

# Enhancing Atomic Instruction Emulation for Cross-ISA Dynamic Binary Translation

Ziyi Zhao, Zhang Jiang,  
Ying Chen, Xiaoli Gong\*  
*Nankai University*



Wenwen Wang  
*University of Georgia*



**UNIVERSITY OF  
GEORGIA**

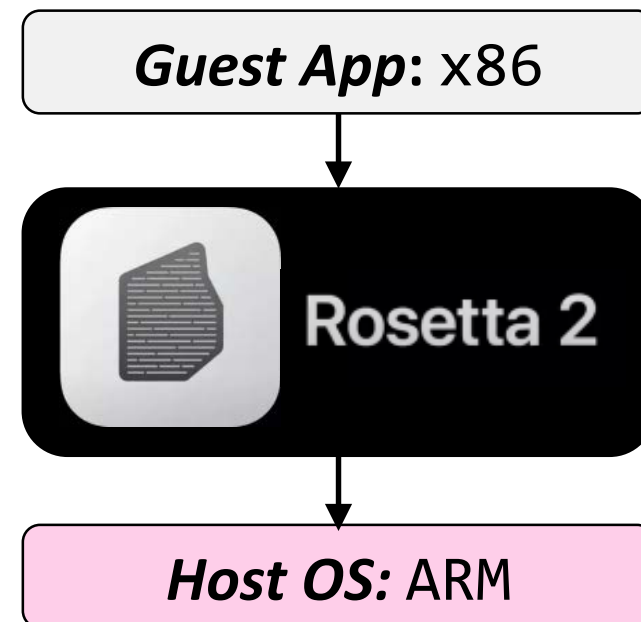
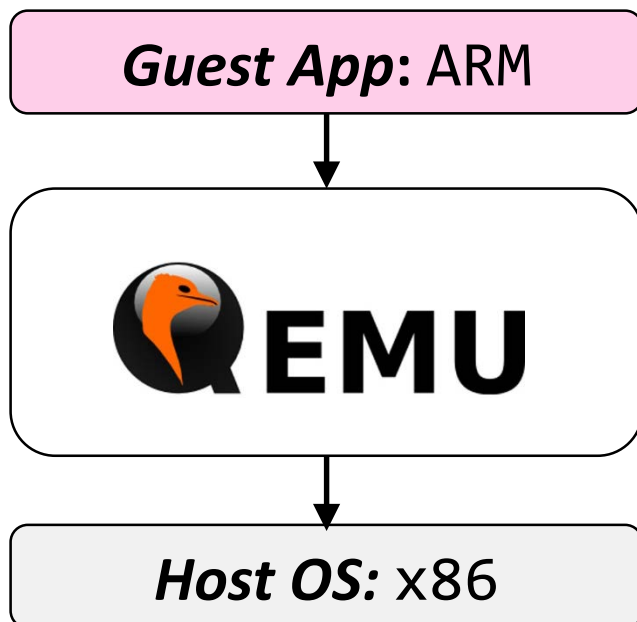
Pen-Chung Yew  
*University of Minnesota  
at Twin Cities*



