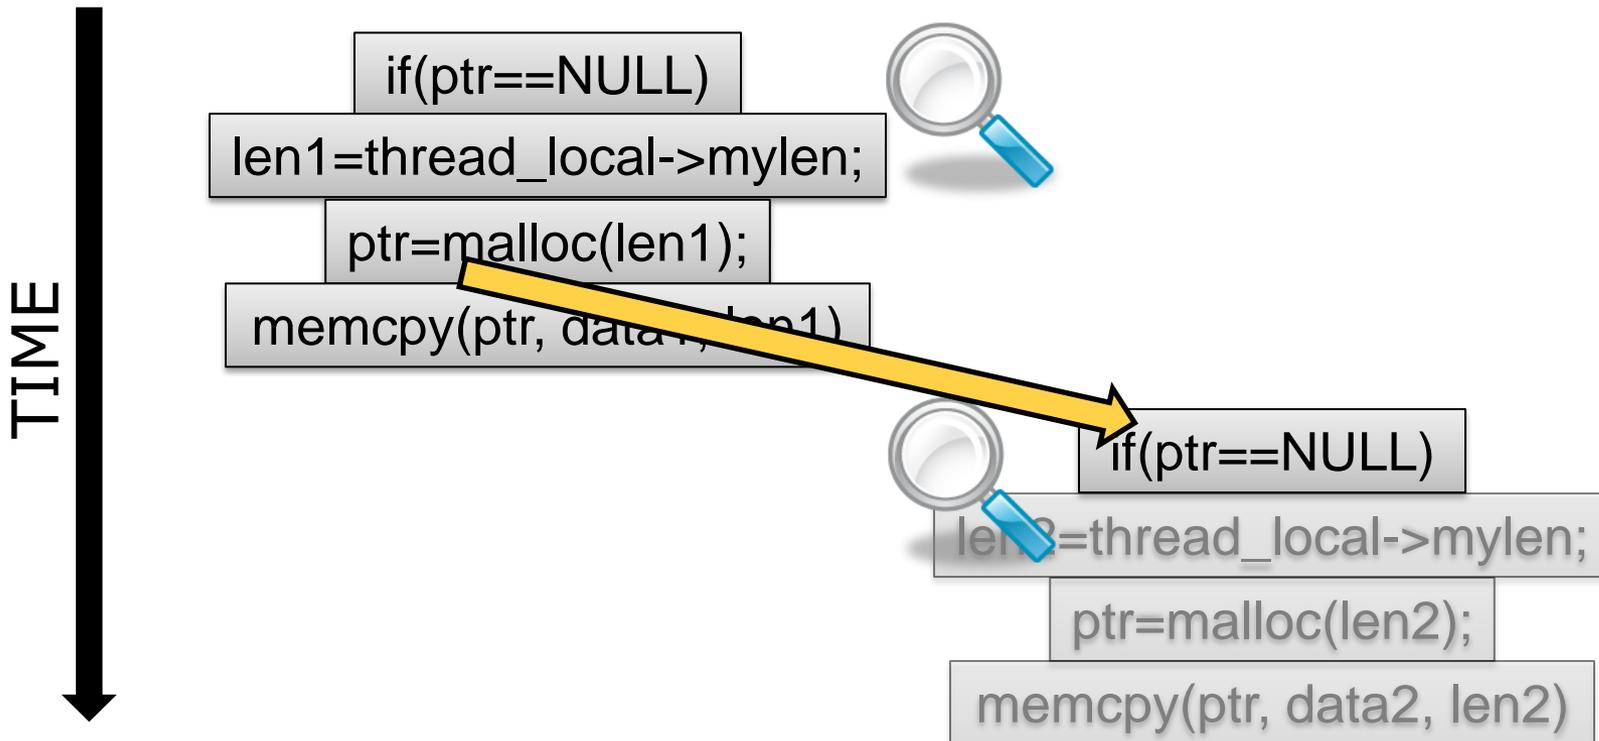
# Inter-thread Sharing is What's Important

“Data races ... are failures in programs that **access and update shared data** in critical sections” – Netzer & Miller, 1992

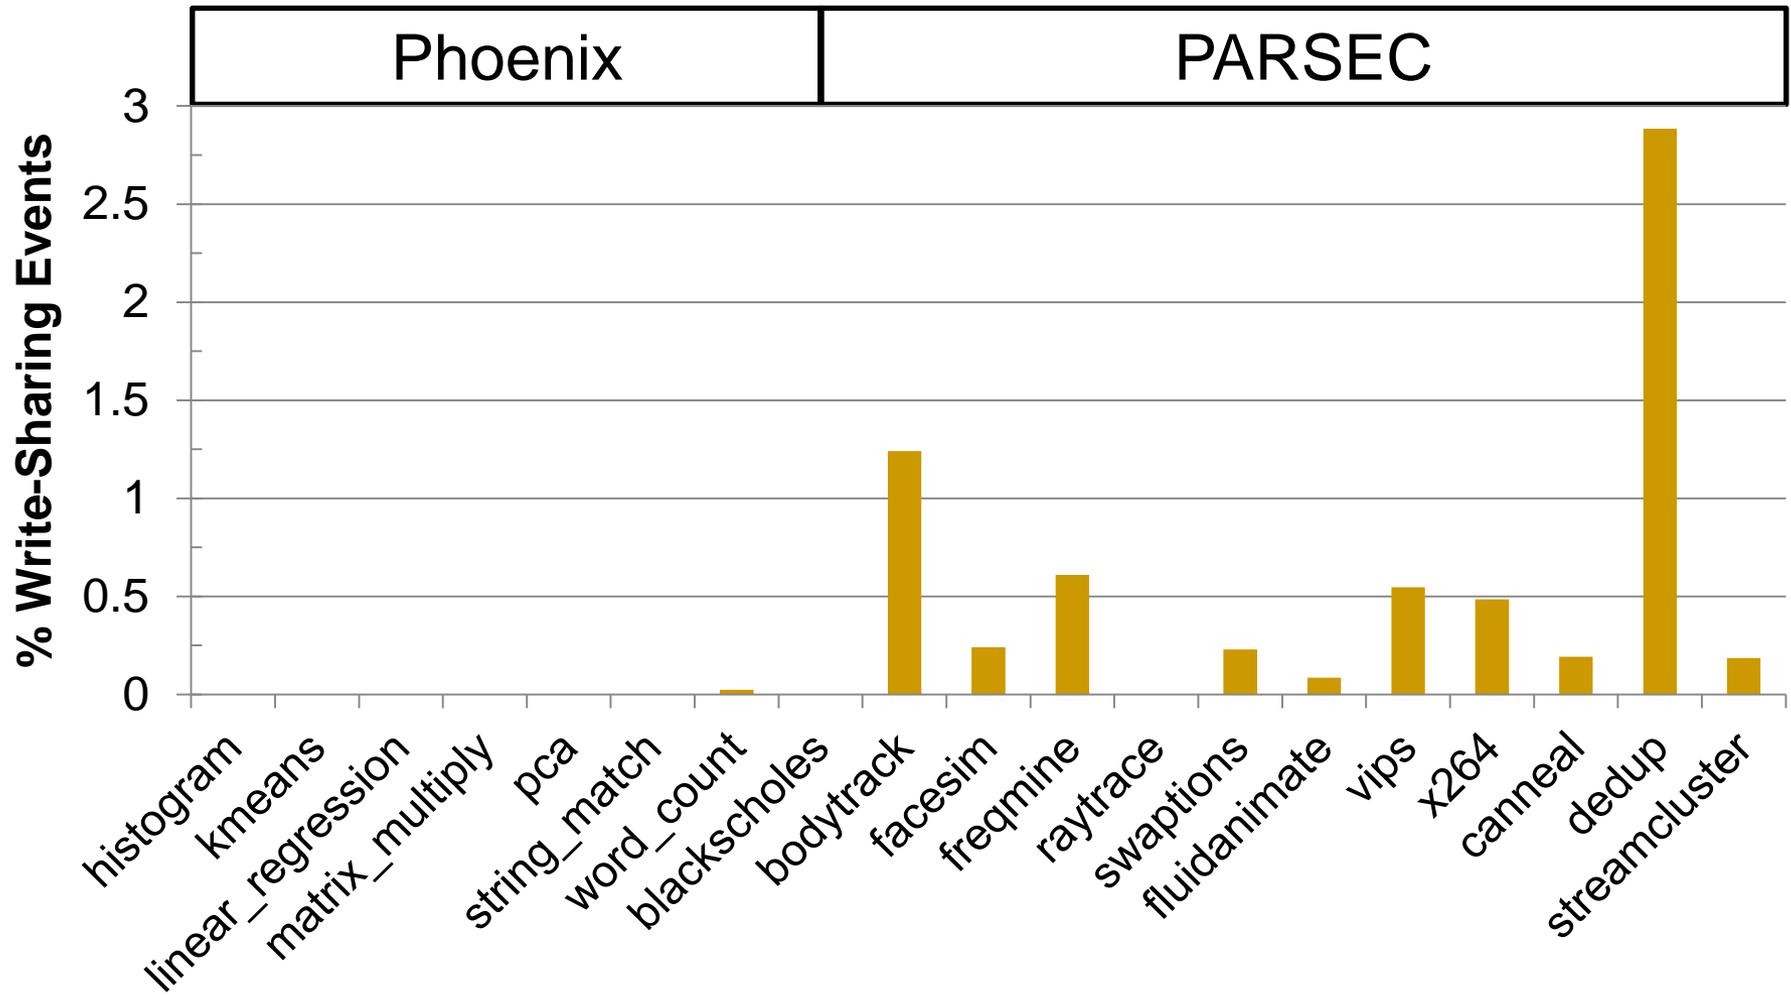


# Inter-thread Sharing is What's Important

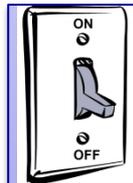
“Data races ... are failures in programs that **access and update shared data** in critical sections” – Netzer & Miller, 1992



# Very Little Inter-Thread Sharing



# Use Demand-Driven Analysis!

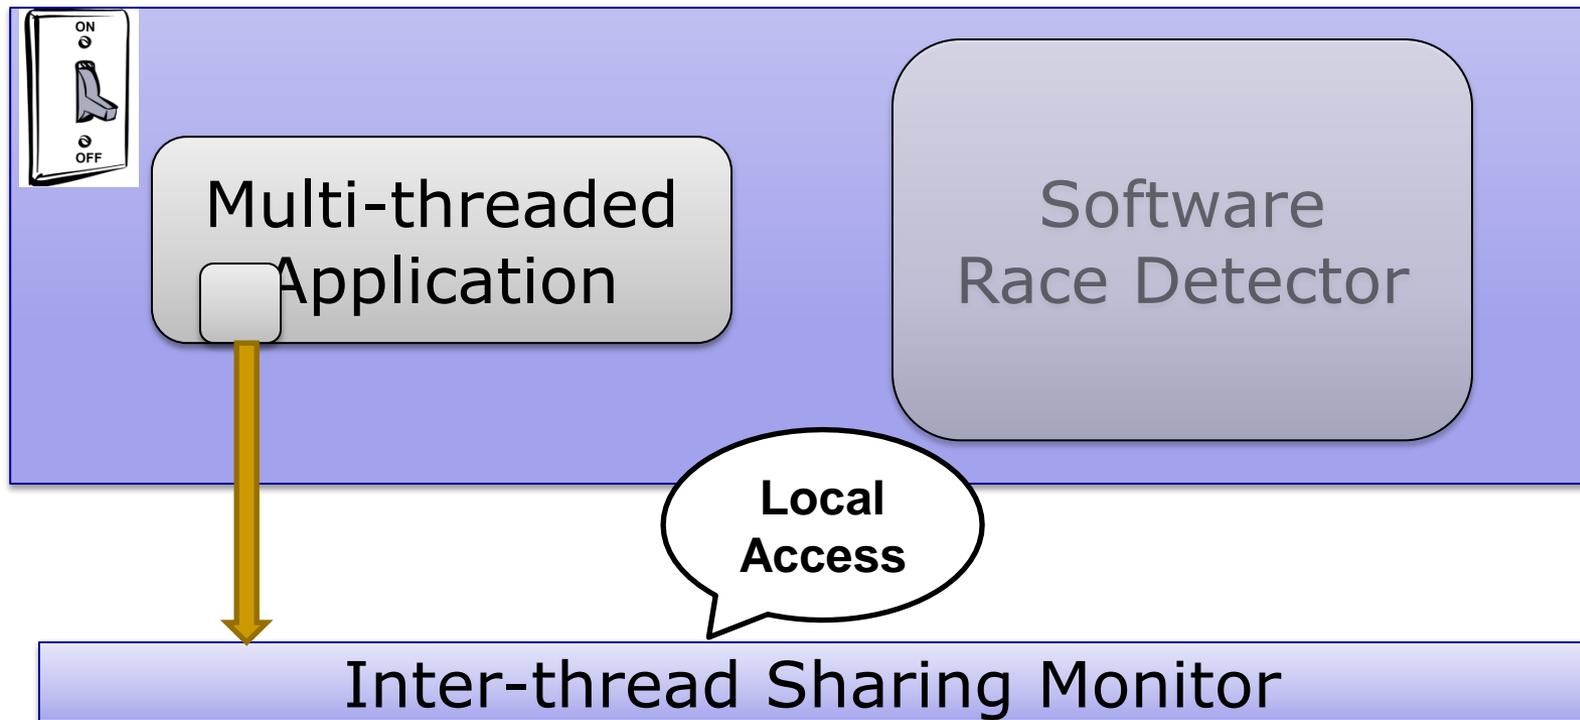


Multi-threaded  
Application

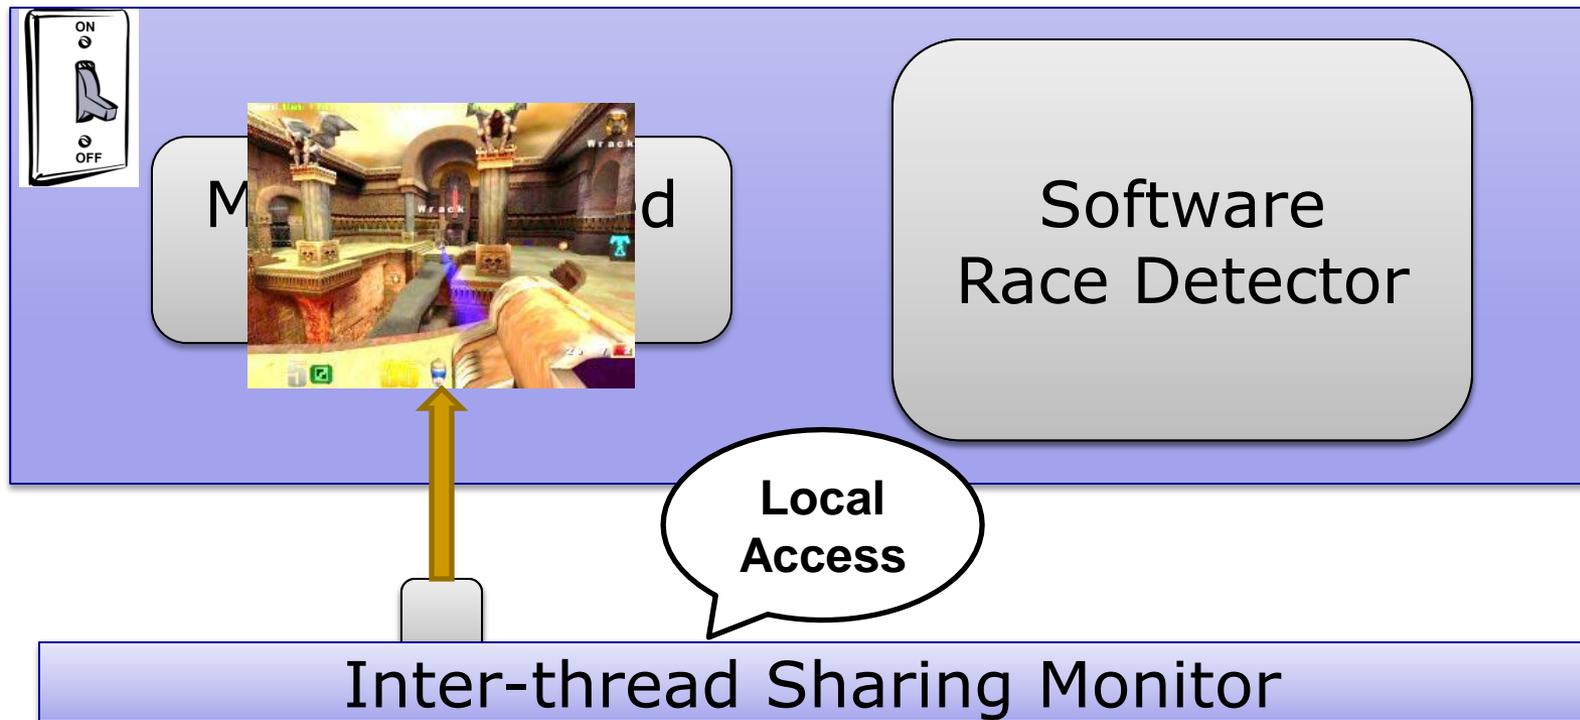
Software  
Race Detector

Inter-thread Sharing Monitor

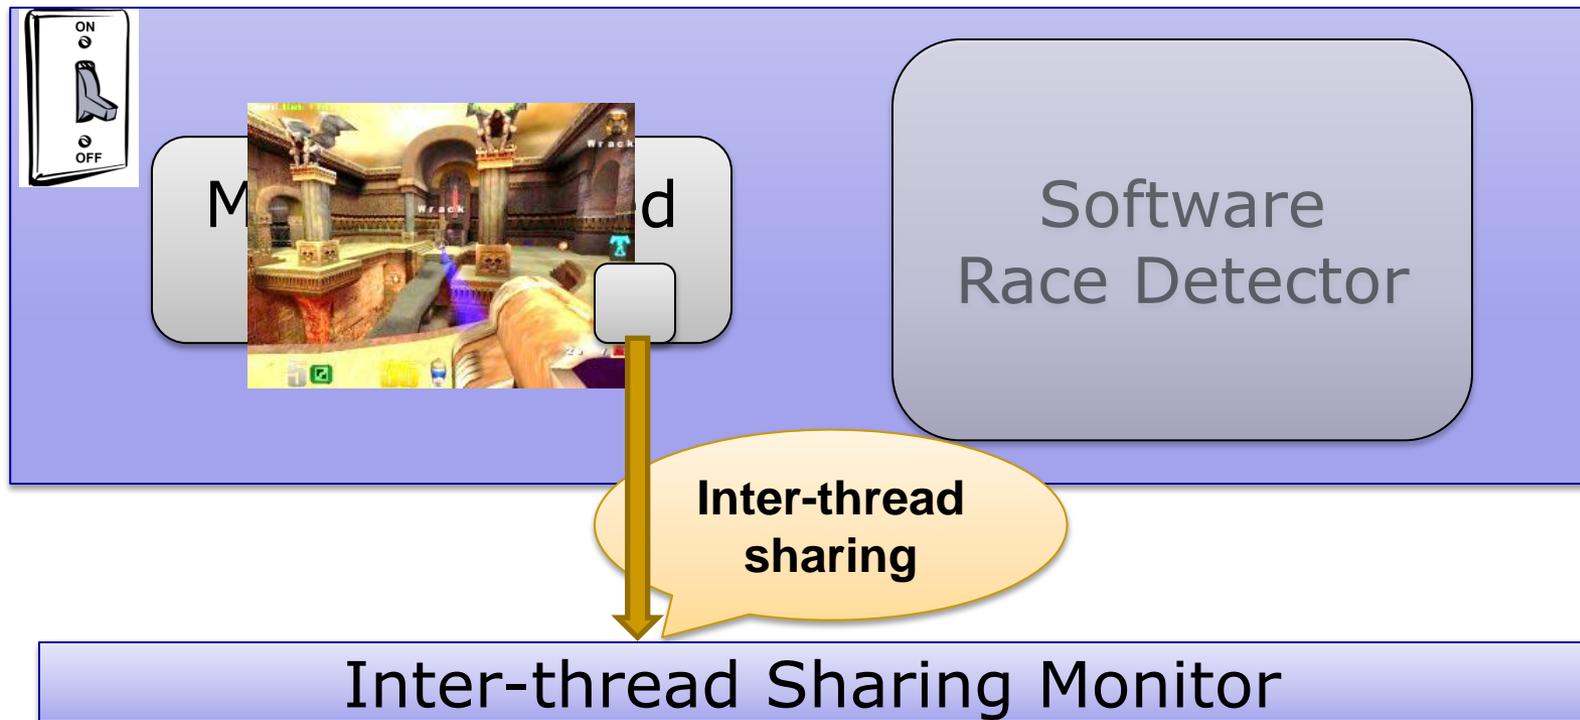
# Use Demand-Driven Analysis!



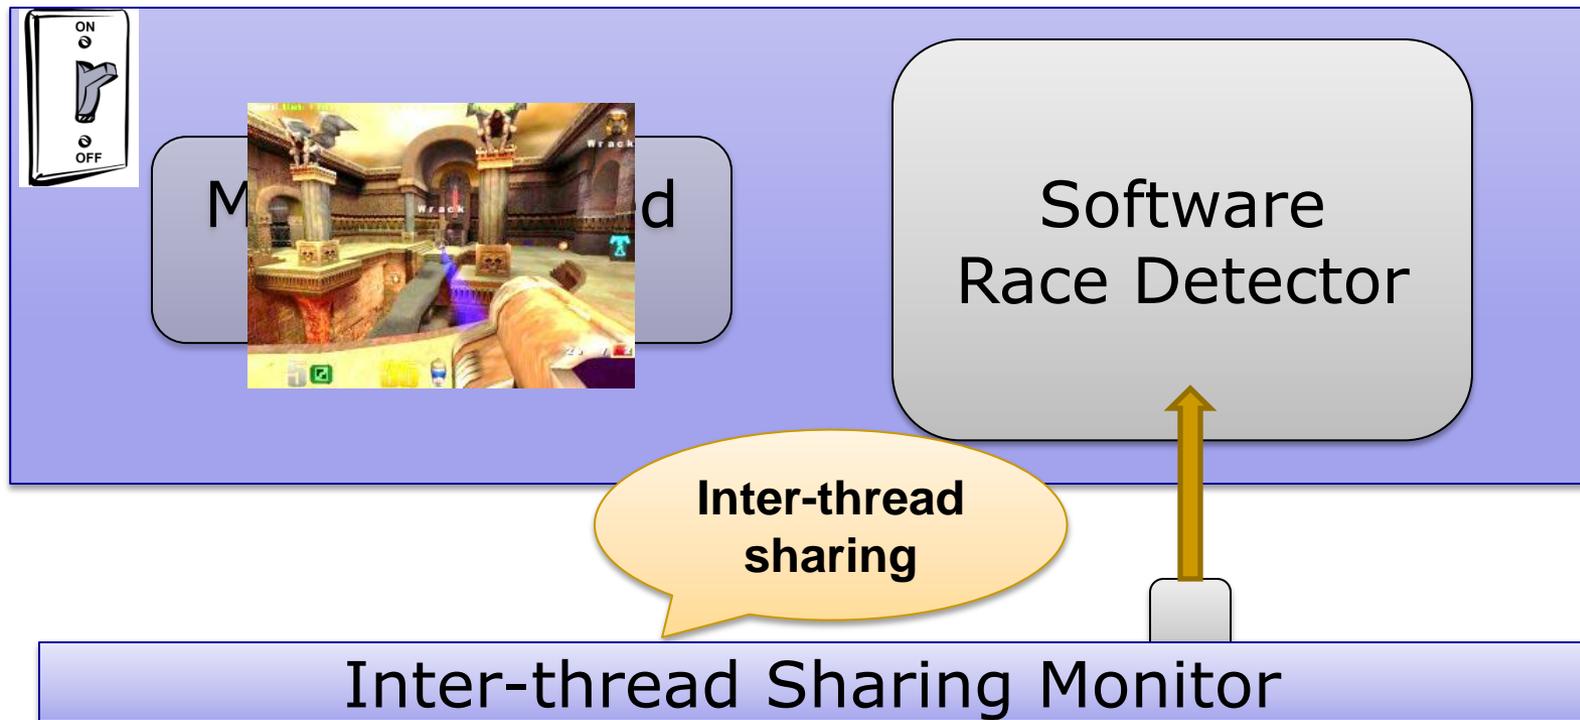
# Use Demand-Driven Analysis!