**UNIVERSITY OF MINNESOTA**  
**Driven to Discover<sup>SM</sup>**

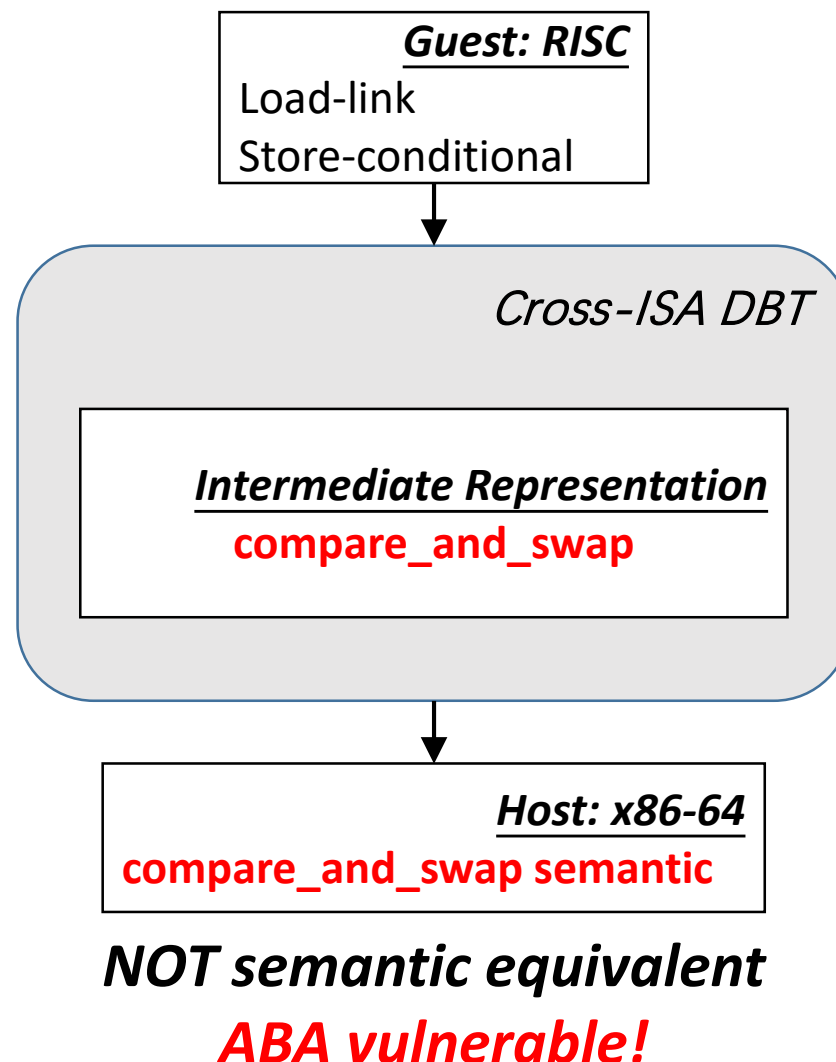
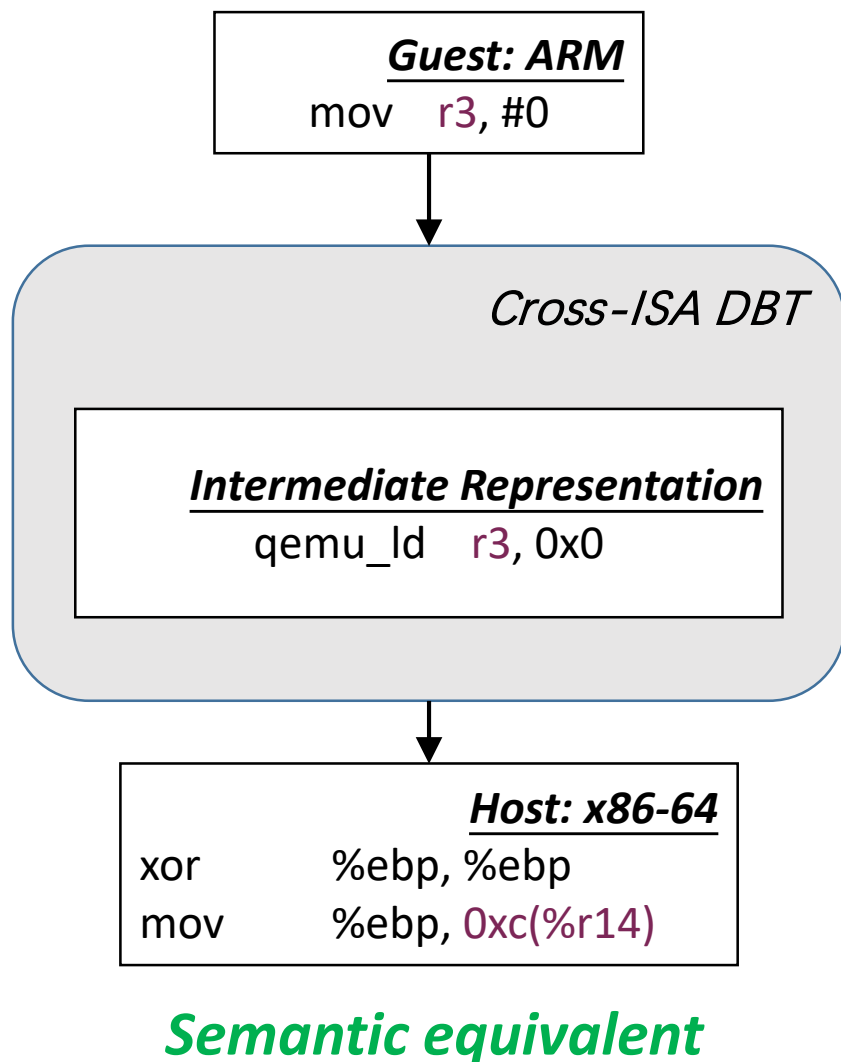
# Cross-ISA Dynamic Binary Translation

- Cross Instruction Set Architecture Dynamic Binary Translator
- A key enabling technology
  - Developing and testing
  - Running application



# Cross-ISA Dynamic Binary Translation

- Fundamental rule: semantic equivalence



# Outline

- Motivation
- Background
  - RISC atomic instruction
  - ABA problem
- Design
  - Challenge
  - Proposed solution
- Evaluation
- Conclusion

# RISC Atomic Instruction: Load-link/Store-conditional

- Load-Link(LL)

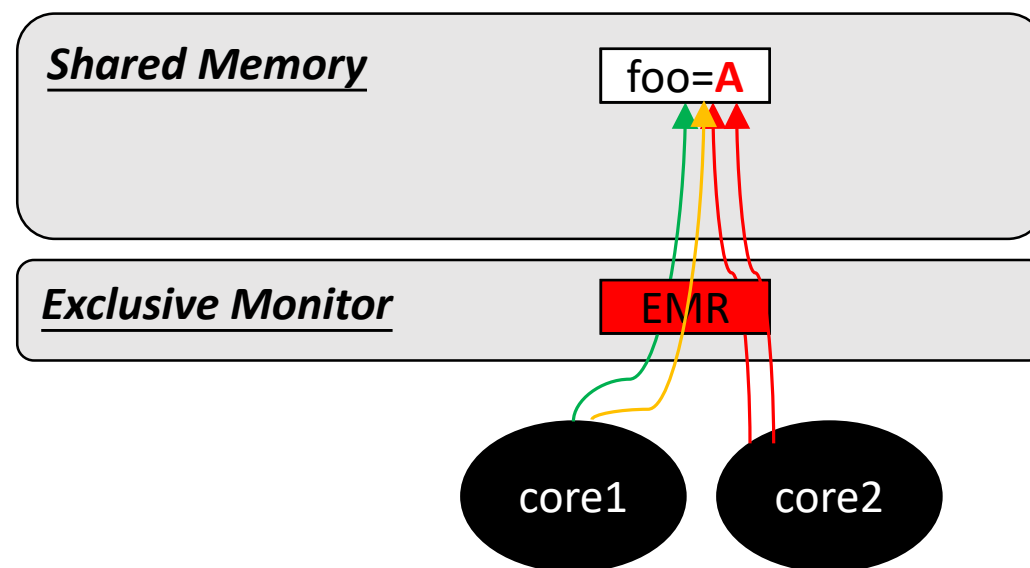
- Set Exclusive Memory Region (EMR)
- Load data