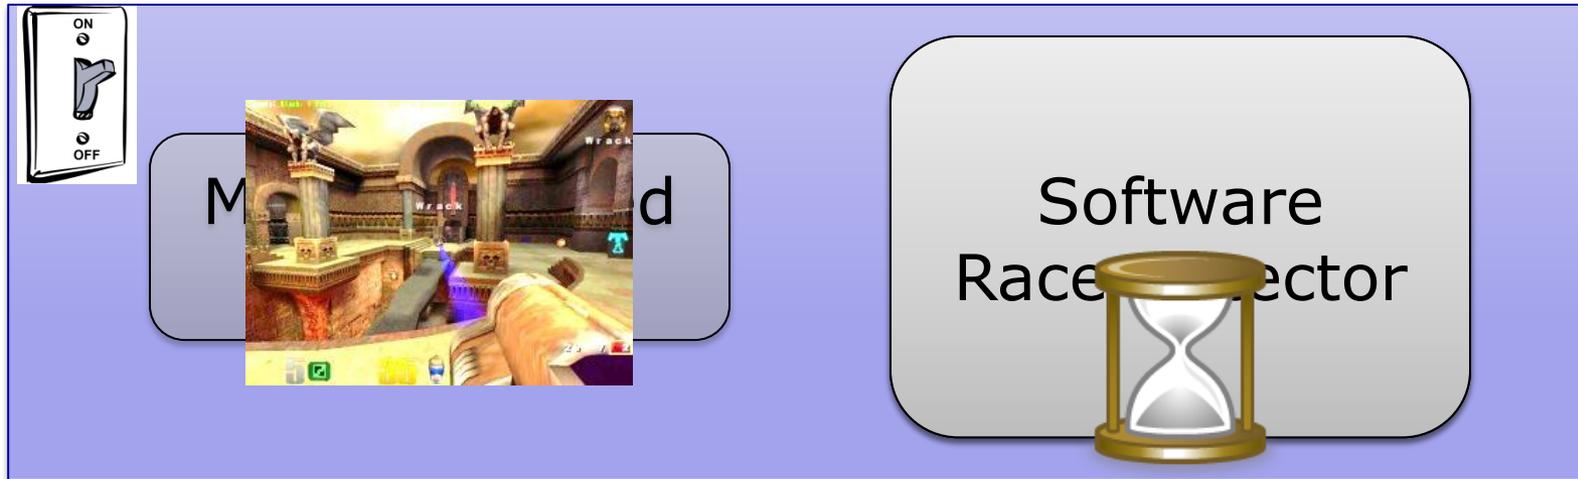
# Use Demand-Driven Analysis!



# Use Demand-Driven Analysis!

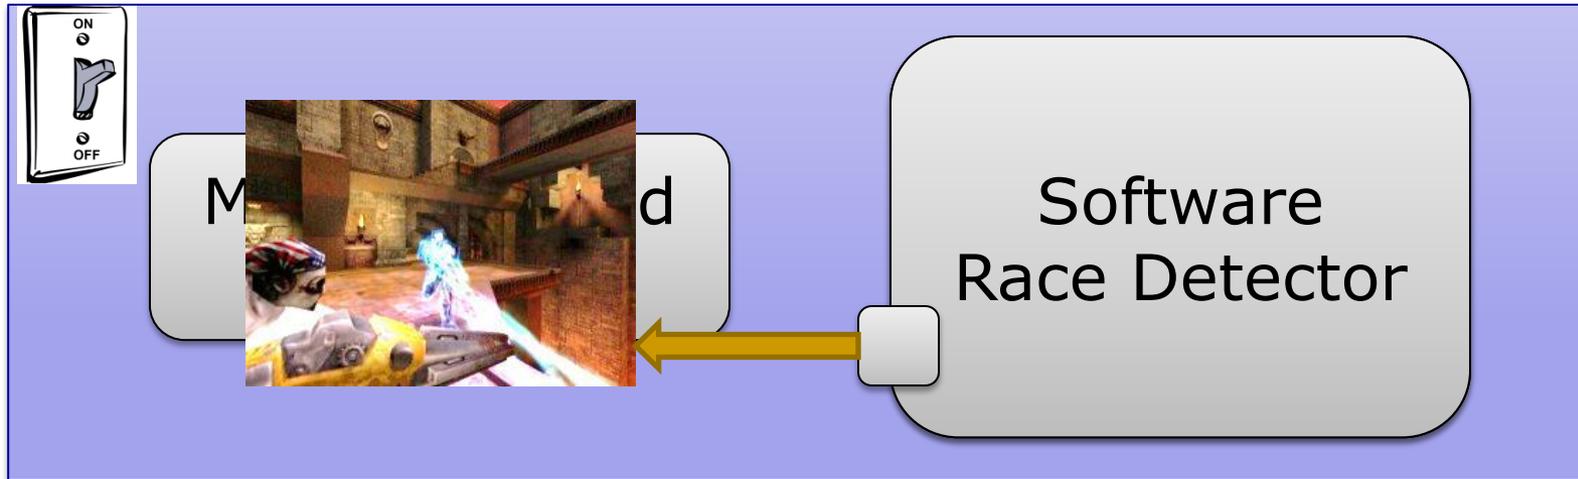


# Use Demand-Driven Analysis!



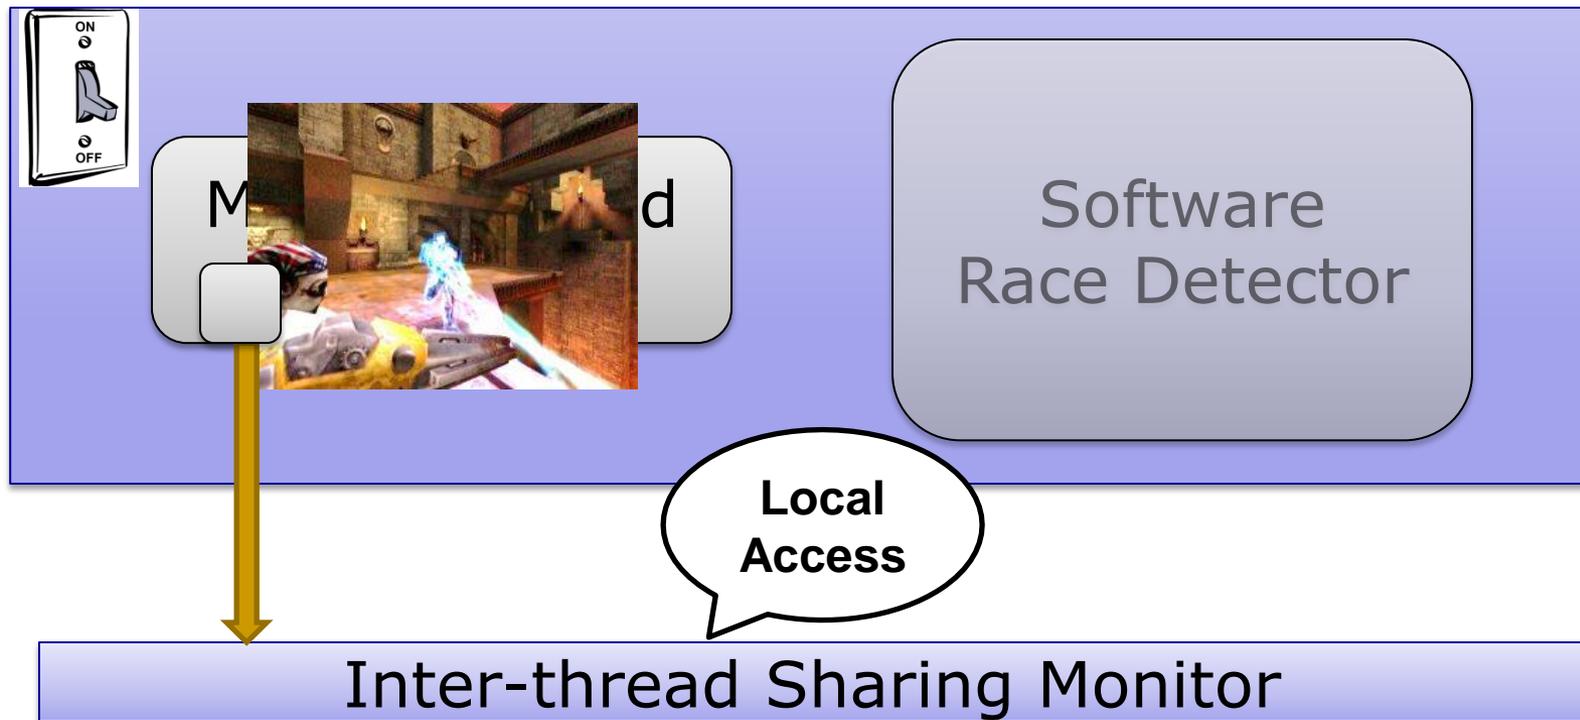
Inter-thread Sharing Monitor

# Use Demand-Driven Analysis!

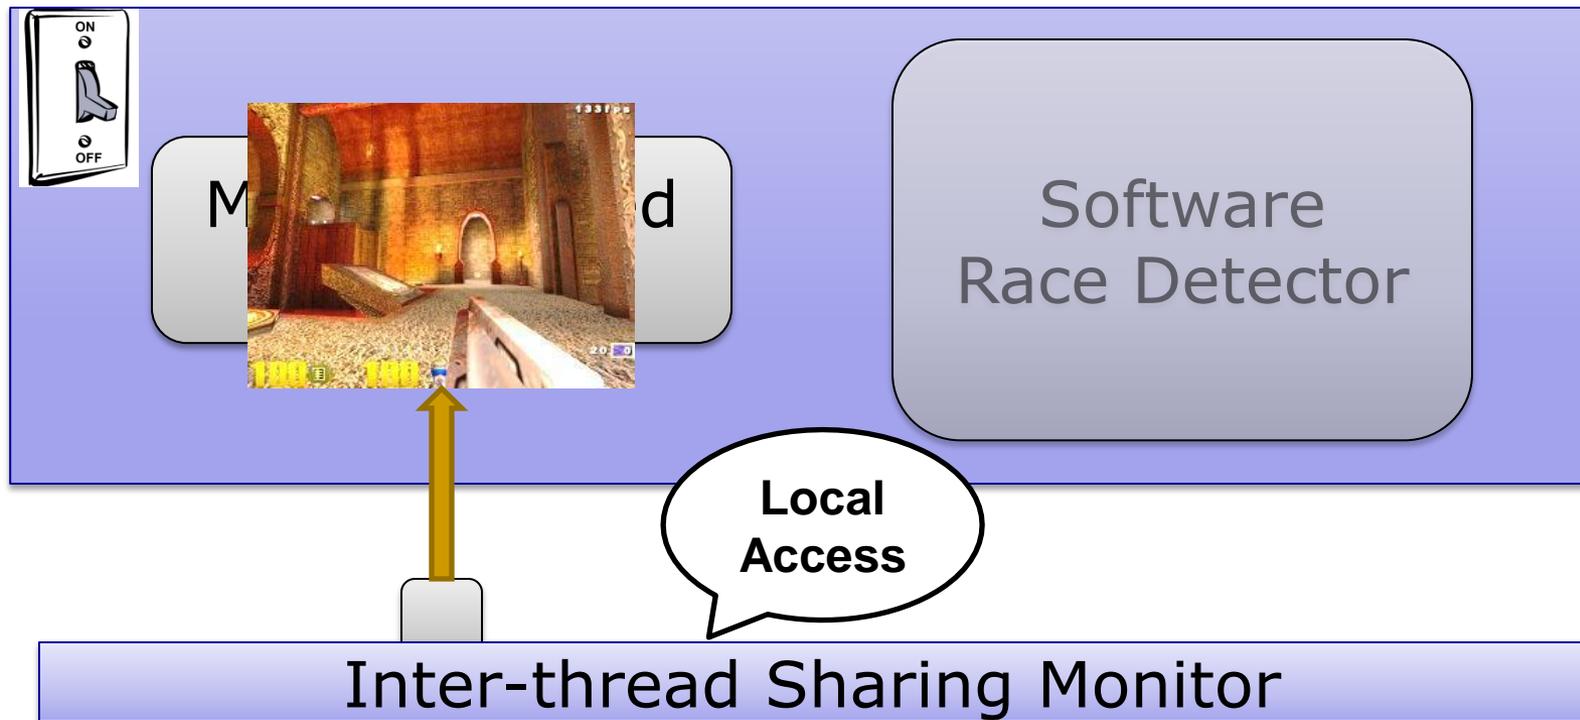


Inter-thread Sharing Monitor

# Use Demand-Driven Analysis!

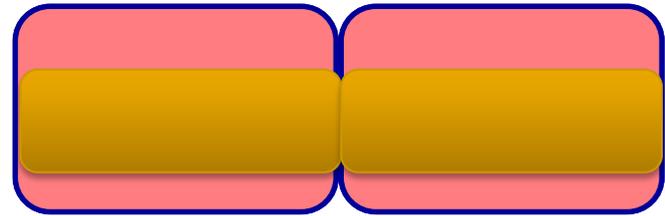
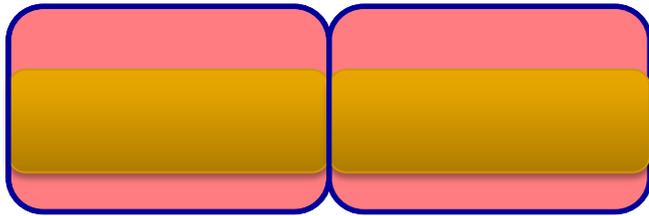


# Use Demand-Driven Analysis!



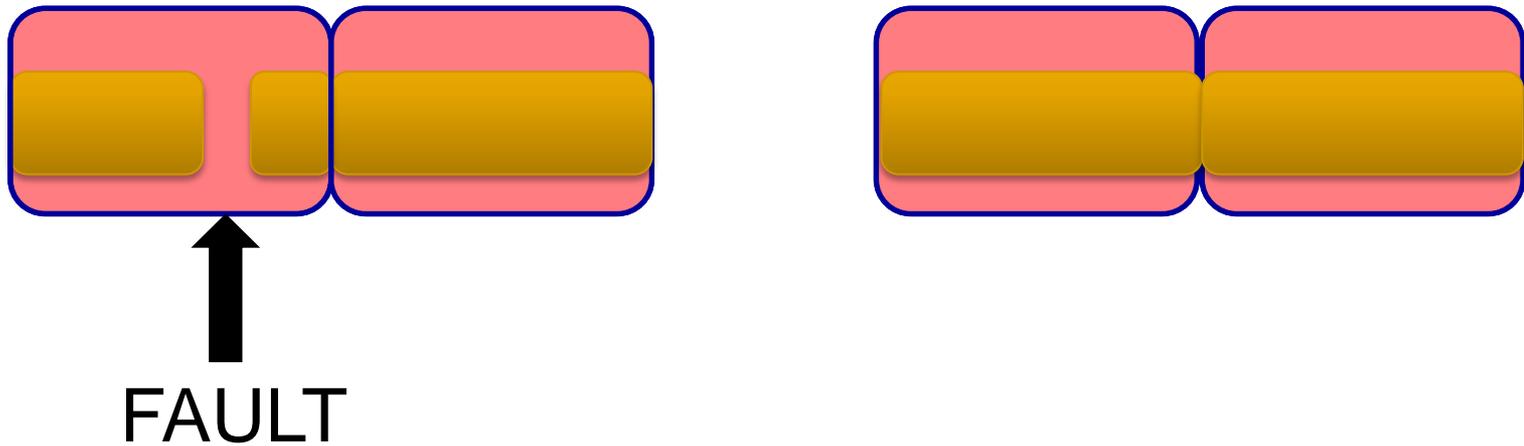
# Finding Inter-thread Sharing

- Virtual Memory Watchpoints?



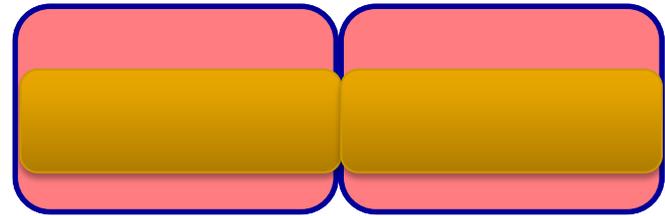
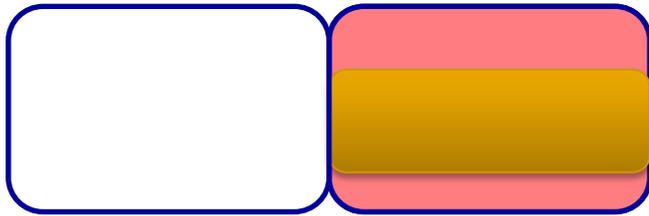
# Finding Inter-thread Sharing

- Virtual Memory Watchpoints?



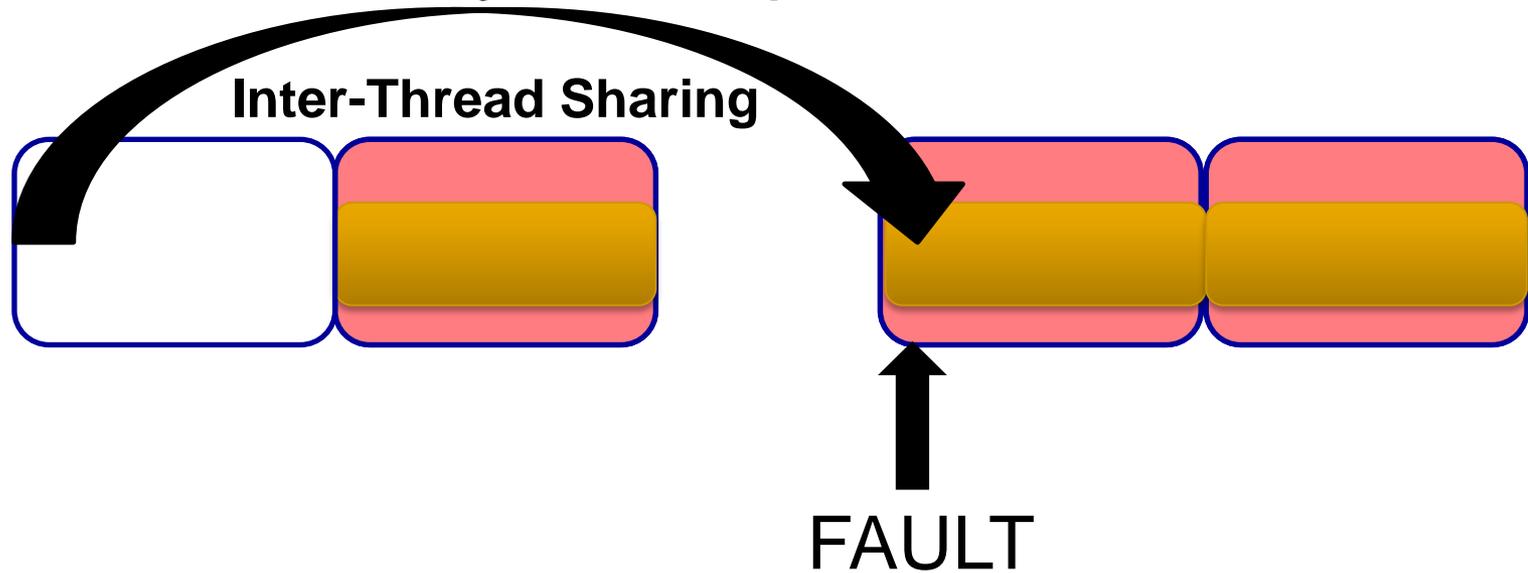
# Finding Inter-thread Sharing

- Virtual Memory Watchpoints?



# Finding Inter-thread Sharing

- Virtual Memory Watchpoints?



# Finding Inter-thread Sharing

- Virtual Memory Watchpoints?



- ~100% of accesses cause page faults

# Finding Inter-thread Sharing

- Virtual Memory Watchpoints?



– ~100% of accesses cause page faults

- Granularity Gap

# Finding Inter-thread Sharing

- Virtual Memory Watchpoints?

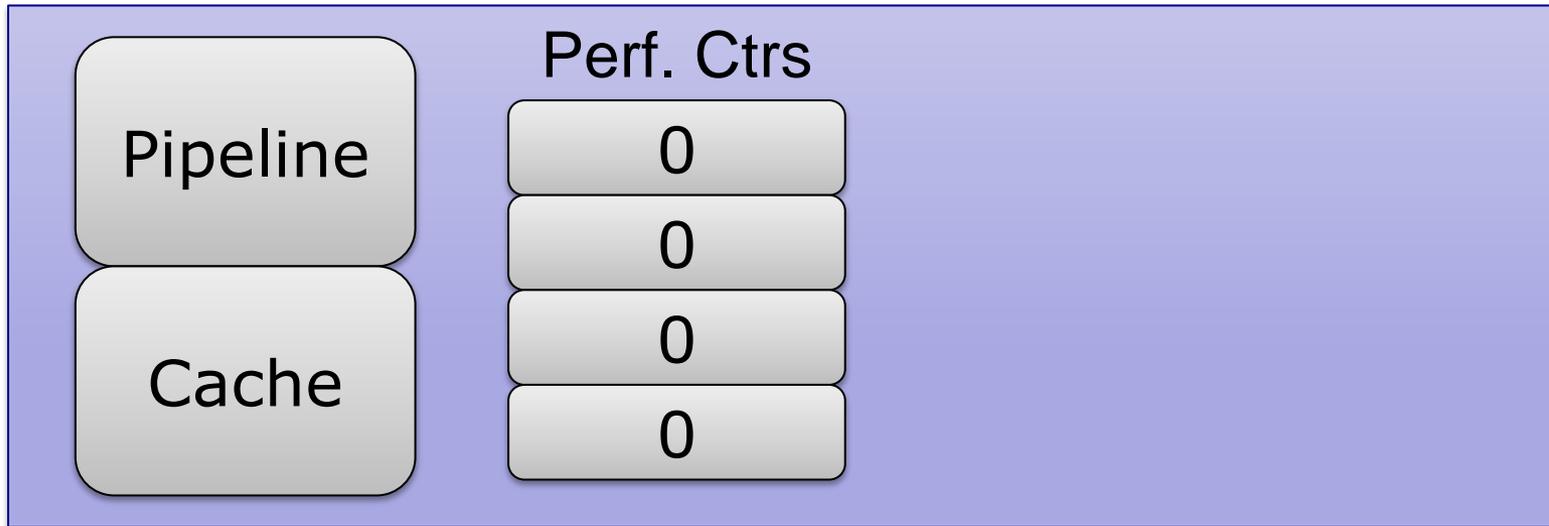


- ~100% of accesses cause page faults

- Granularity Gap
- Per-process not per-thread
- Must go through the kernel on faults
- Syscalls for setting/removing meta-data