- Sore-Conditional(SC)

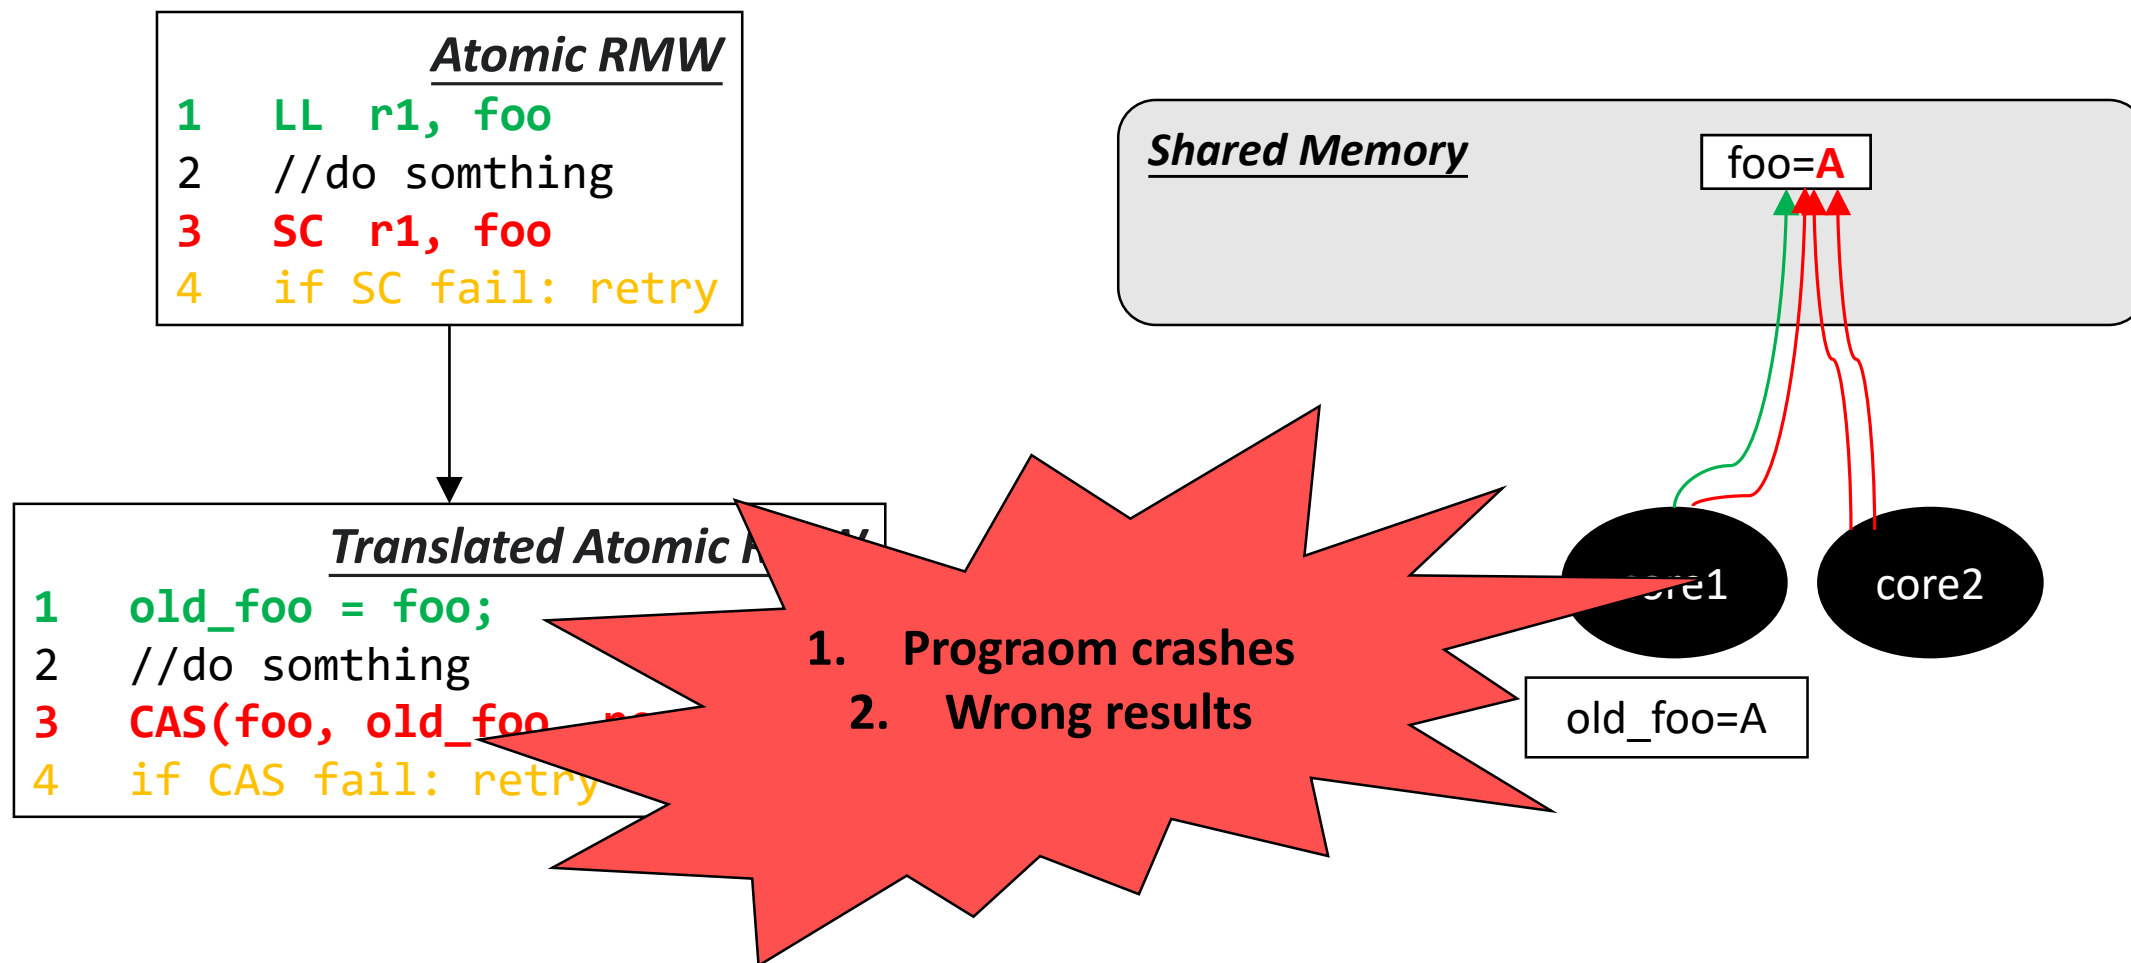
- Check EMR
- Write data

```

Atomic RMW
1  LL  r1, foo
2  //do something
3  SC  r1, foo
4  if SC fail: retry
  