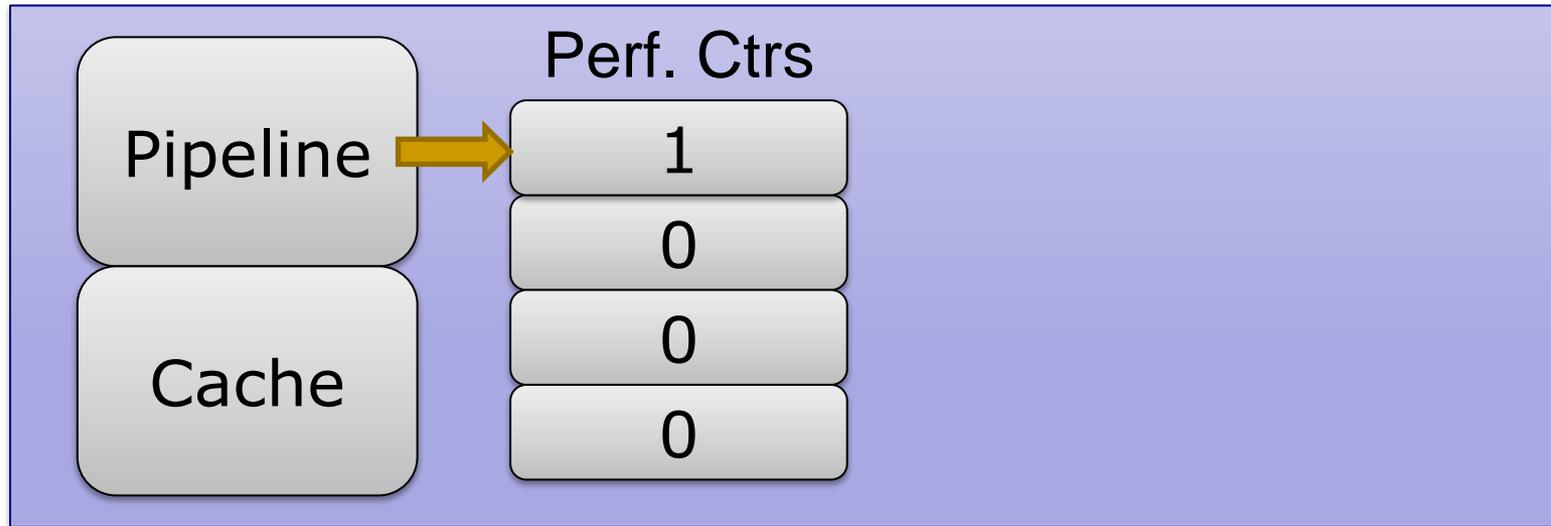
# Hardware Sharing Detector

- Hardware Performance Counters



# Hardware Sharing Detector

- Hardware Performance Counters



# Hardware Sharing Detector

- Hardware Performance Counters



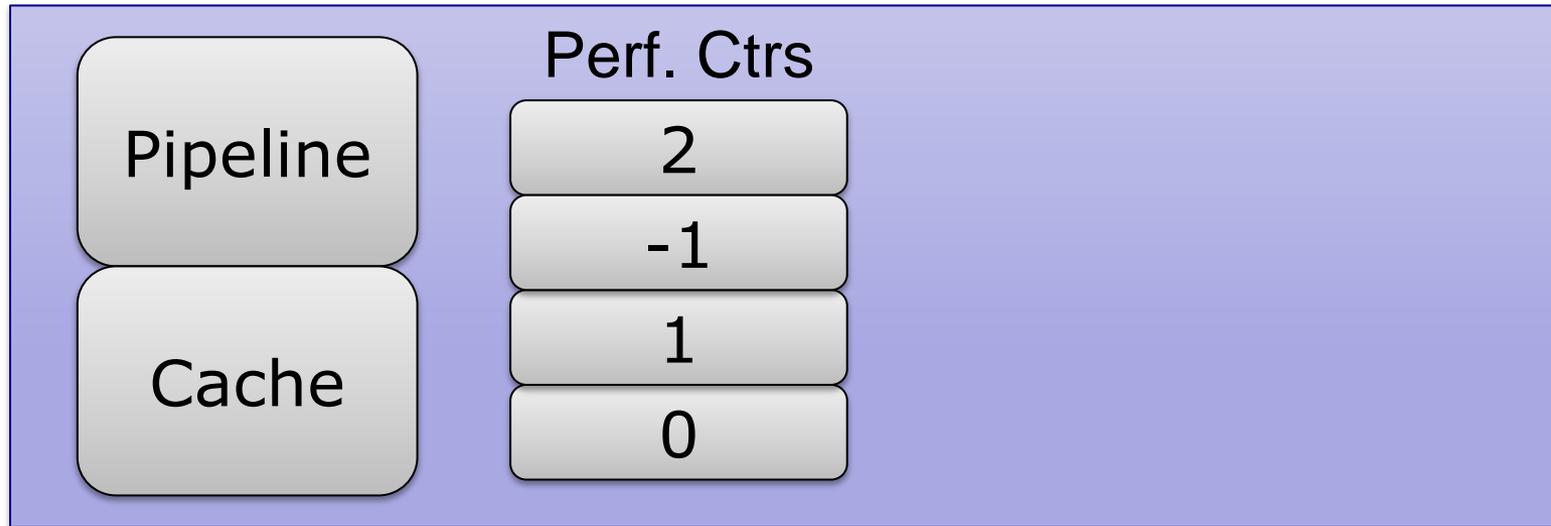
# Hardware Sharing Detector

- Hardware Performance Counters



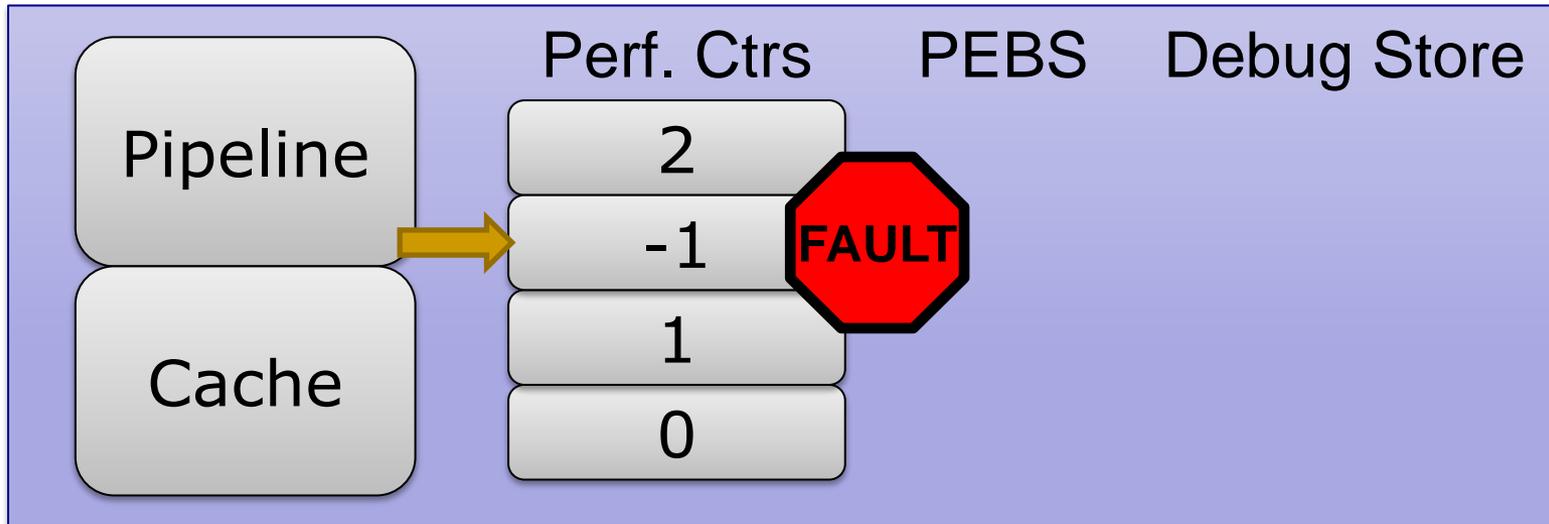
# Hardware Sharing Detector

- Hardware Performance Counters



# Hardware Sharing Detector

- Hardware Performance Counters



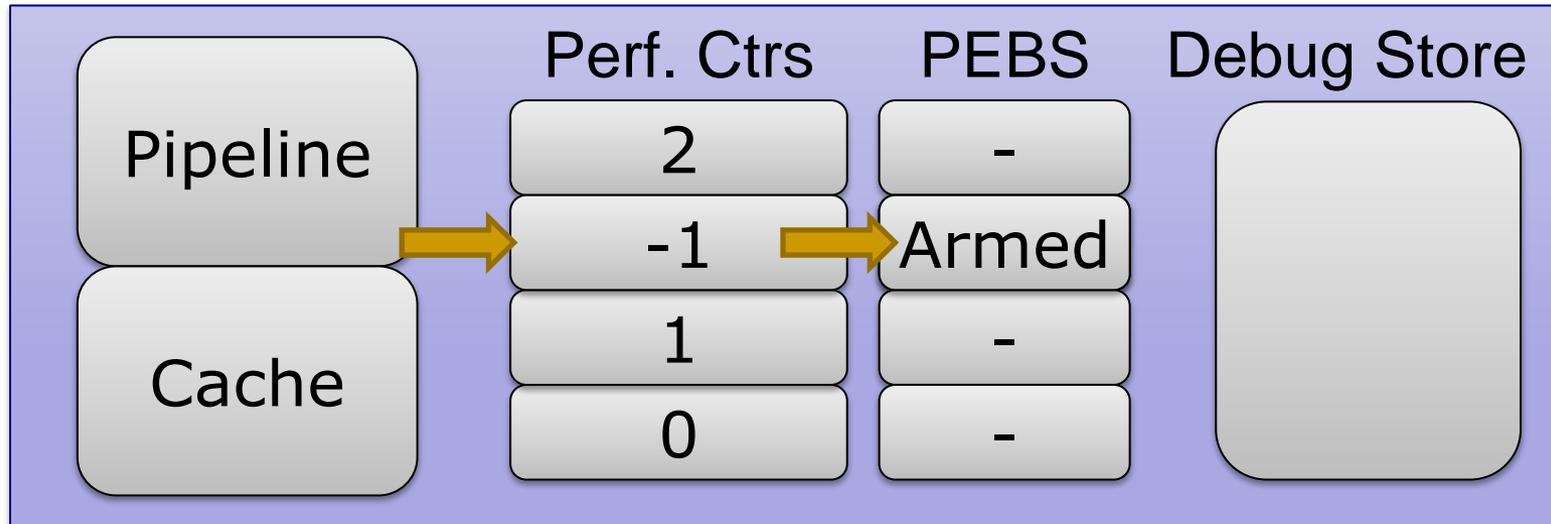
# Hardware Sharing Detector

- Hardware Performance Counters

	Perf. Ctrs	PEBS	Debug Store
Pipeline	2	-	
	-1	-	
Cache	1	-	
	0	-	

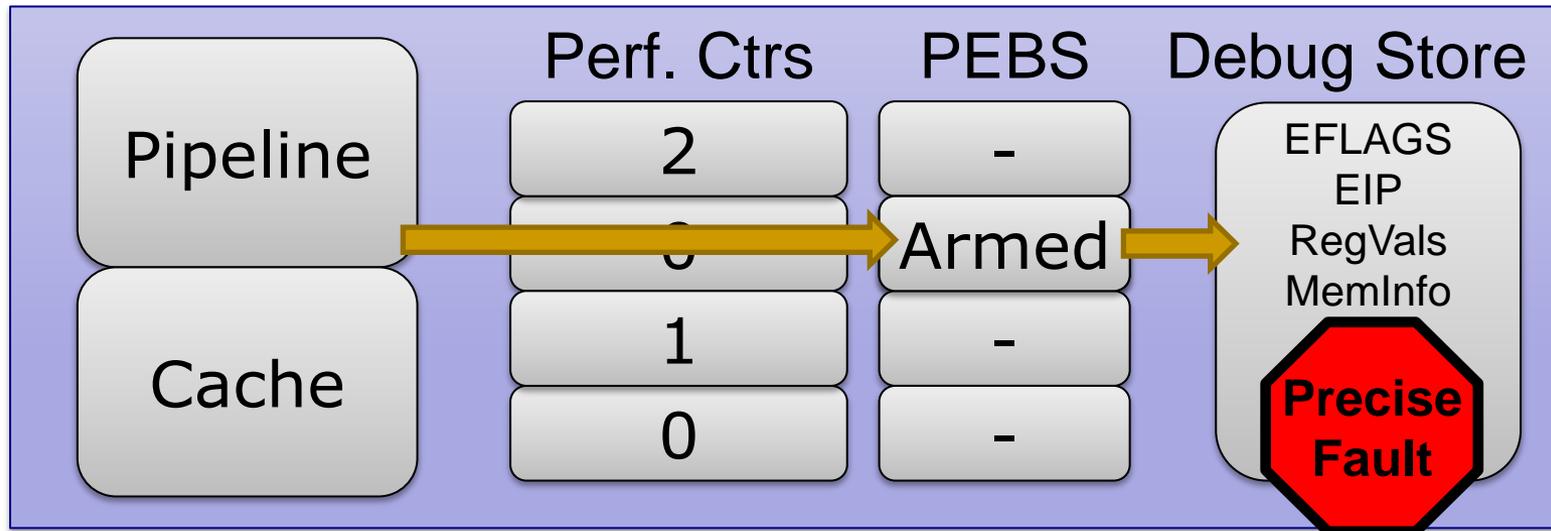
# Hardware Sharing Detector

- Hardware Performance Counters



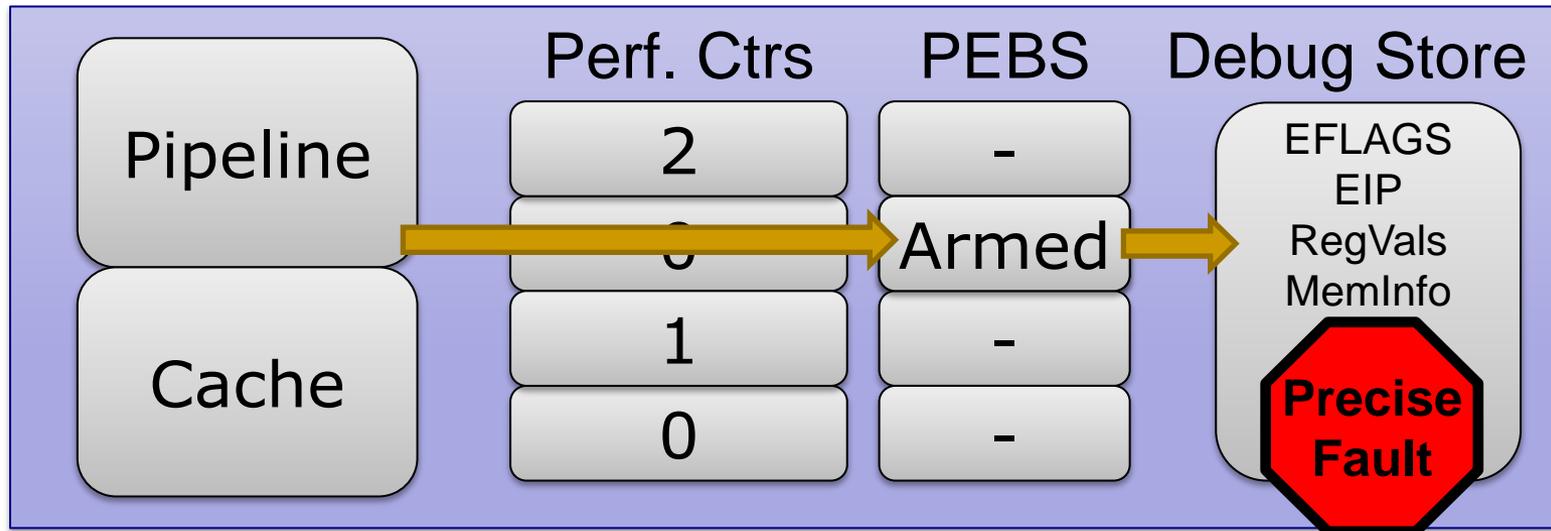
# Hardware Sharing Detector

- Hardware Performance Counters

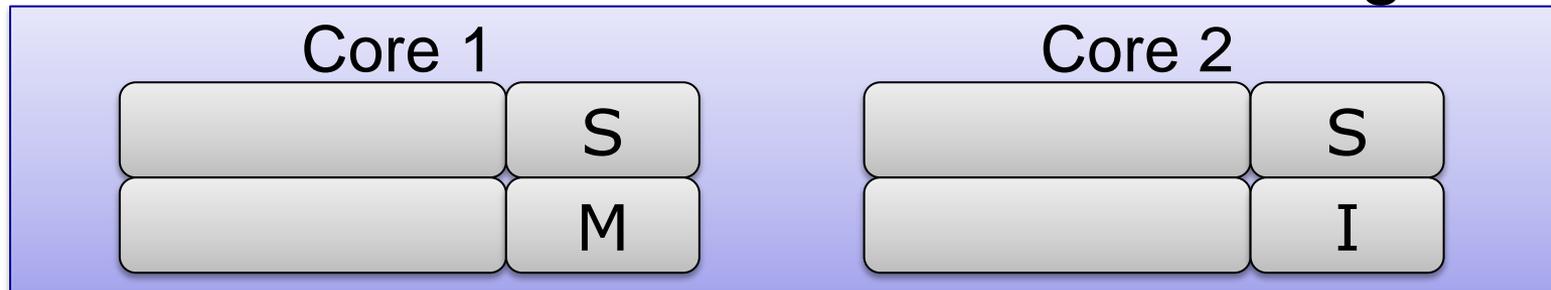


# Hardware Sharing Detector

- Hardware Performance Counters

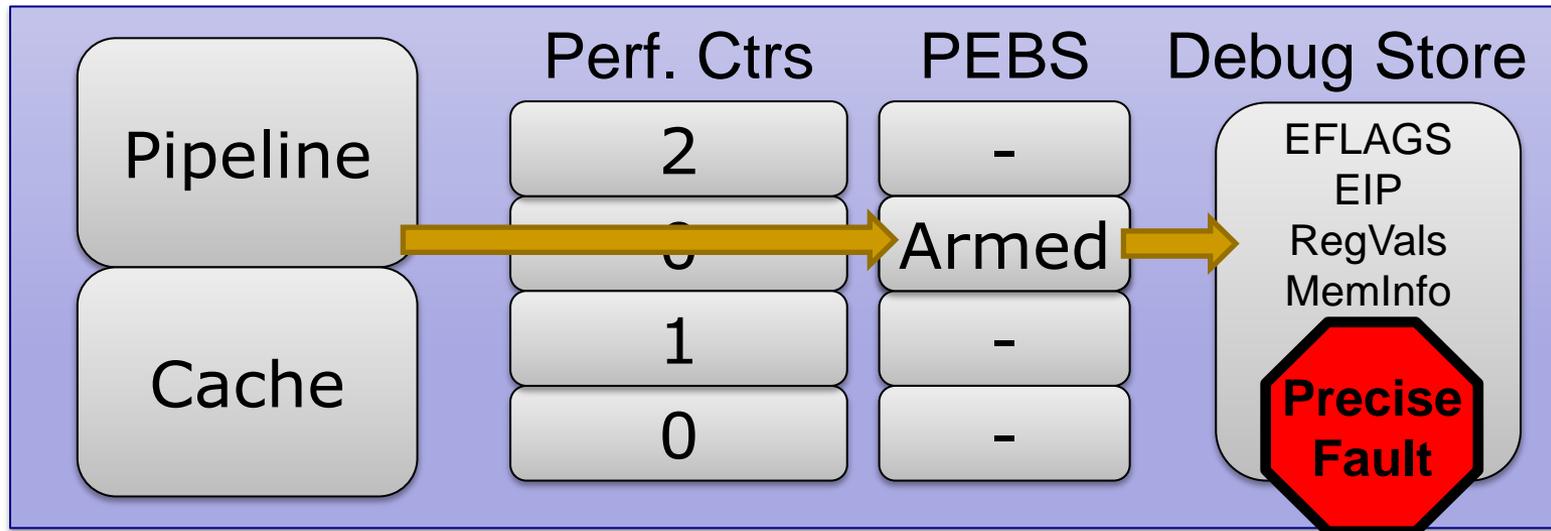


- Intel's HITM event: W→R Data Sharing

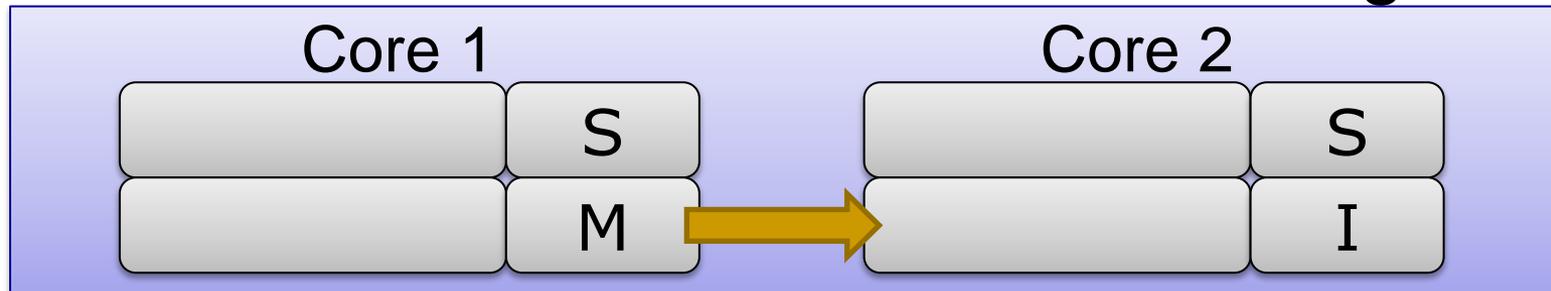


# Hardware Sharing Detector

- Hardware Performance Counters

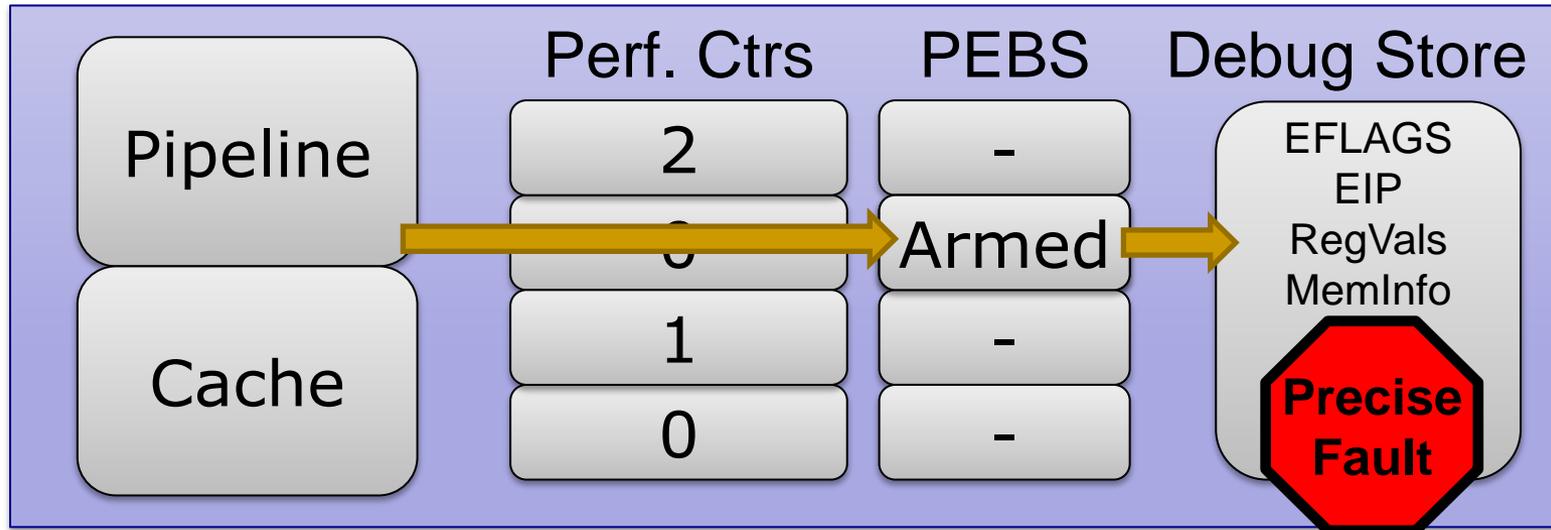


- Intel's HITM event: W→R Data Sharing

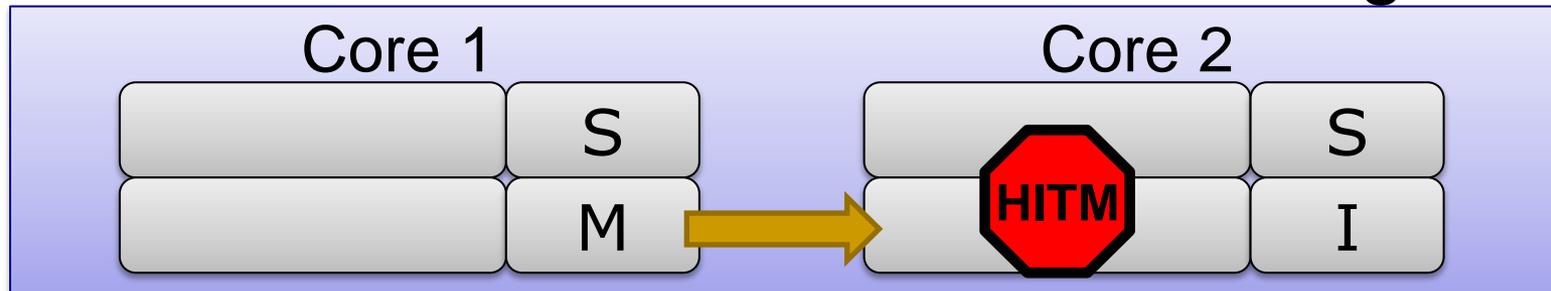


# Hardware Sharing Detector

- Hardware Performance Counters



- Intel's HITM event: W→R Data Sharing



---

# Potential Accuracy & Perf. Problems

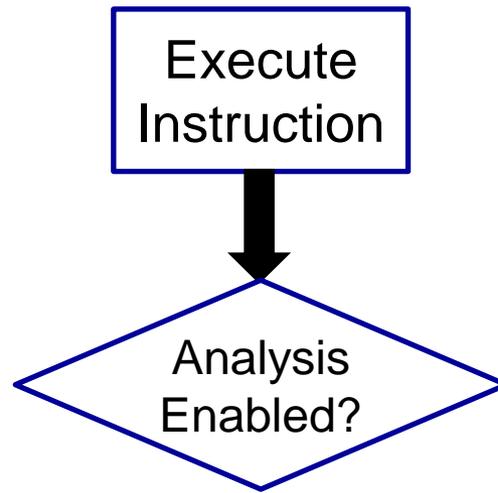
- Limitations of Performance Counters
  - HITM only finds  $W \rightarrow R$  Data Sharing
  - Hardware prefetcher events aren't counted
- Limitations of Cache Events
  - SMT sharing can't be counted
  - Cache eviction causes missed events
  - False sharing, etc...
- PEBS events still go through the kernel