```




# ABA Problem




# Design Goal: Atomicity of LL/SC

- Strong atomicity

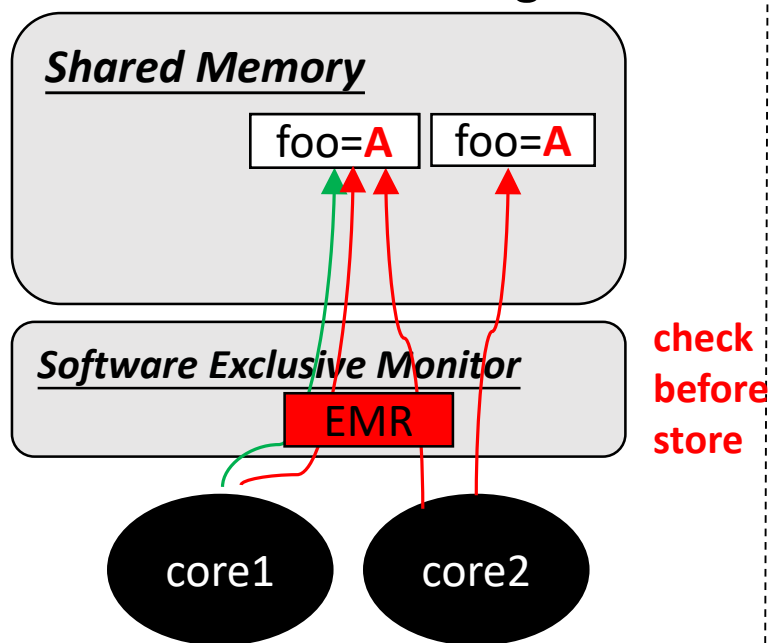


- Weak atomicity



# Previous Ideas to Correctly Emulate LL/SC


## Software Checking



All stores are instrumented!!  
HUGE overhead

Can be fast with lock-free  
hash table

## Let MMU to do the check

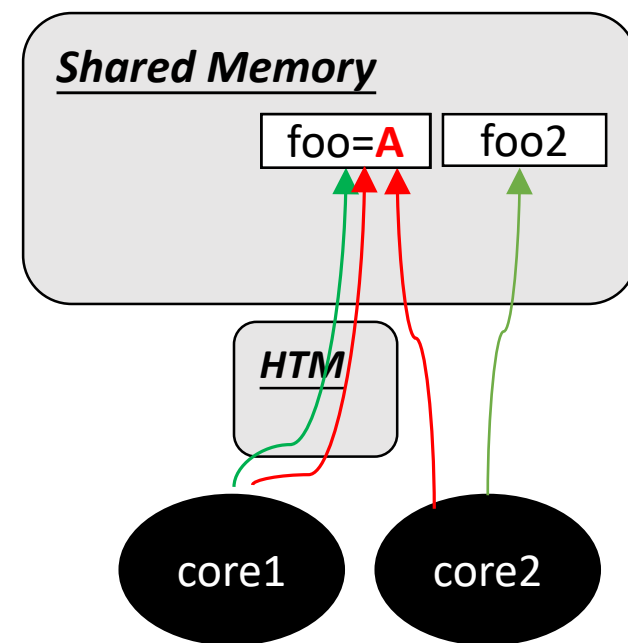


Only stores to the RO region are  
intercepted

**BUT REALLY?**

But false sharing + large  
syscall overhead!

## Hardware Transaction Memory



Fast and correct!

Could be challenging to  
implement in DBT!




# A. HST: Hash-table Based Store Test

- Challenge: store instrumentation overhead
  - store instructions account for up to 20% instructions, highly concurrent
  - **Be Fast**: lightweight checking procedure
  - **Be Scalable**: no locking
- Solution: Lock-free hash table

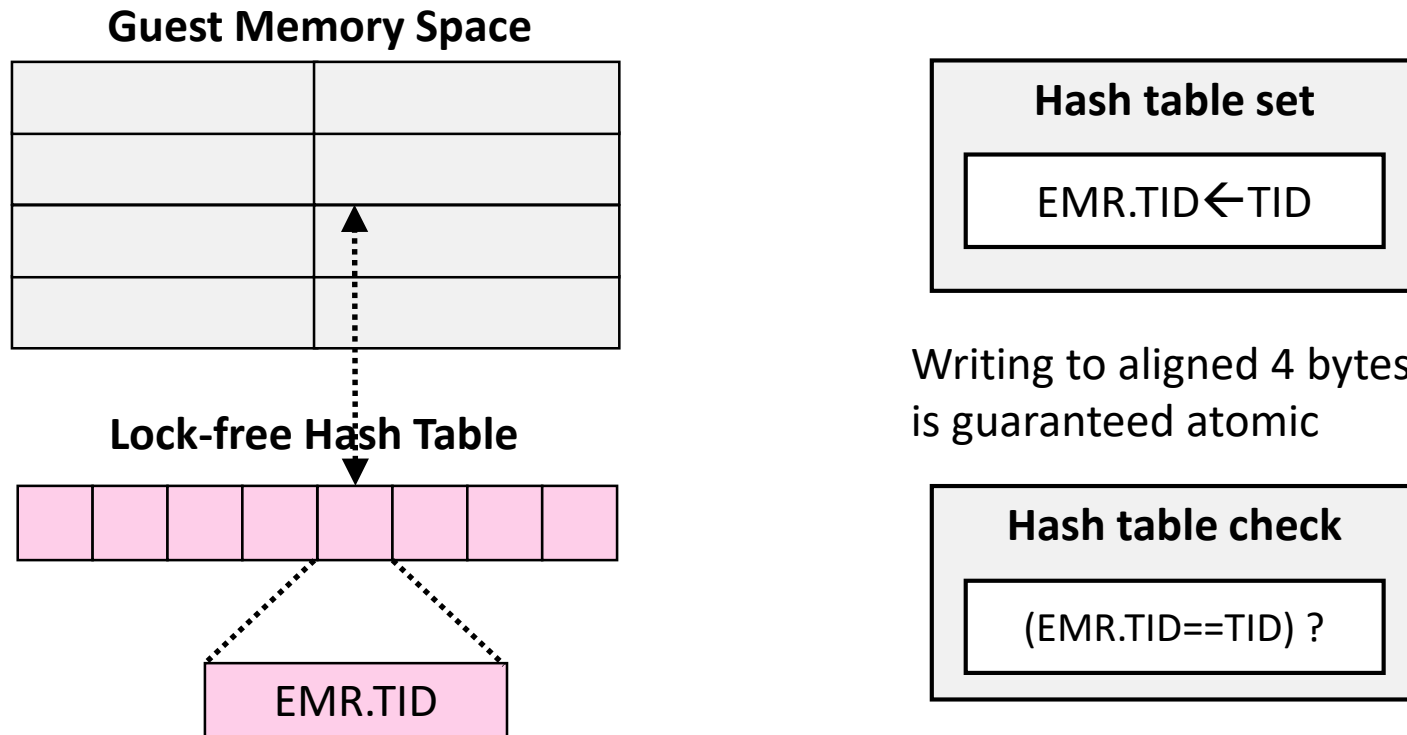
<u>Translated Store</u>	
1	find_EMR(mem)
2	lock
3	update_EMR
4	unlock
5	store reg, mem