---

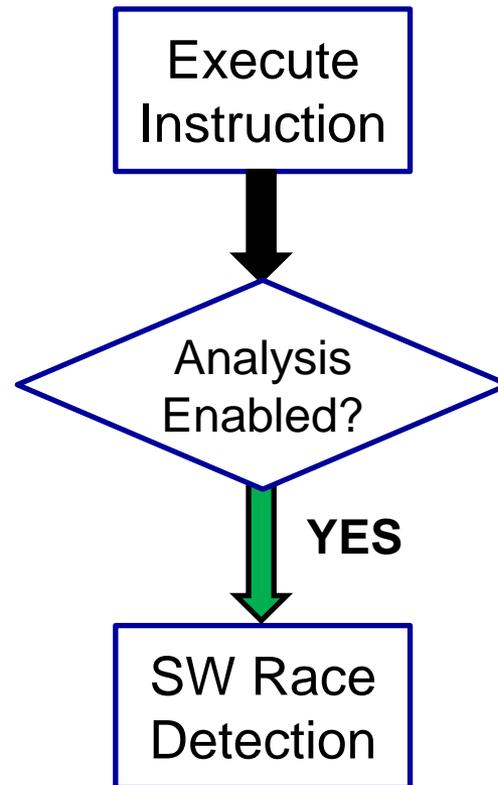
# Demand-Driven Analysis on Real HW

Execute  
Instruction

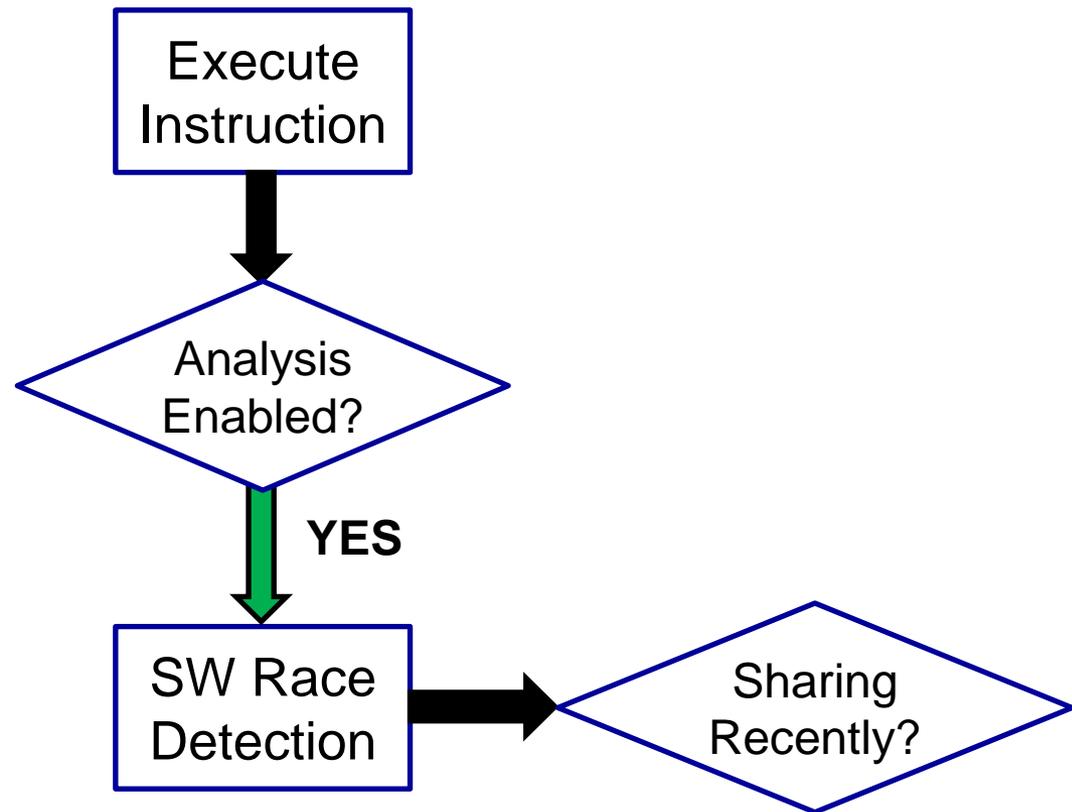
# Demand-Driven Analysis on Real HW



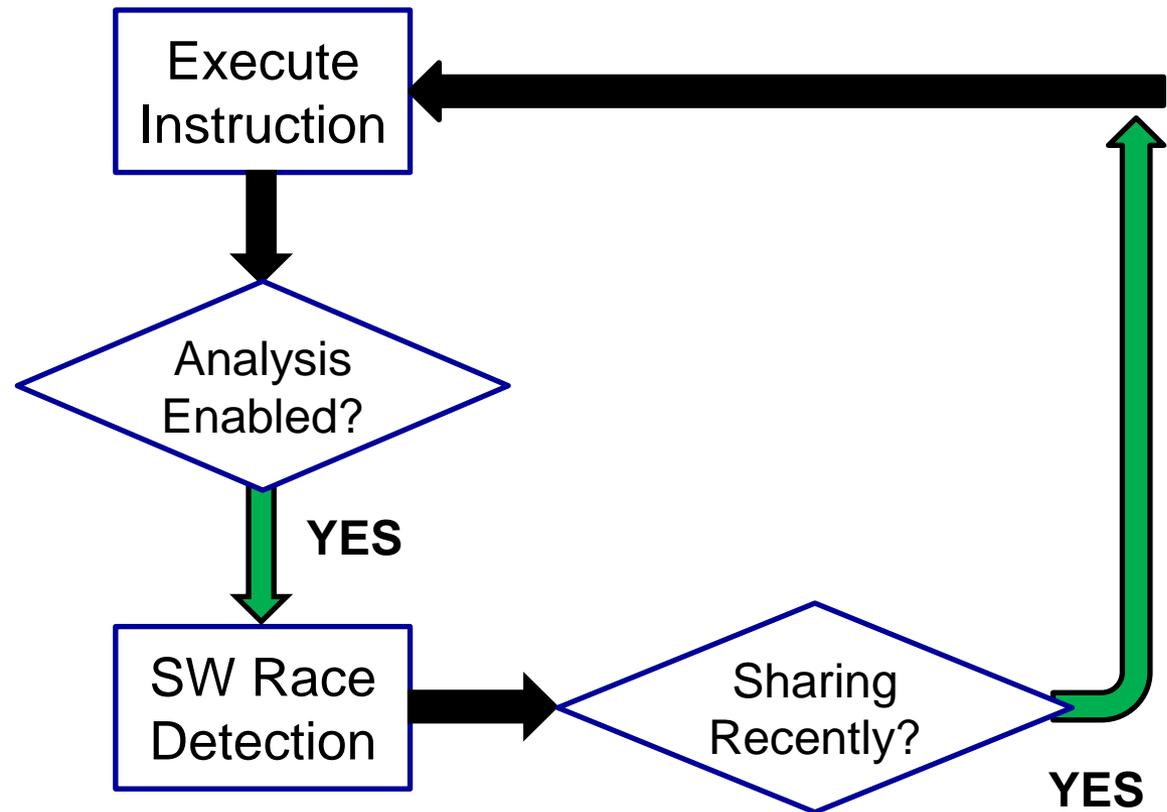
# Demand-Driven Analysis on Real HW



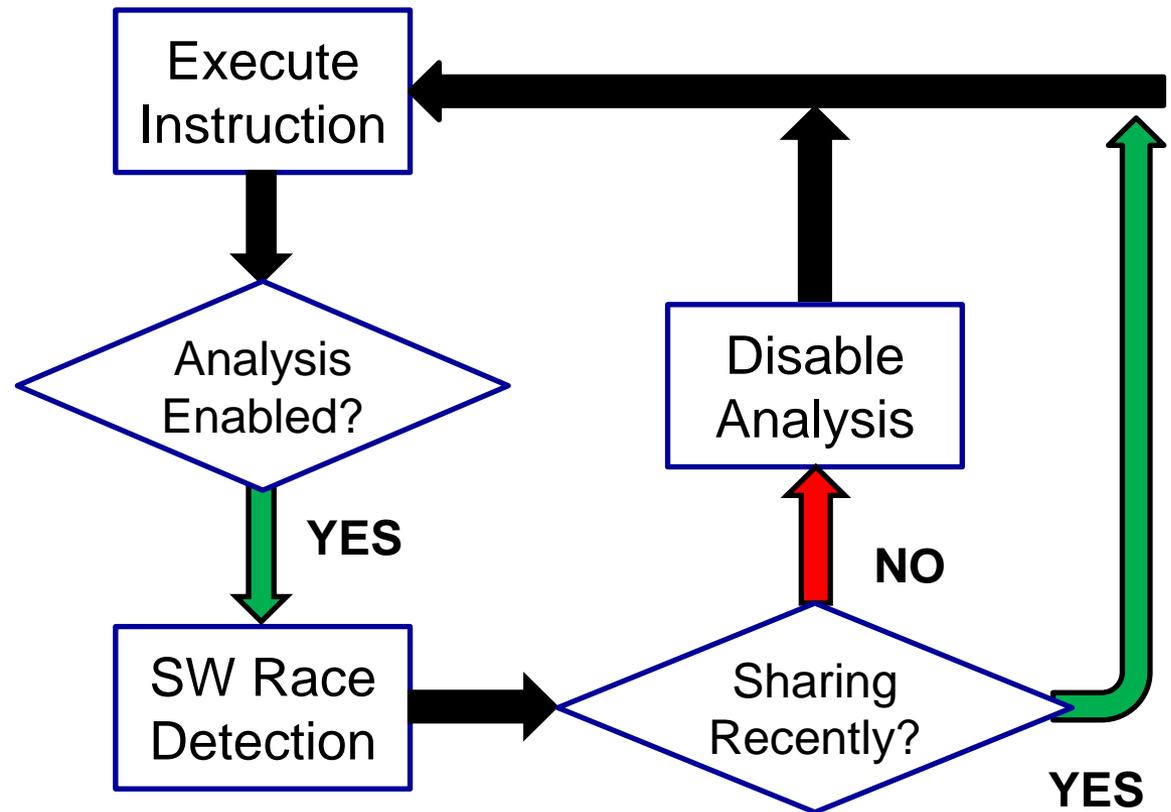
# Demand-Driven Analysis on Real HW



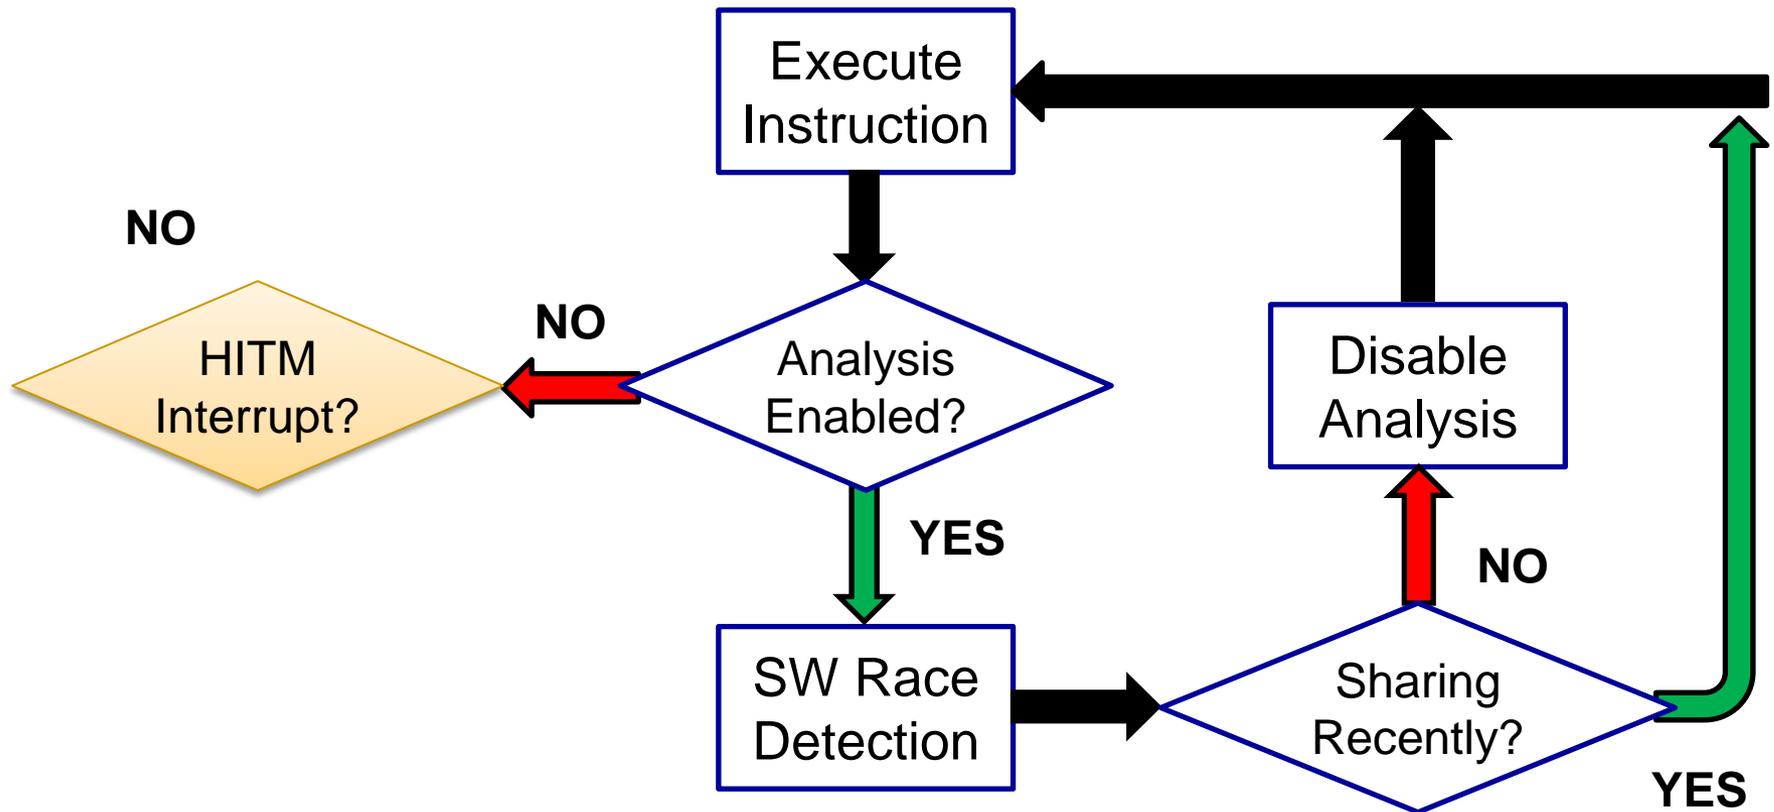
# Demand-Driven Analysis on Real HW



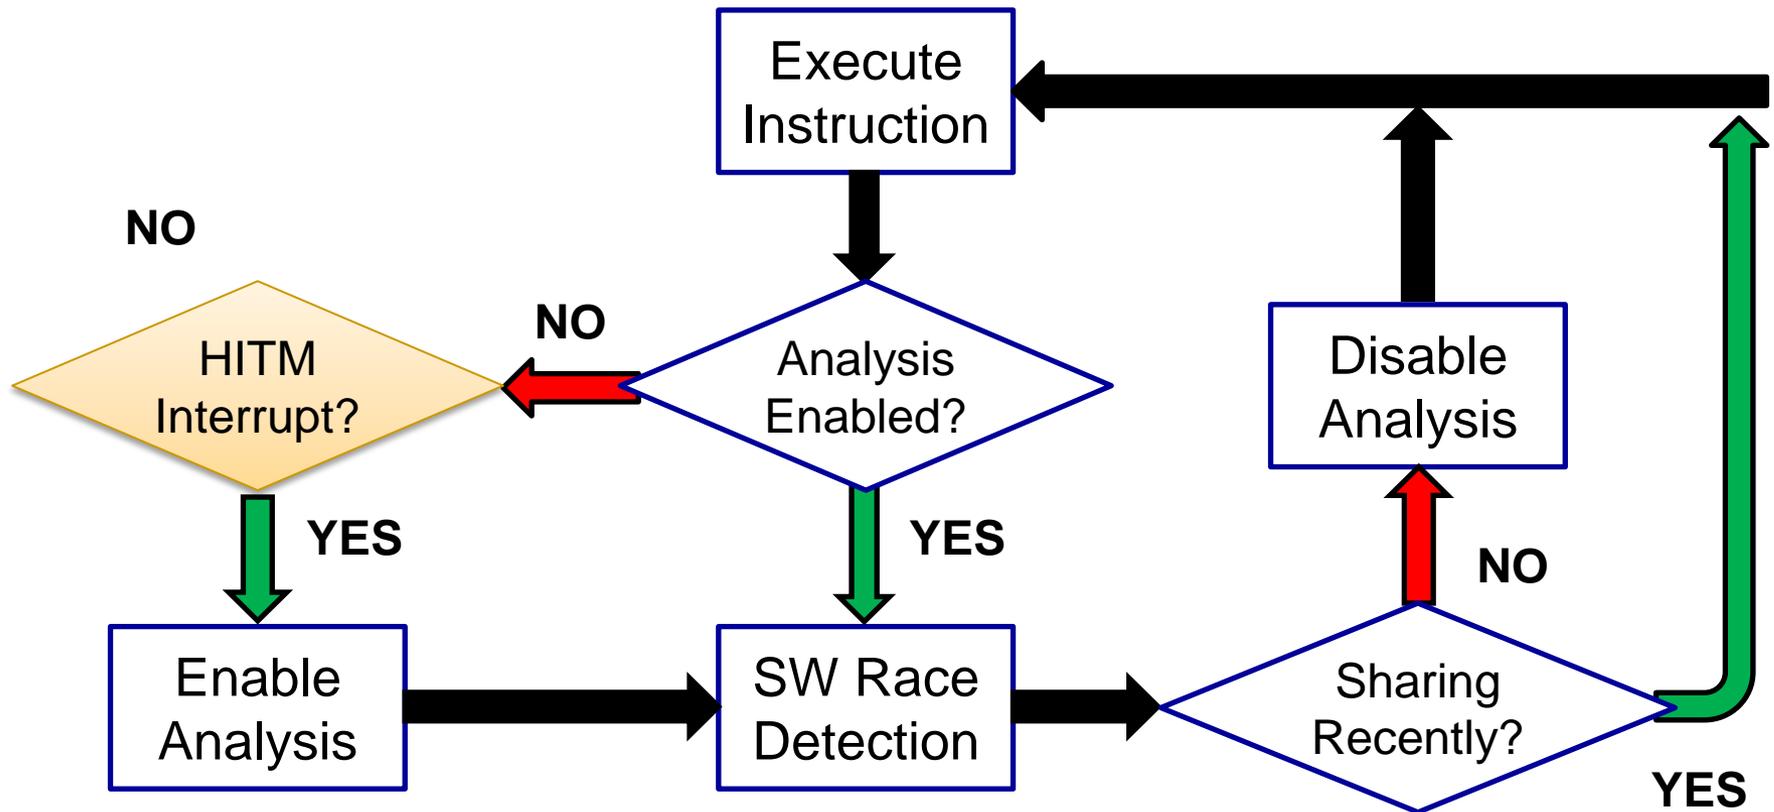
# Demand-Driven Analysis on Real HW



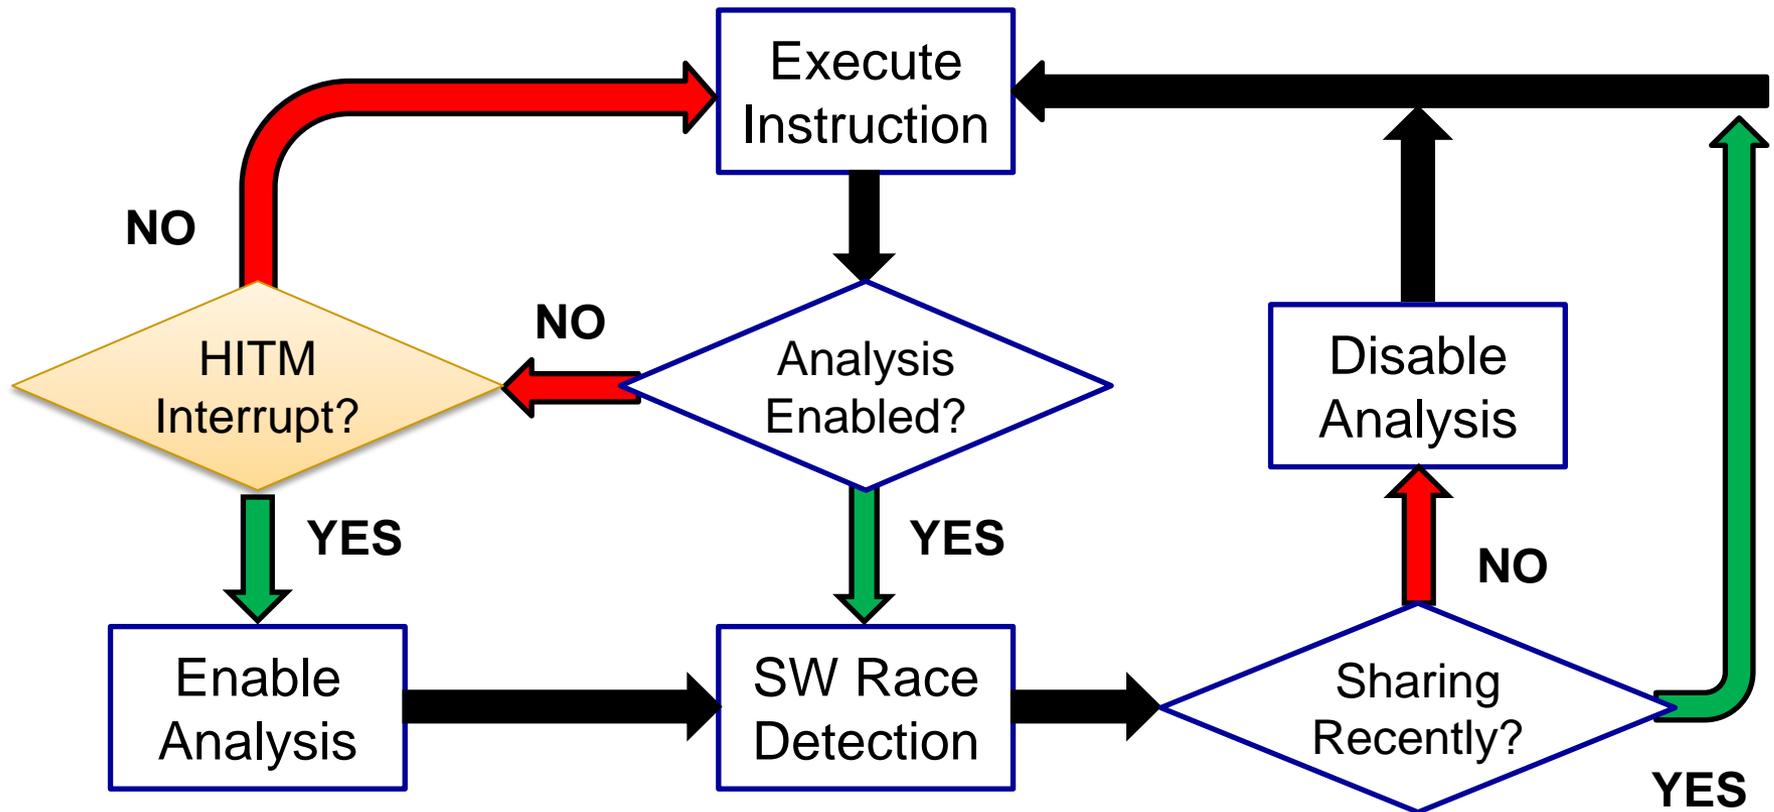
# Demand-Driven Analysis on Real HW



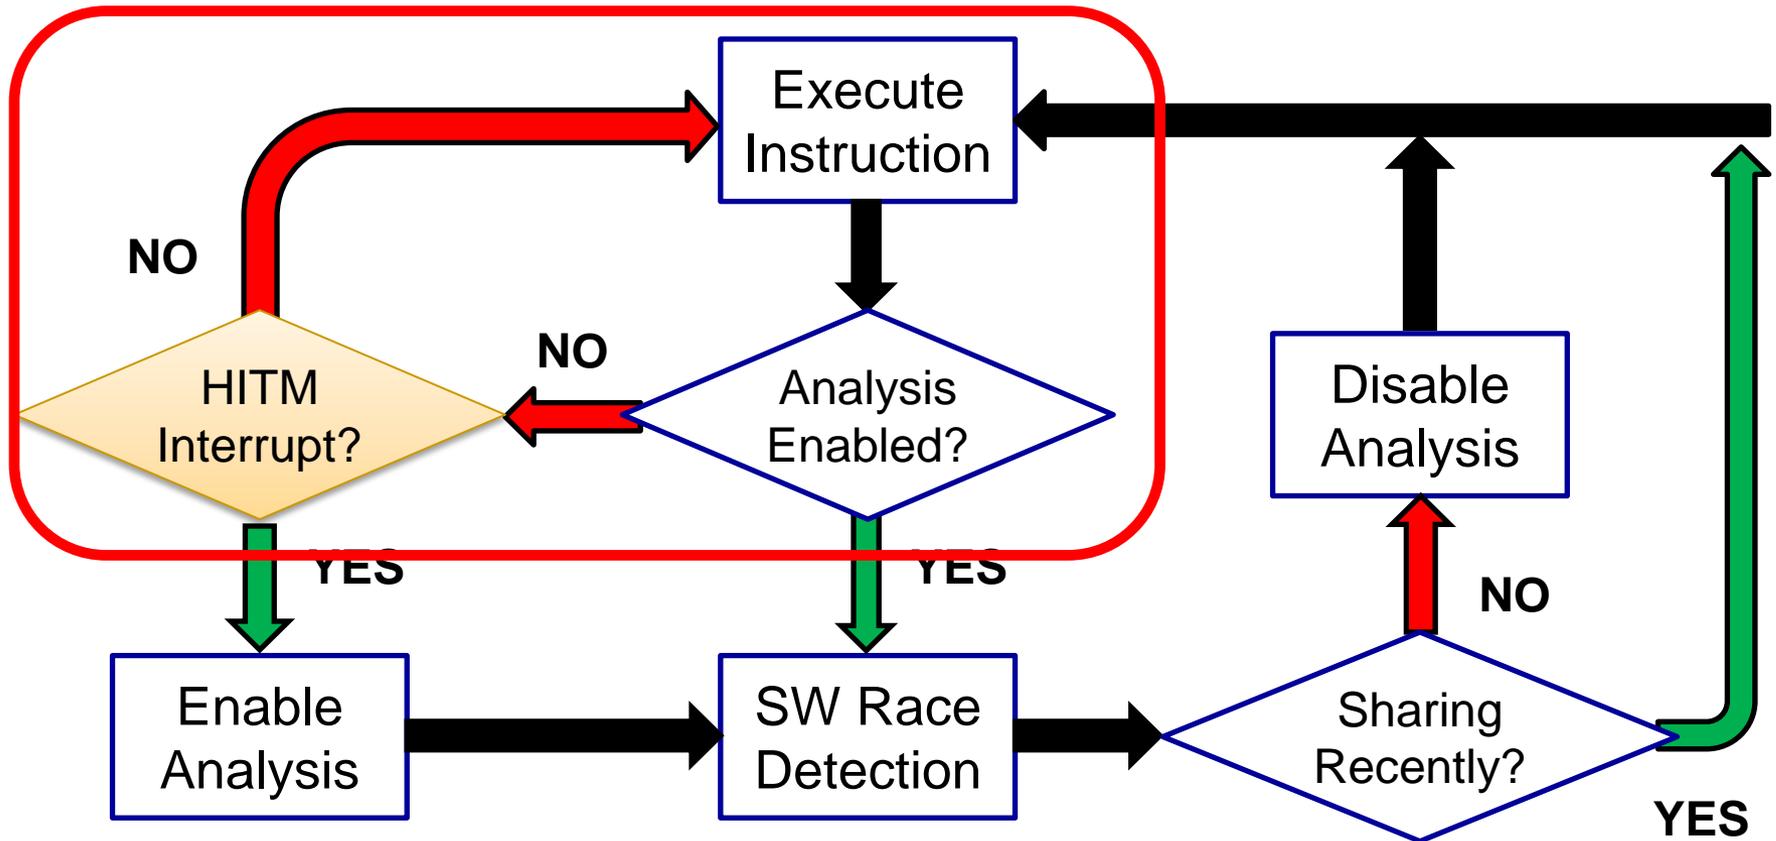
# Demand-Driven Analysis on Real HW



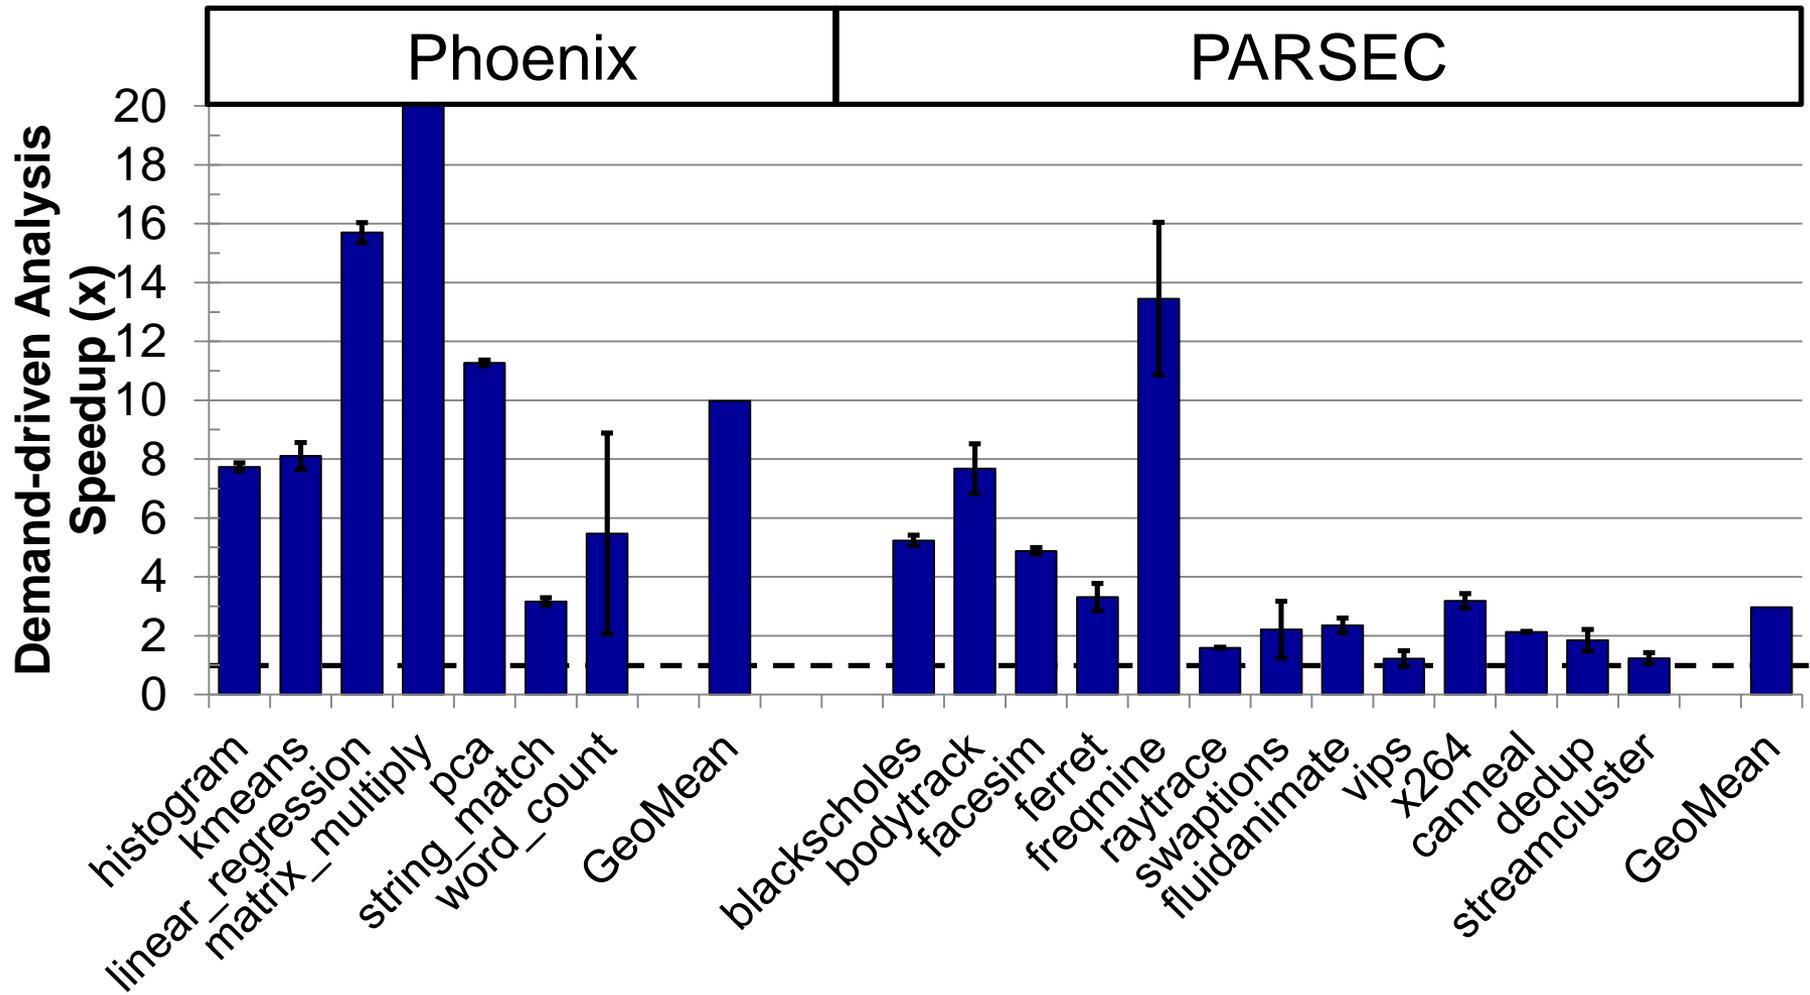
# Demand-Driven Analysis on Real HW



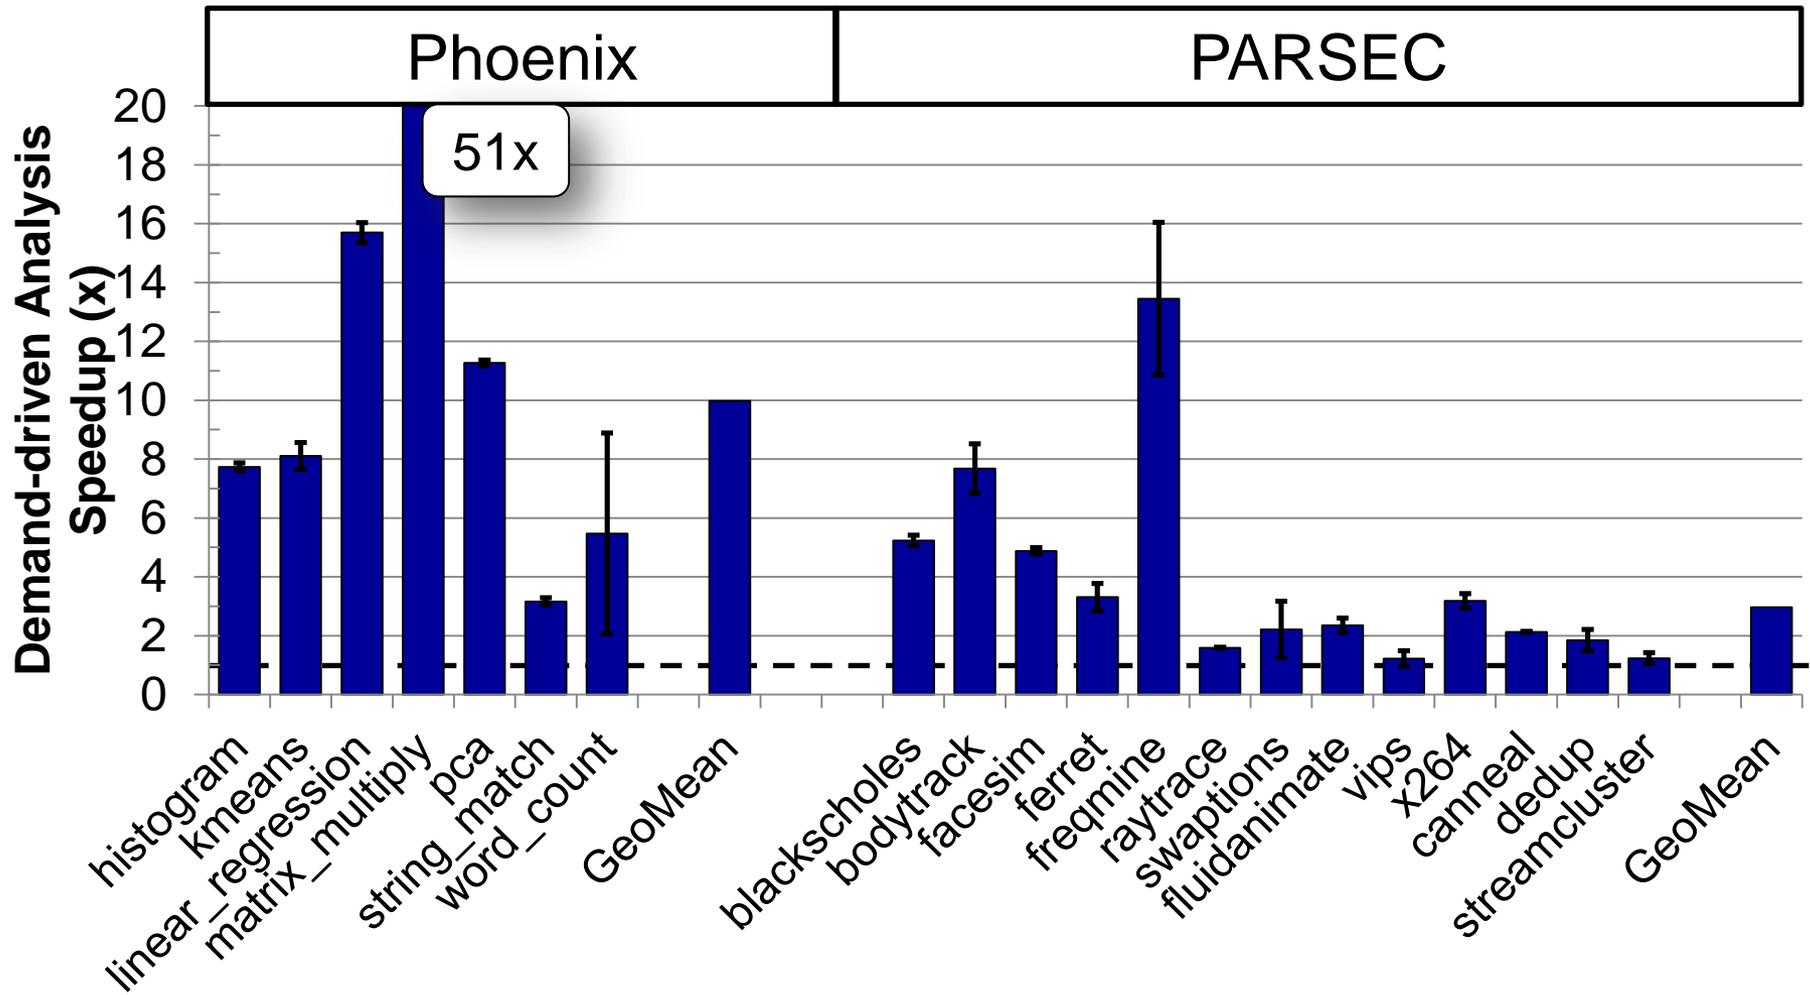
# Demand-Driven Analysis on Real HW



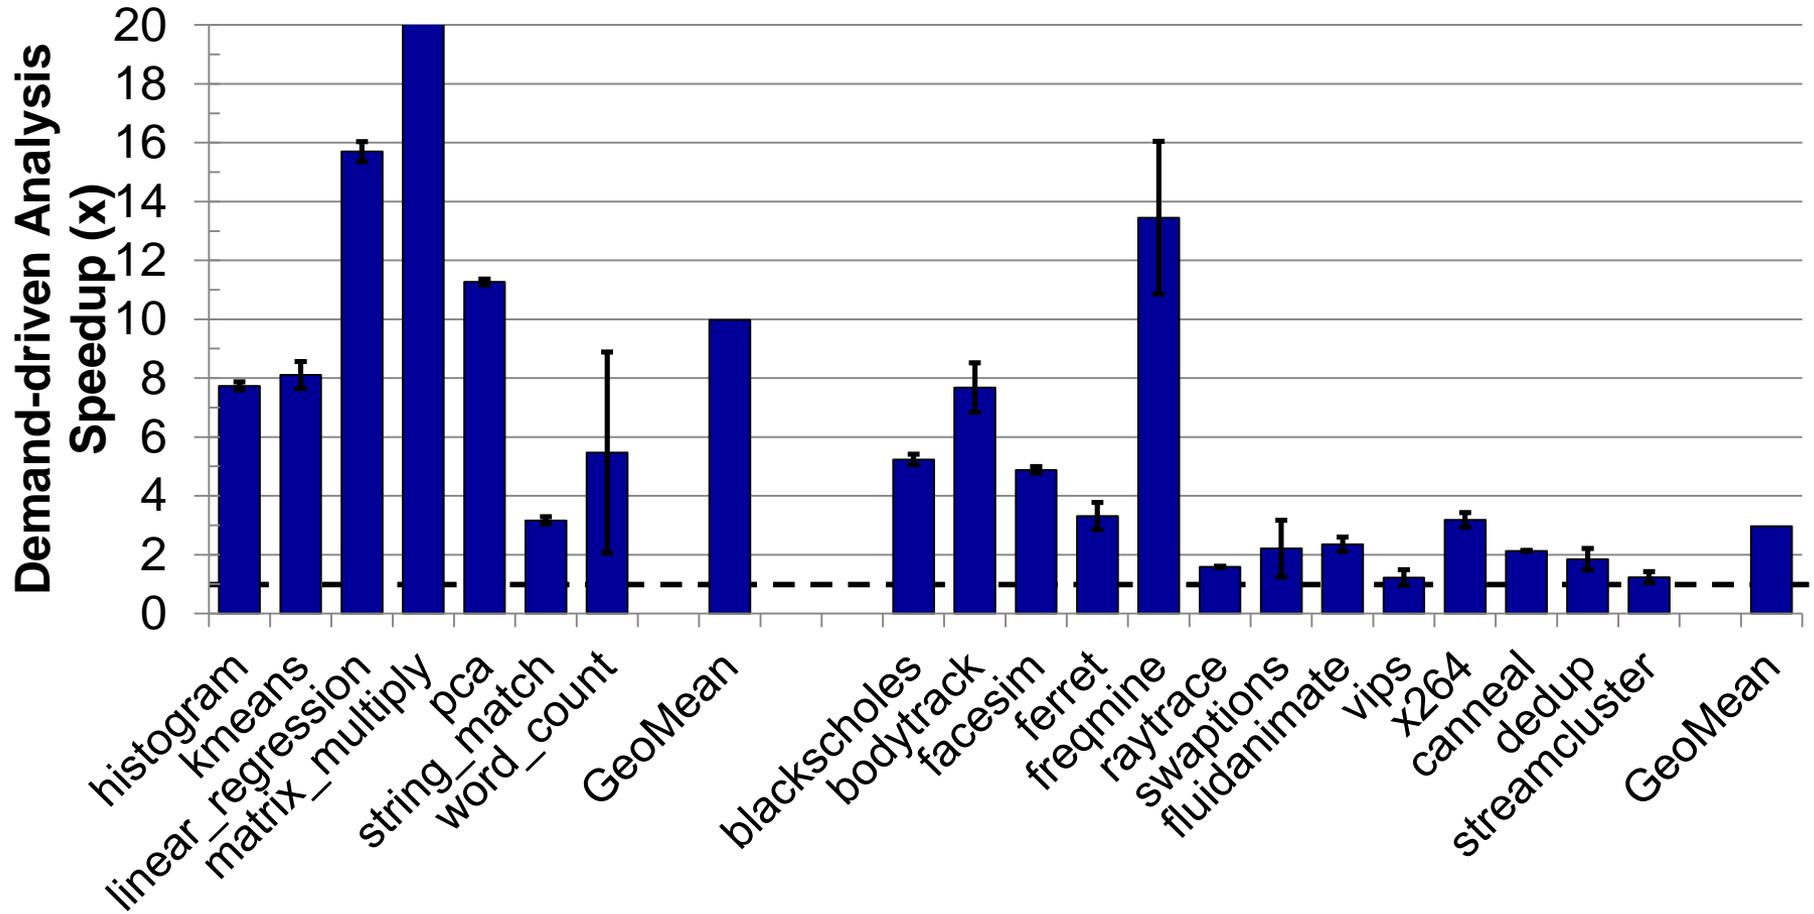
# Performance Increases



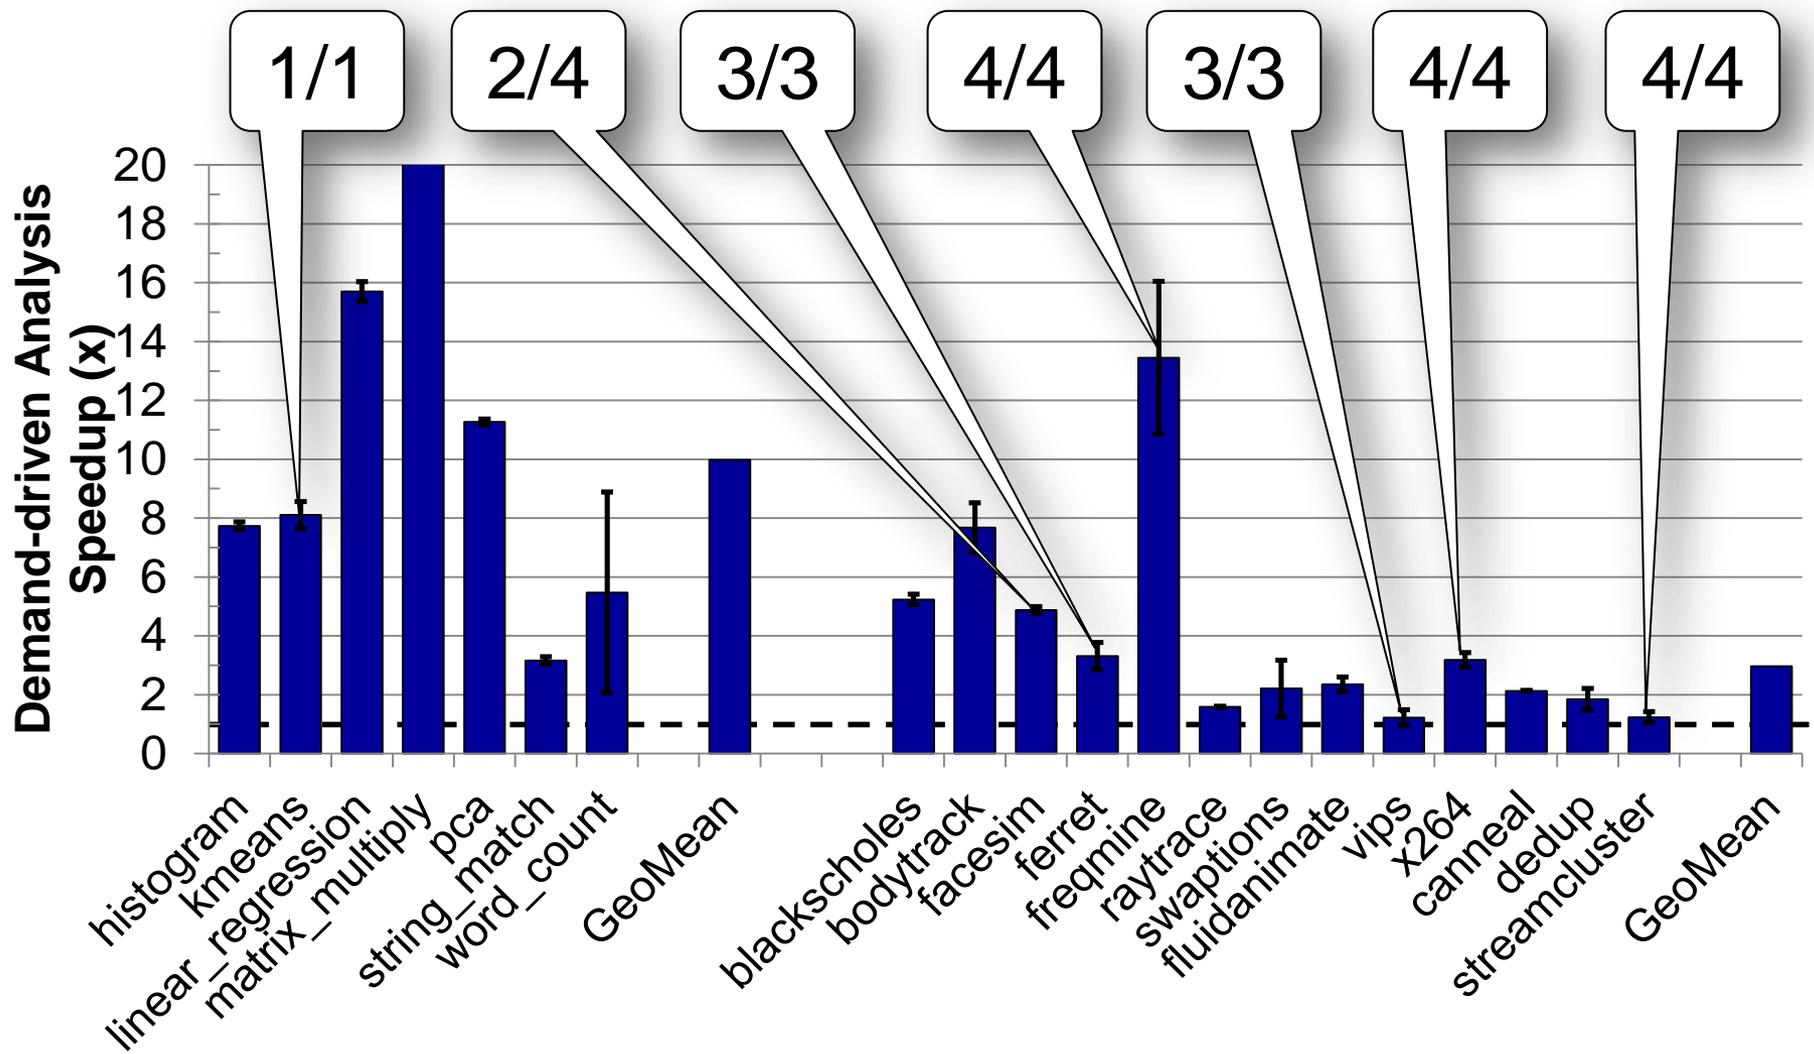