20%~45%  
instrumentation  
overhead



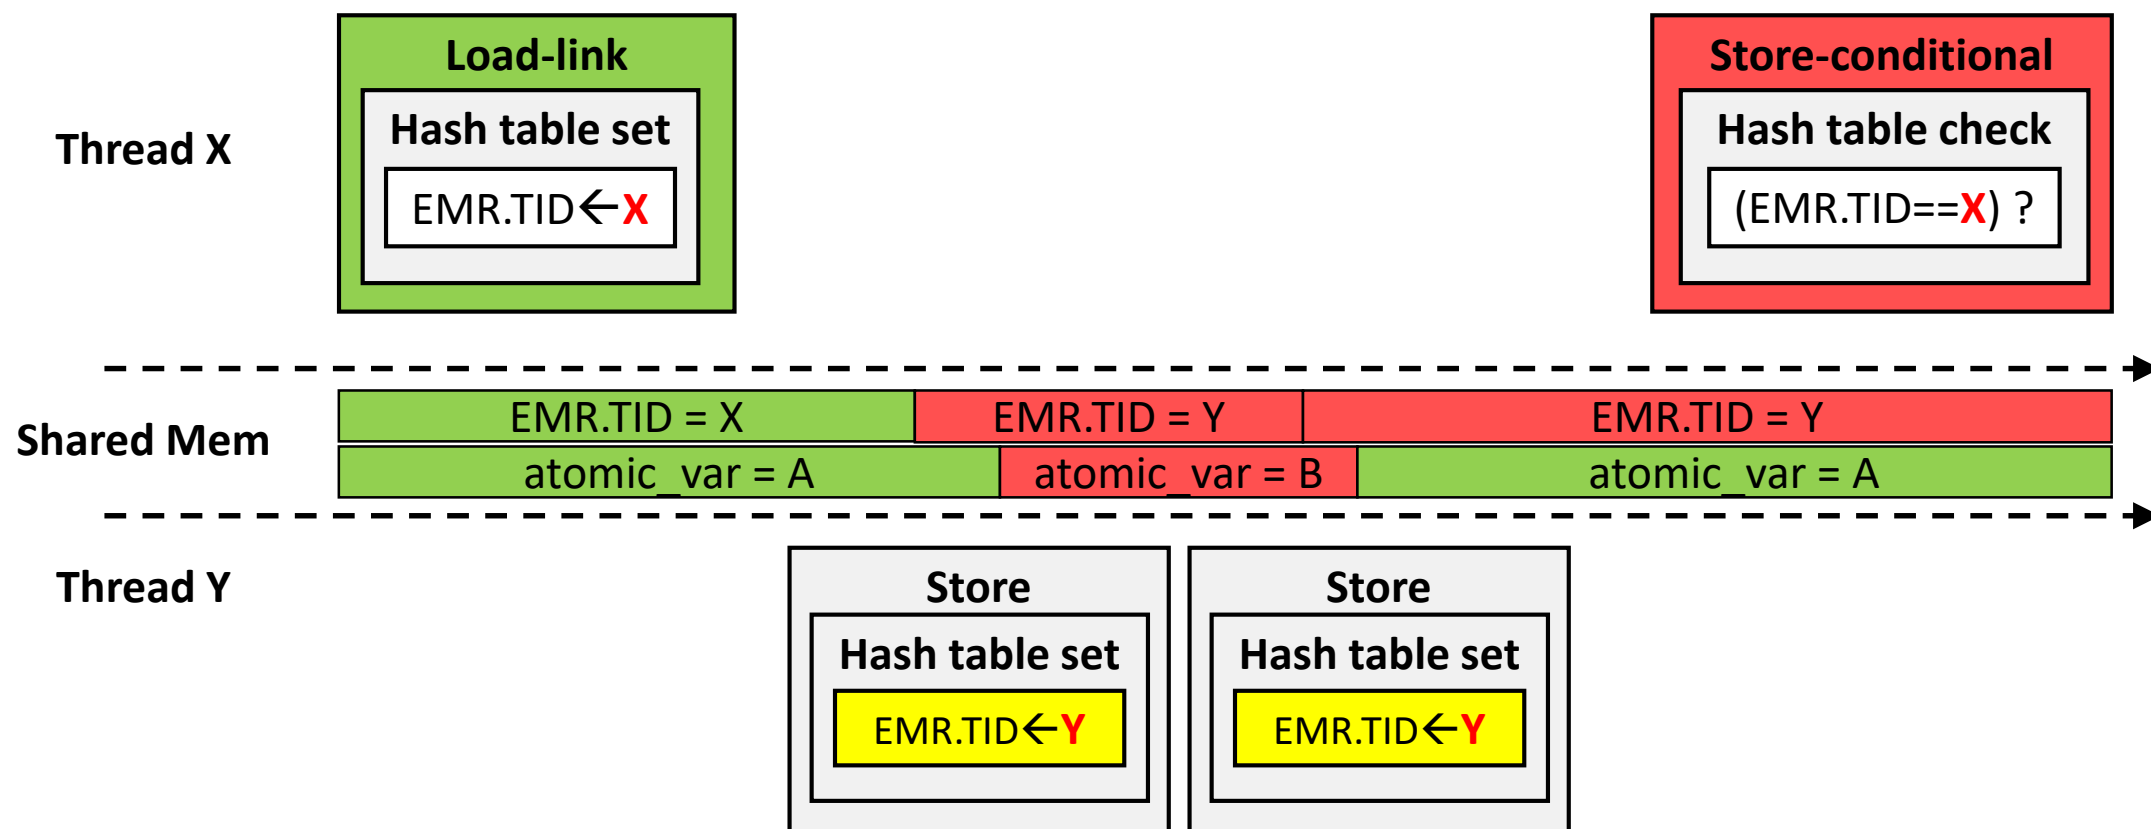
# A. HST: Hash-table Based Store Test

- Challenge: store instrumentation overhead
- Solution: Lock-free hash table



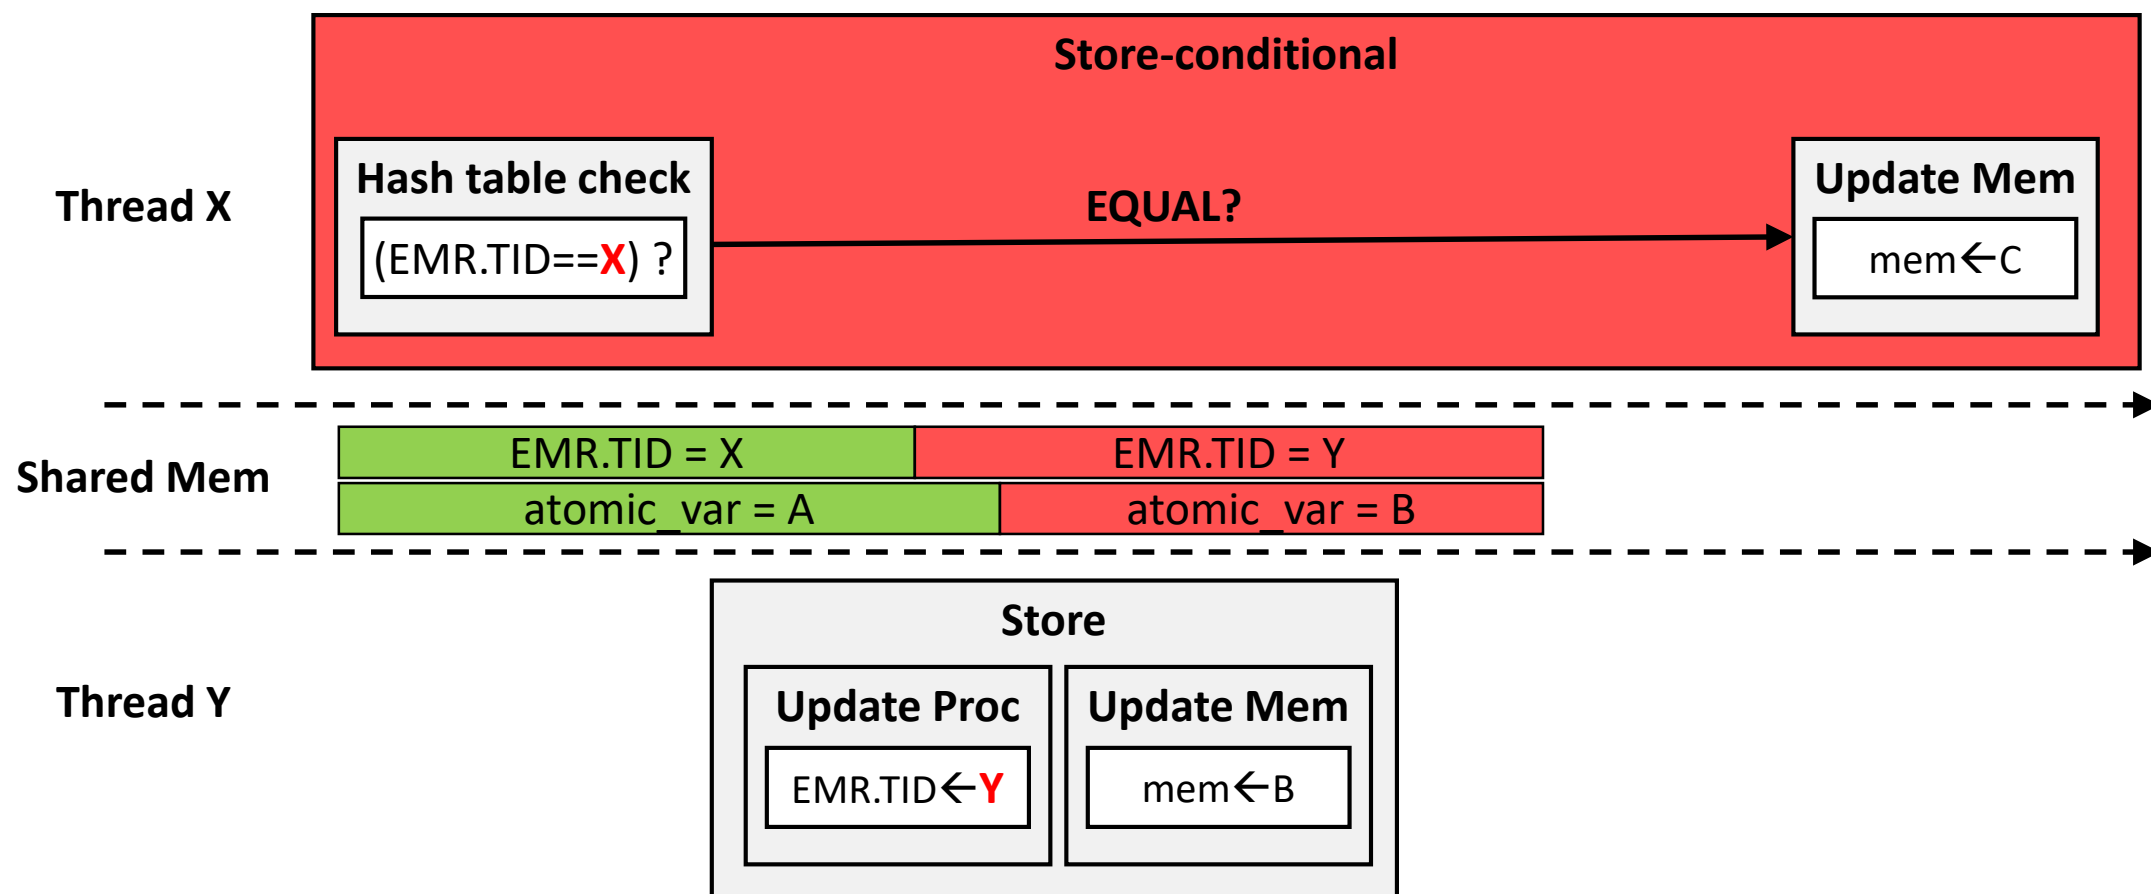
# A. HST: Hash-table Based Store Test

- Challenge: store instrumentation overhead
- Solution: Lock-free hash table



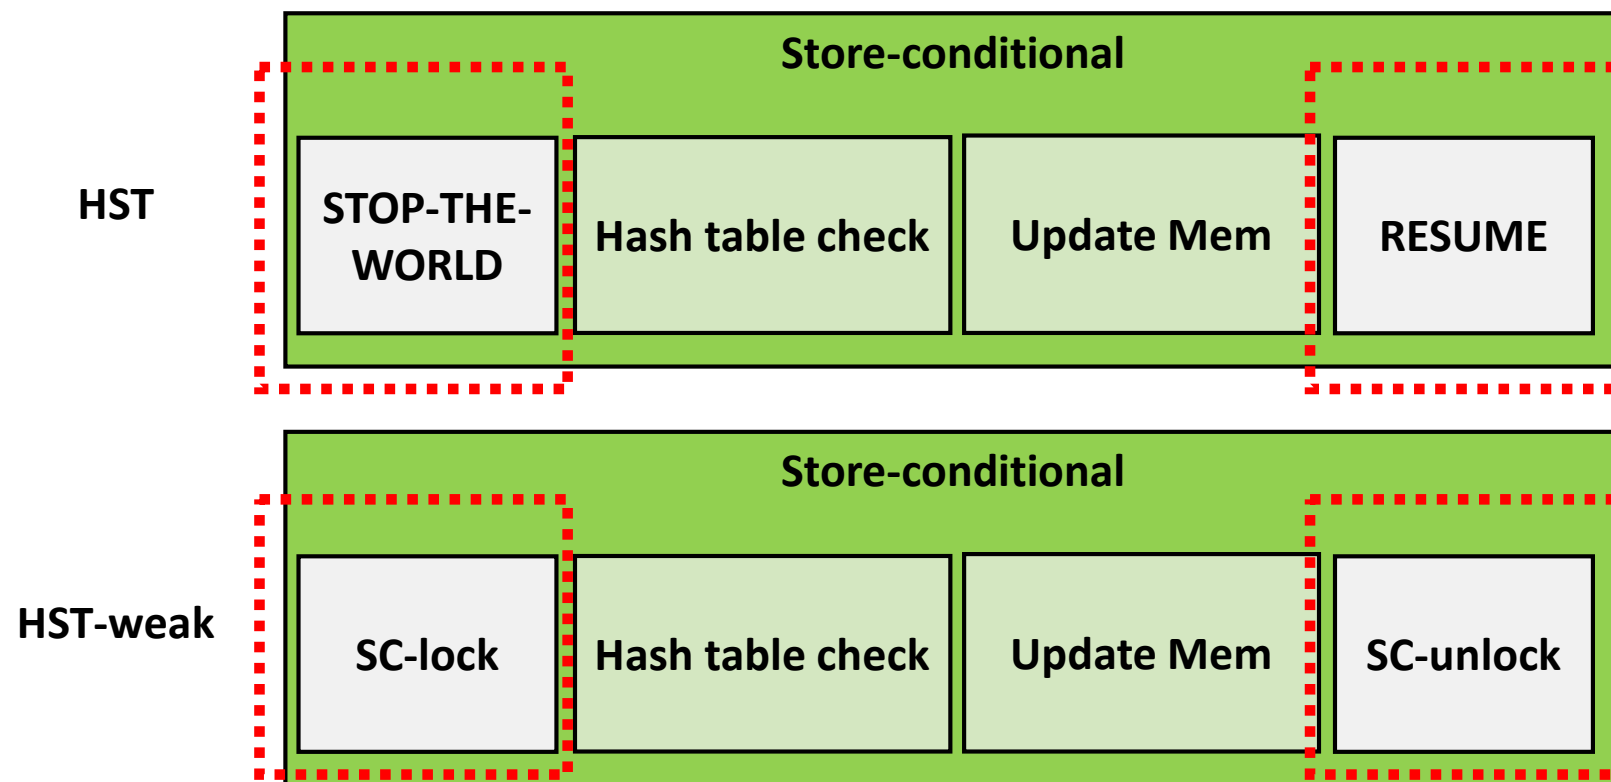
# A. HST: Hash-table Based Sore Test

- Challenge: race condition between SC and store