# Performance Increases



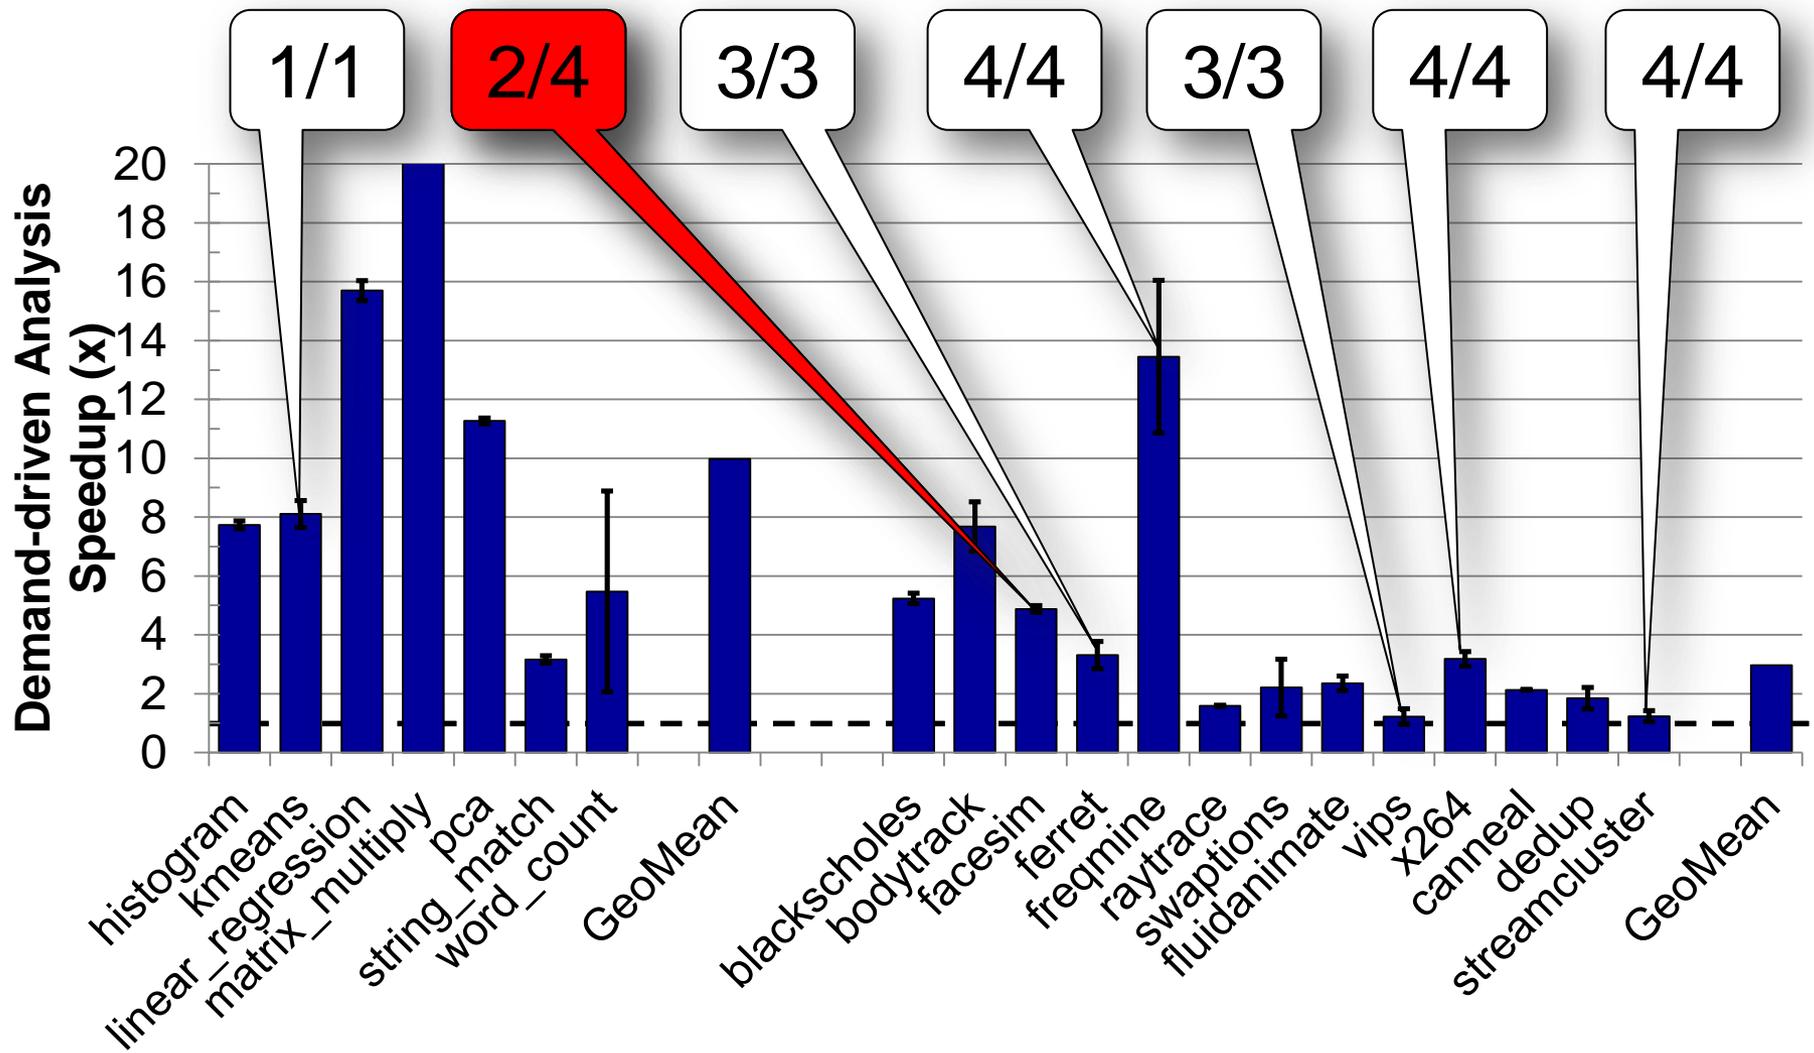
# Demand-Driven Analysis Accuracy



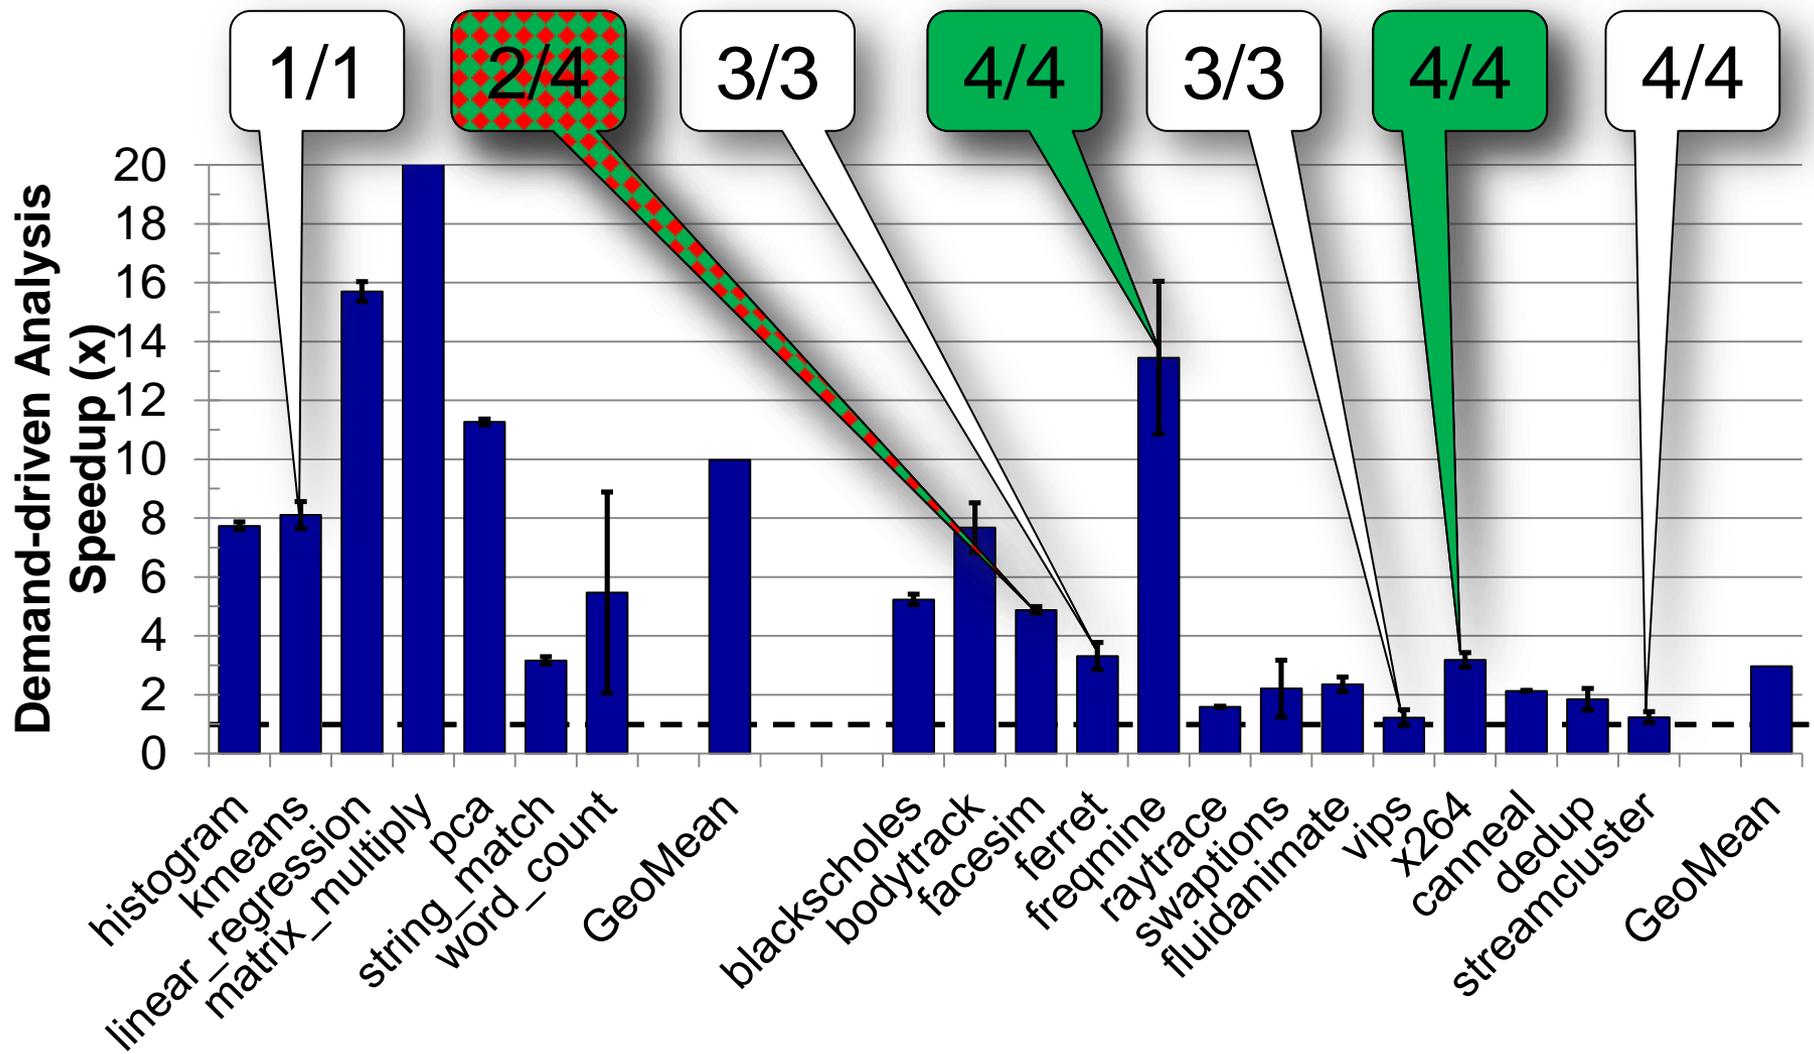
# Demand-Driven Analysis Accuracy



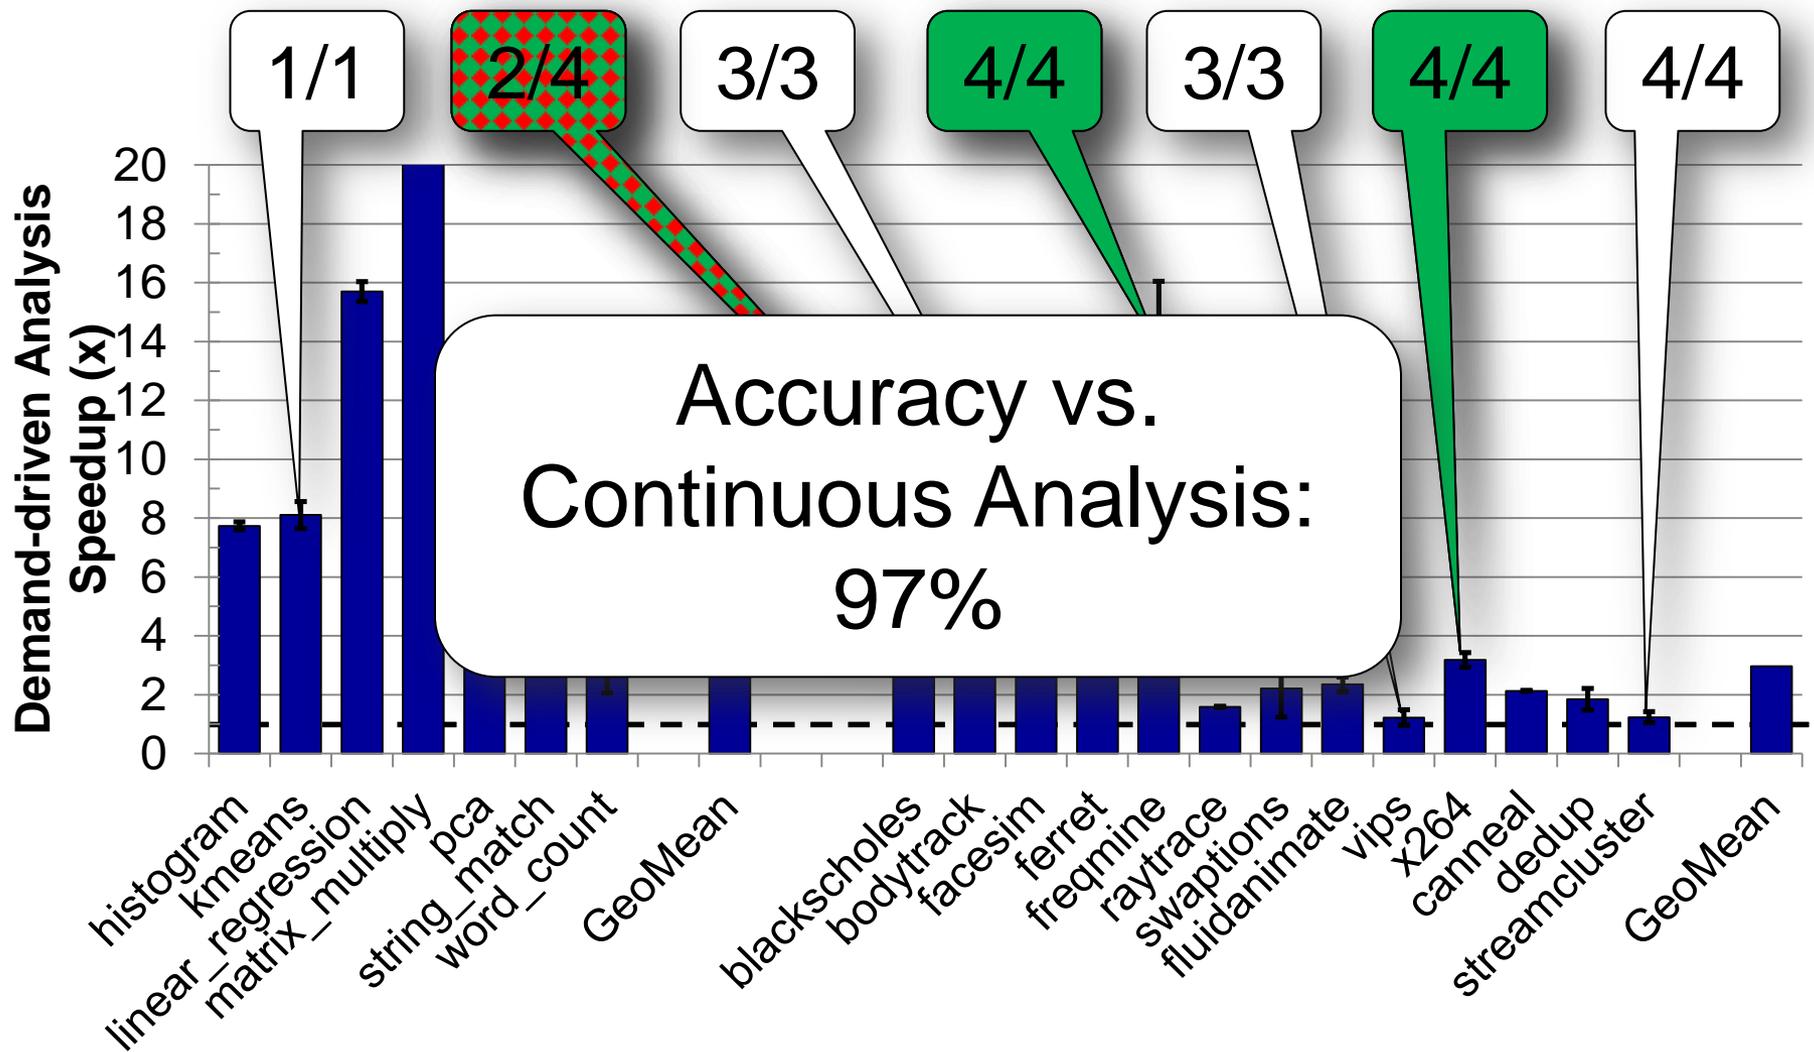
# Demand-Driven Analysis Accuracy



# Demand-Driven Analysis Accuracy



# Demand-Driven Analysis Accuracy



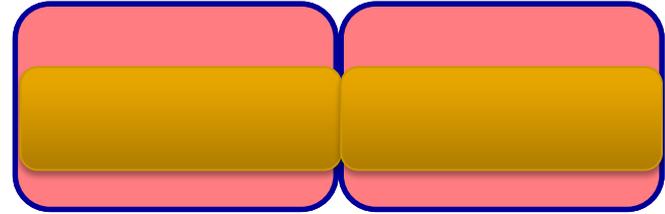
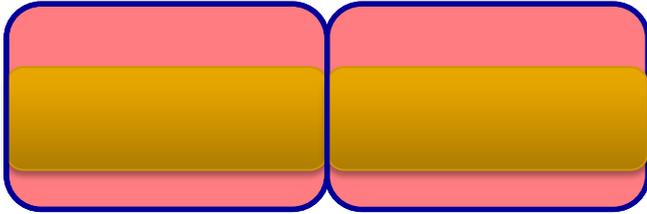
---

# Outline

- Problem Statement
- Background Information
  - Demand-Driven Dynamic Dataflow Analysis
- Proposed Solutions
  - Demand-Driven Data Race Detection
  - Unlimited Hardware Watchpoints

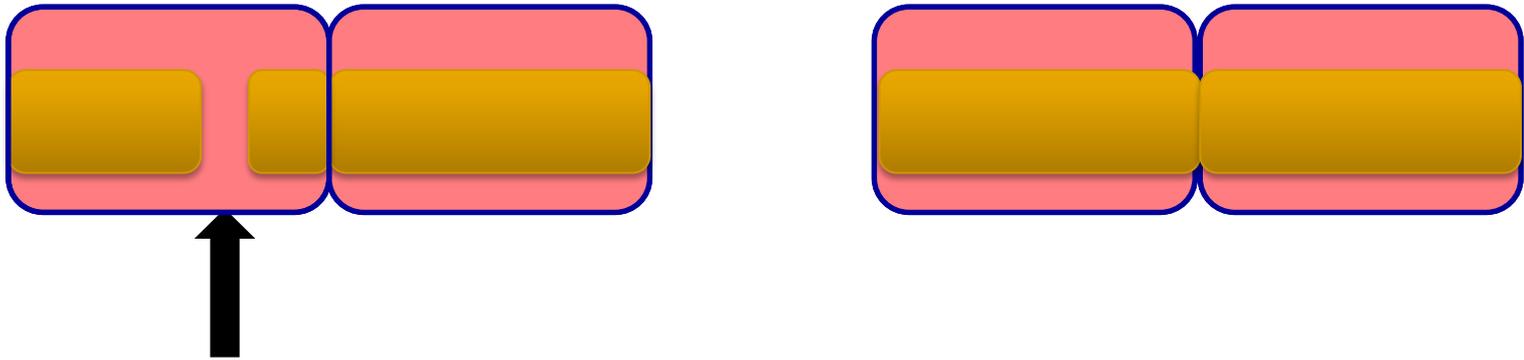
# Watchpoints Globally Useful

- Byte/Word Accurate and Per-Thread



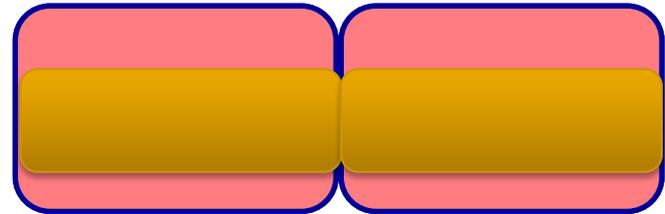
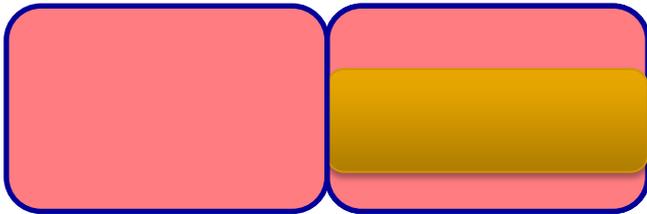
# Watchpoints Globally Useful

- Byte/Word Accurate and Per-Thread



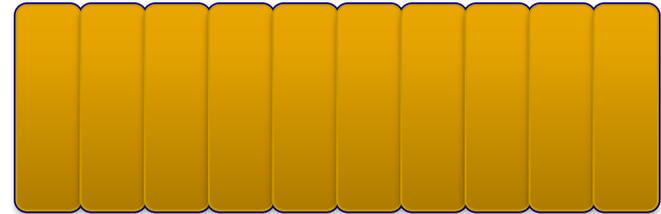
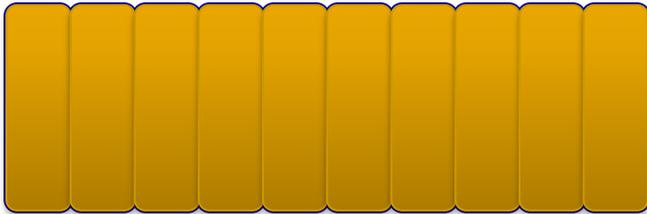
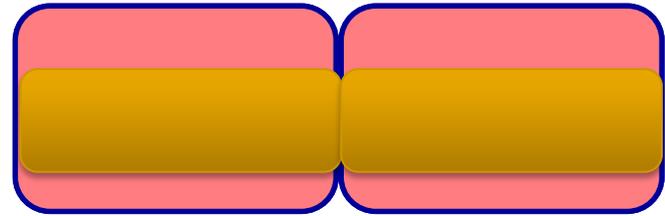
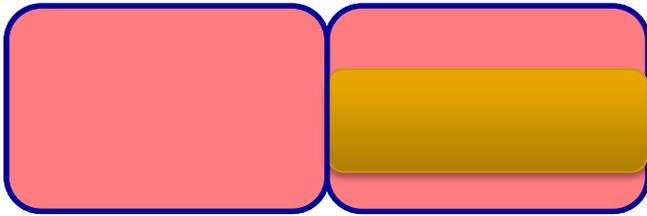
# Watchpoints Globally Useful

- Byte/Word Accurate and Per-Thread



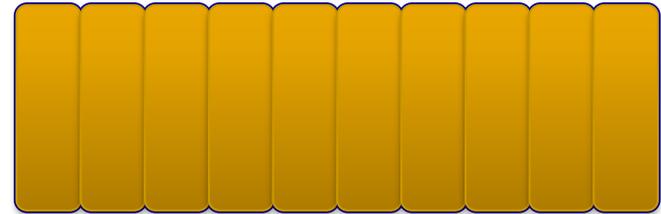
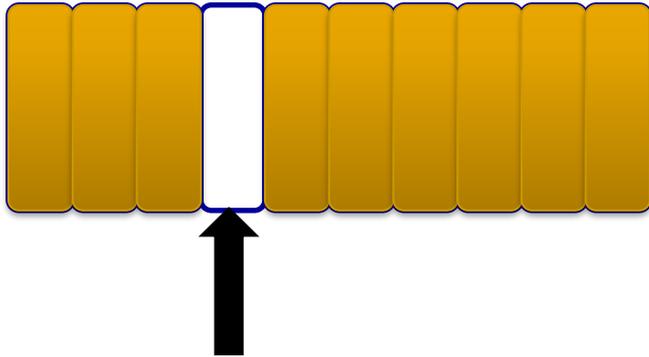
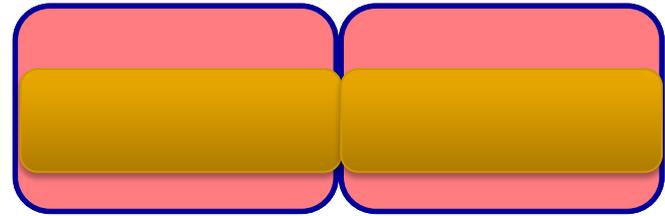
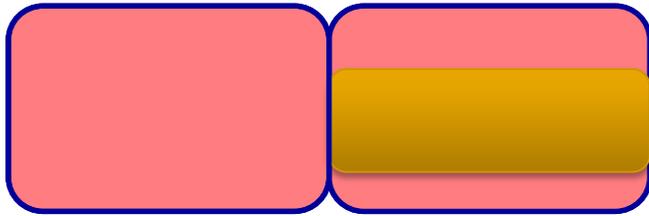
# Watchpoints Globally Useful

- Byte/Word Accurate and Per-Thread



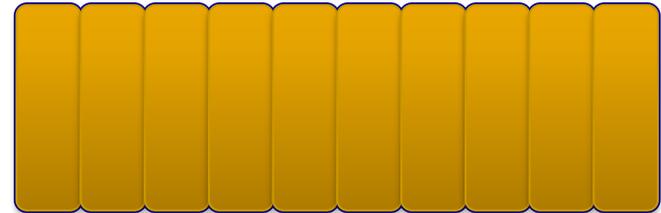
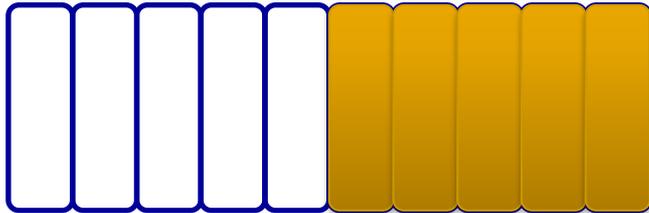
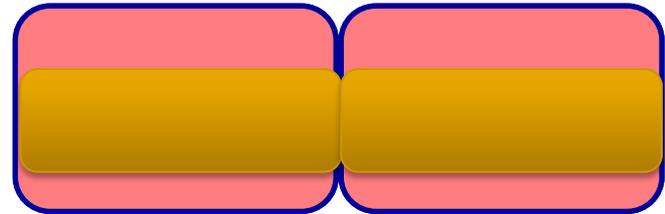
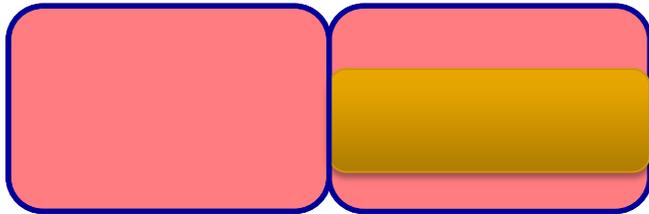
# Watchpoints Globally Useful

- Byte/Word Accurate and Per-Thread



# Watchpoints Globally Useful

- Byte/Word Accurate and Per-Thread



---

# Watchpoint-Based Software Analyses

- Taint analysis / Dynamic dataflow analysis
- Data Race Detection
- Deterministic Execution
- Canary-Based Bounds Checking
- Speculative Program Optimization
- Hybrid Transactional Memory

# Challenges

- Some analyses require watchpoint ranges



- Better stored as base + length

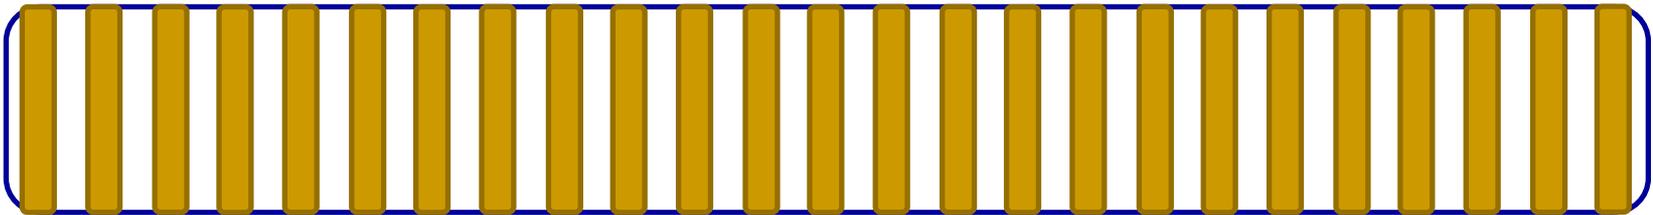
# Challenges

- Some analyses require watchpoint ranges



- Better stored as base + length

- Some need large # of small watchpoints



- Better stored as bitmaps

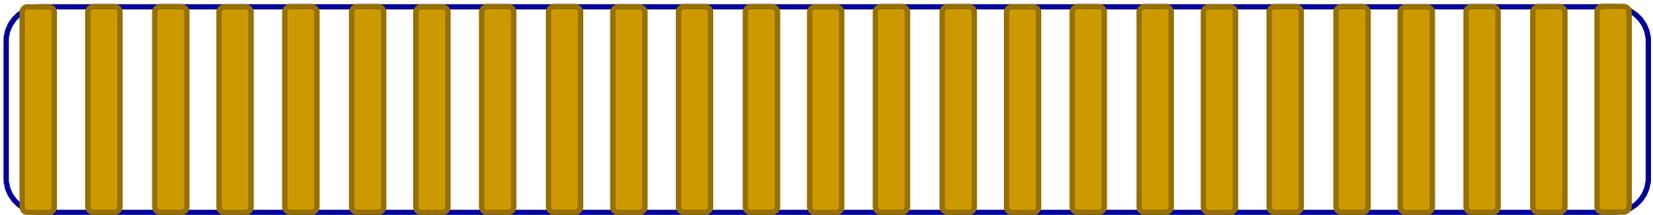
# Challenges

- Some analyses require watchpoint ranges



- Better stored as base + length

- Some need large # of small watchpoints

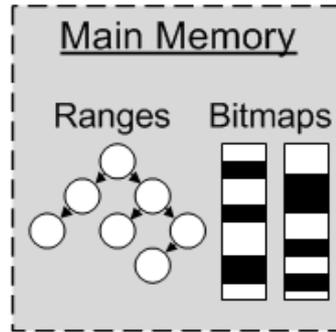


- Better stored as bitmaps

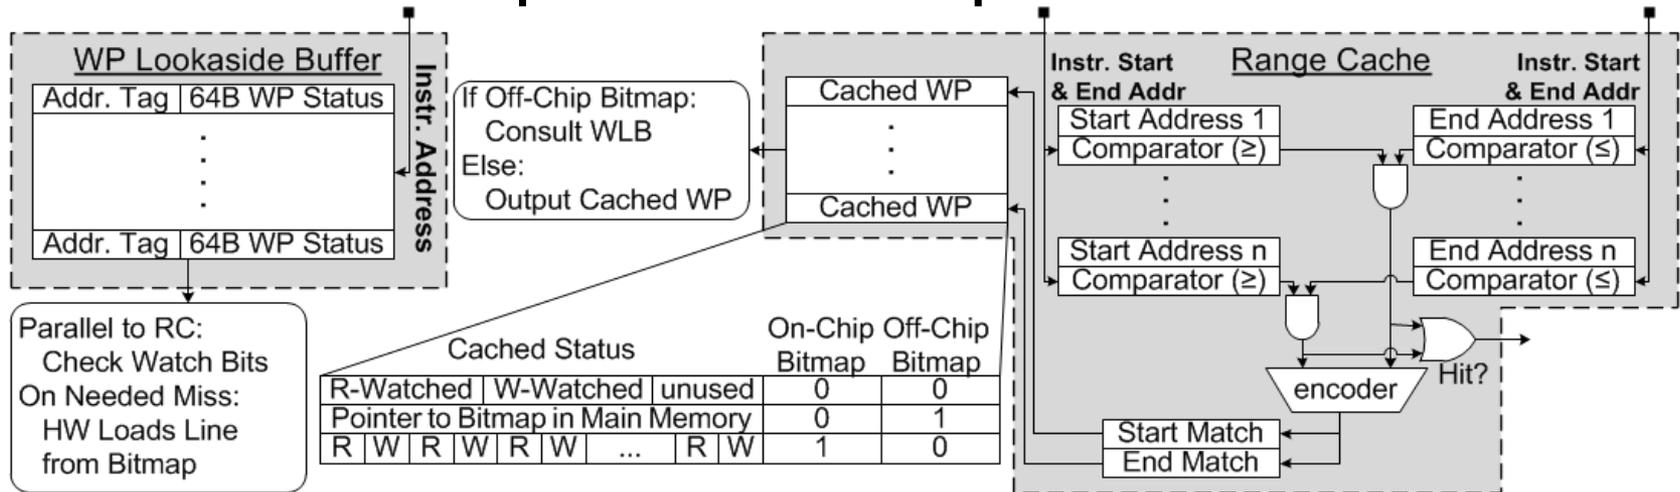
- Need a large number

# The Best of Both Worlds

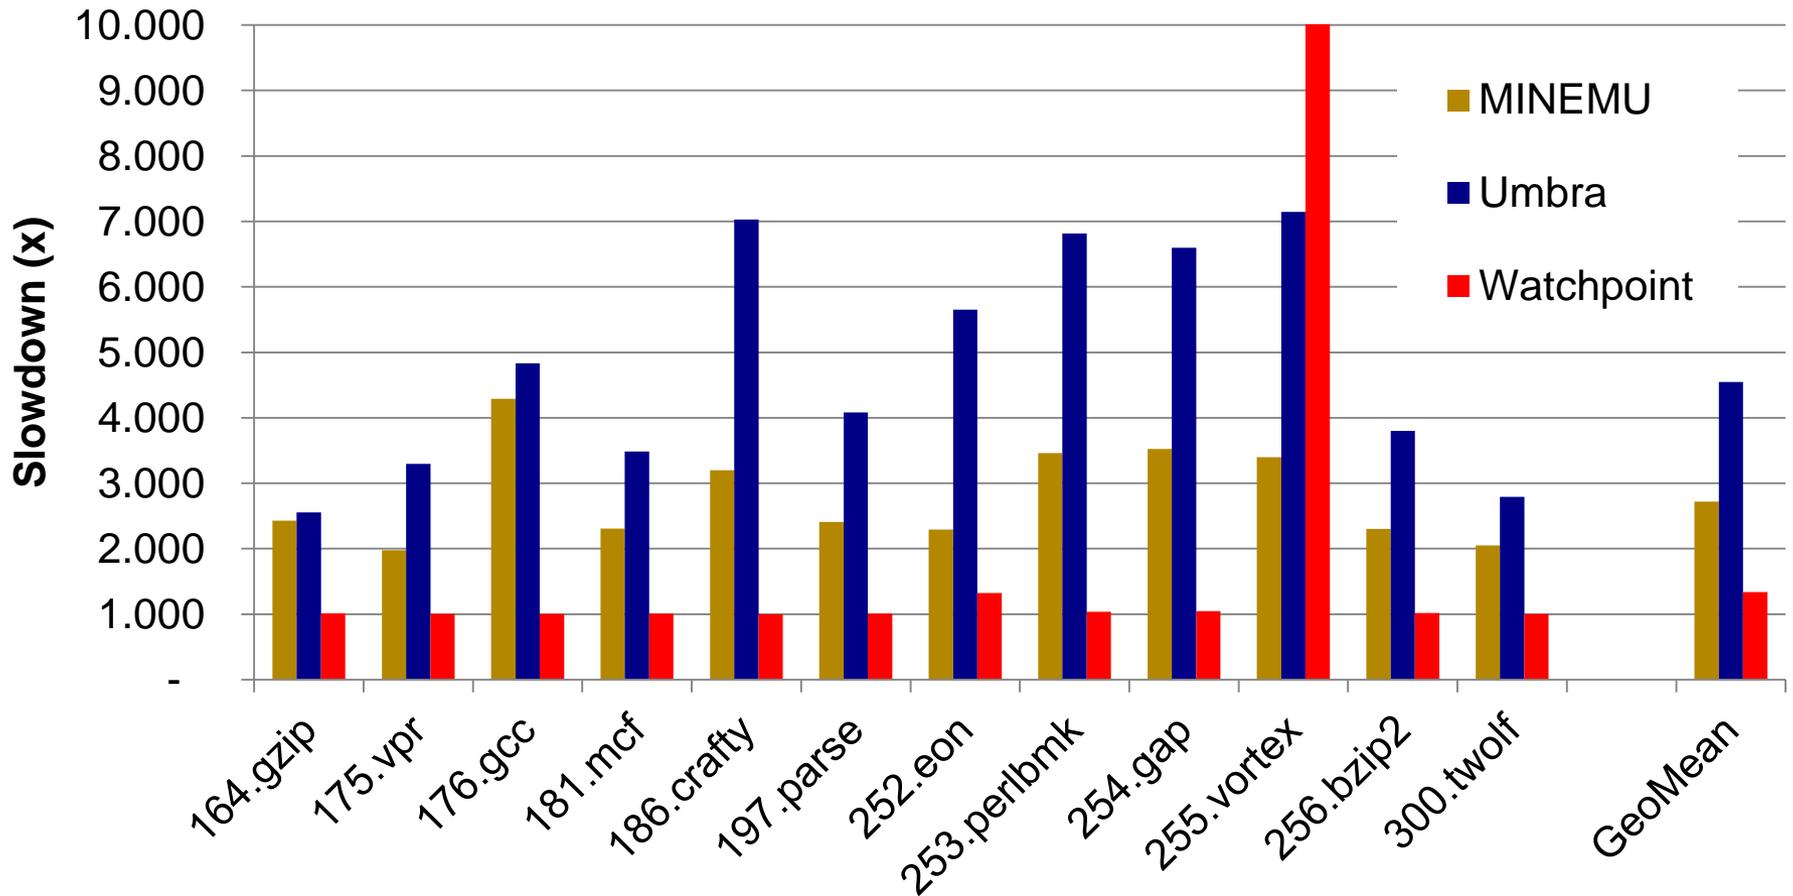
- Store Watchpoints in Main Memory



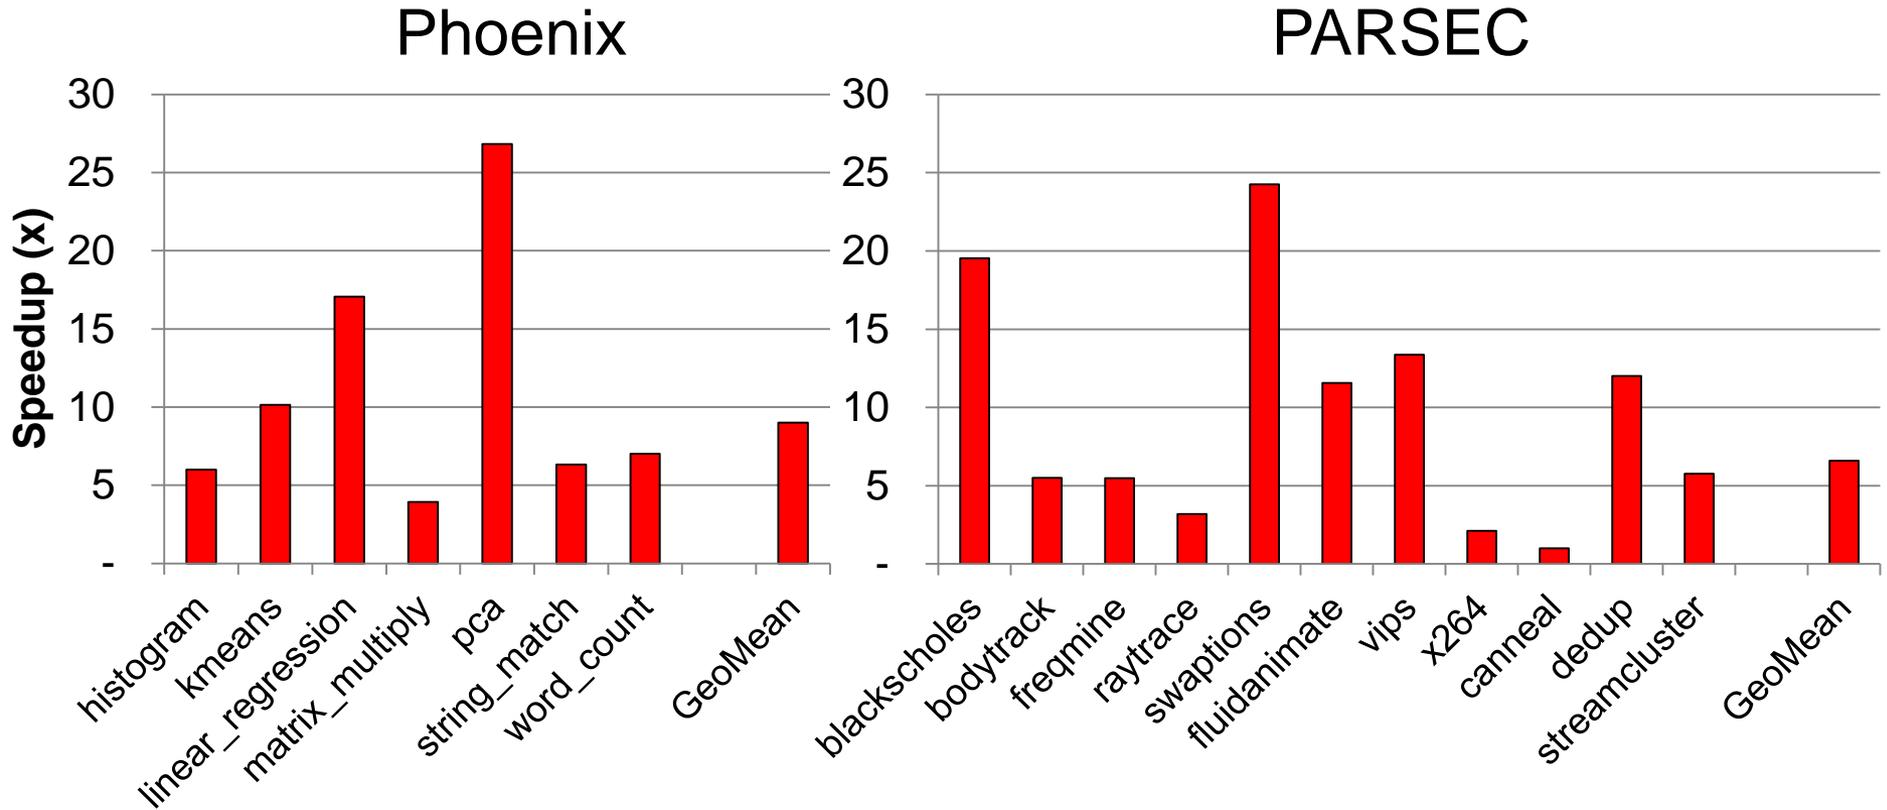
- Cache watchpoints on-chip



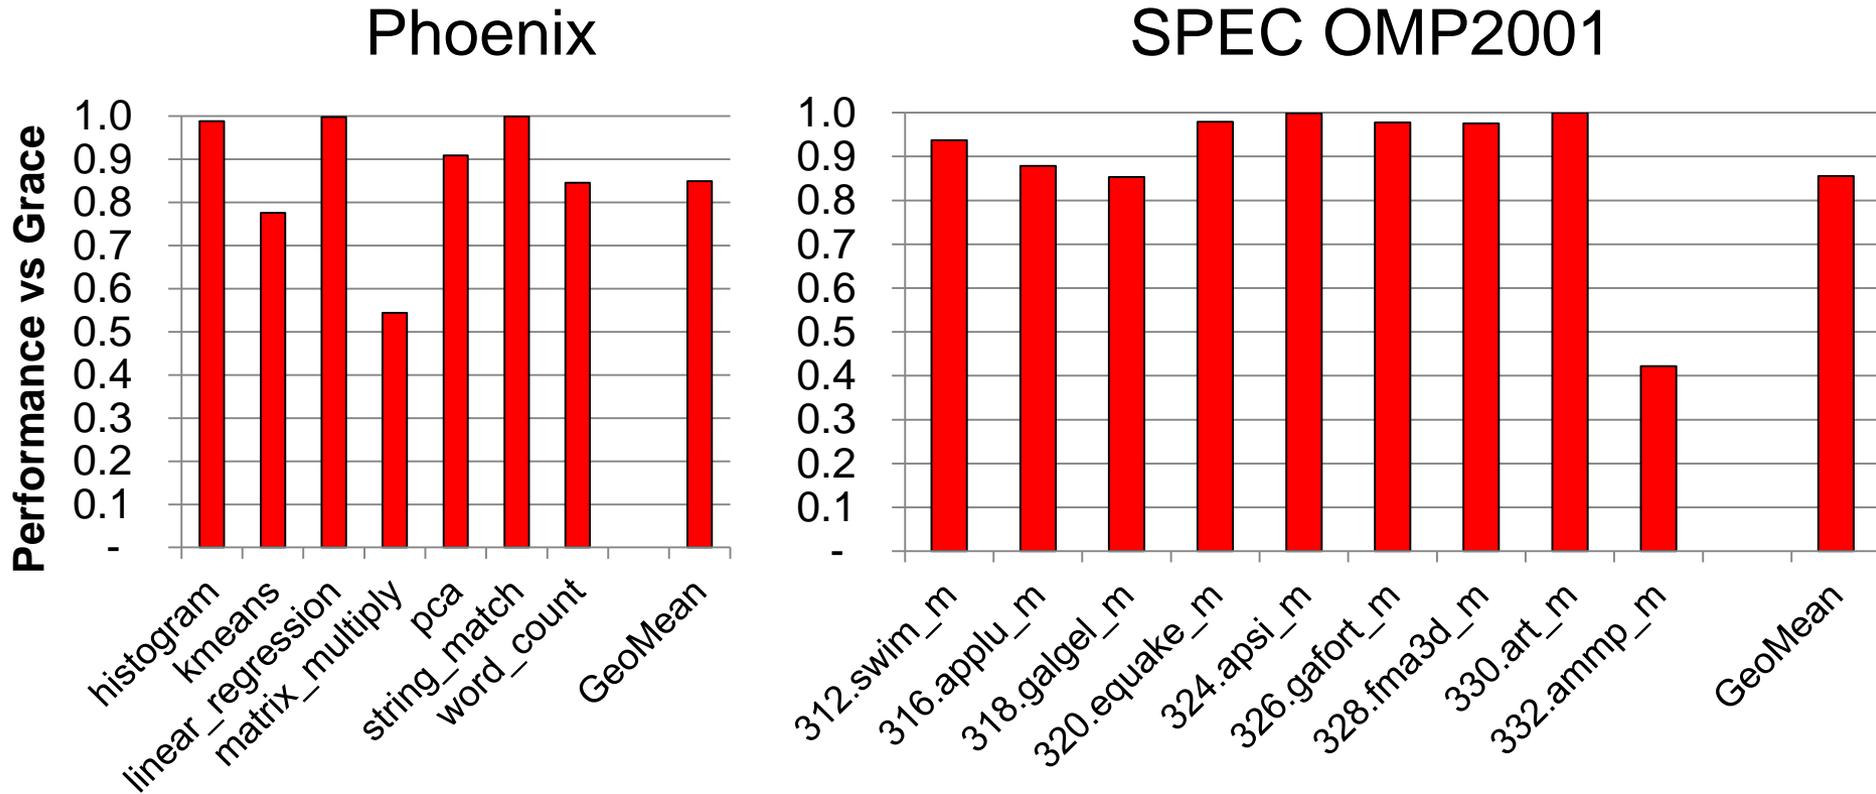
# Demand-Driven Taint Analysis



# Watchpoint-Based Data Race Detection



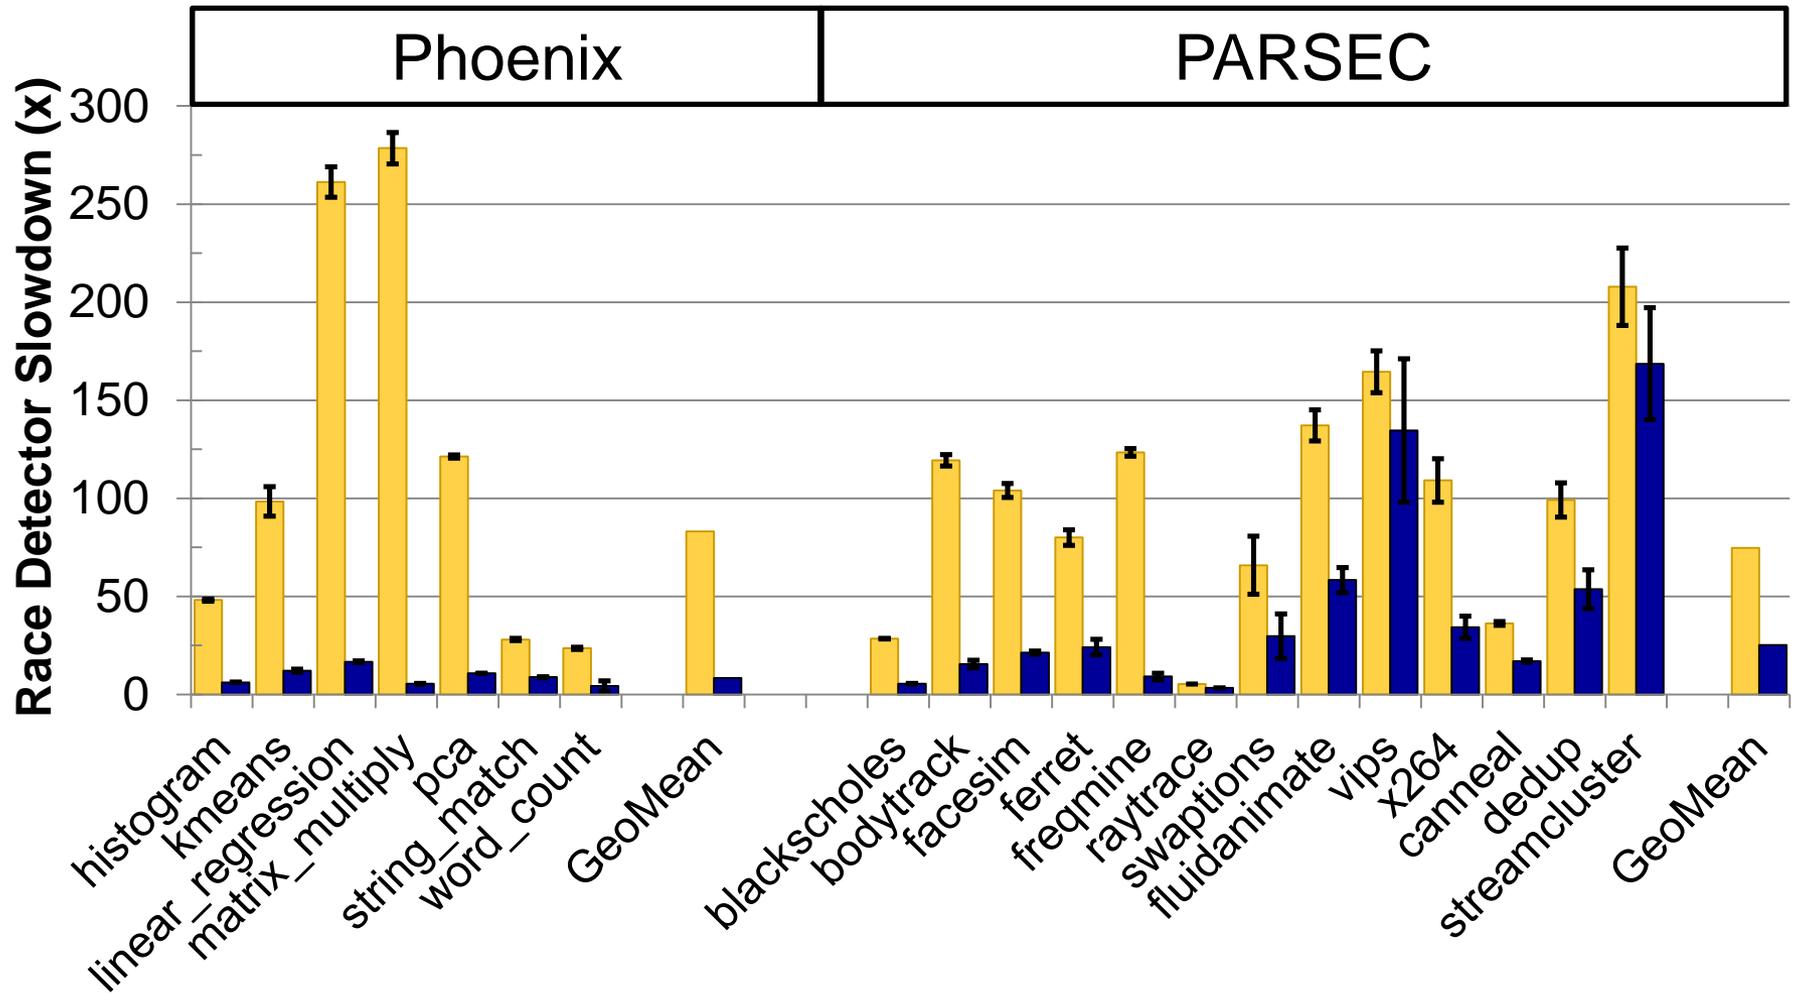
# Watchpoint-Based Grace System



---

# BACKUP SLIDES

# Performance Difference



# Accuracy on Real Hardware

	kmeans	facesim	ferret	freqmine	vips	x264	streamcluster
W→W	1/1 (100%)	0/1 (0%)	-	-	1/1 (100%)	-	1/1 (100%)
R→W	-	0/1 (0%)	2/2 (100%)	2/2 (100%)	1/1 (100%)	3/3 (100%)	1/1 (100%)
W→R	-	2/2 (100%)	1/1 (100%)	2/2 (100%)	1/1 (100%)	3/3/ (100%)	1/1 (100%)

	Spider Monkey-0	Spider Monkey-1	Spider Monkey-2	NSPR-1	Memcached-1	Apache-1
W→W	9/9 (100%)	1/1 (100%)	1/1 (100%)	3/3 (100%)	-	1/1 (100%)
R→W	3/3 (100%)	-	1/1 (100%)	1/1 (100%)	1/1 (100%)	7/7 (100%)
W→R	8/8 (100%)	1/1 (100%)	2/2 (100%)	4/4 (100%)	-	2/2 (100%)

---

# Width Test