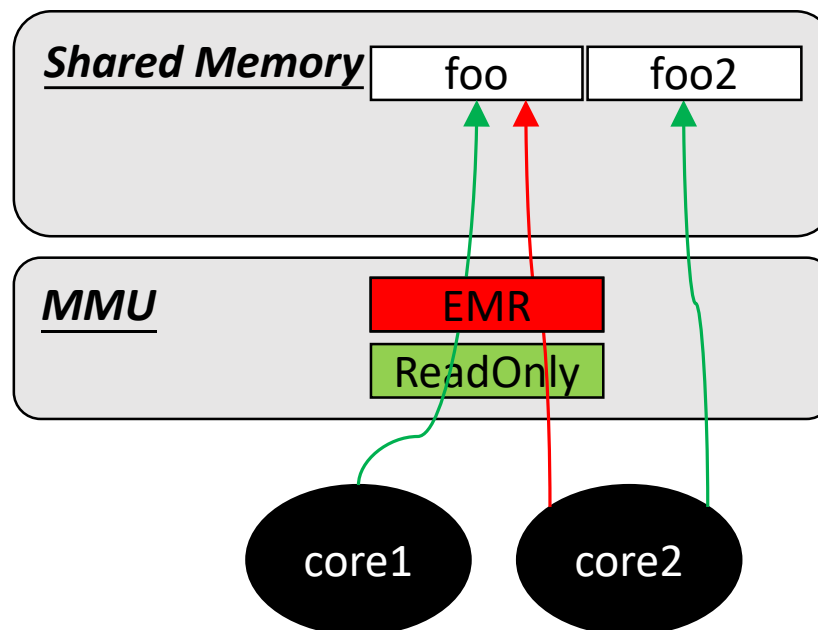
# A. HST: Hash-table Based Sore Test

- Solution: Guaranteeing the atomicity of SC
  - HST: stop-the-world → strong atomicity
  - HST-weak: SC-mutex-lock → weak atomicity



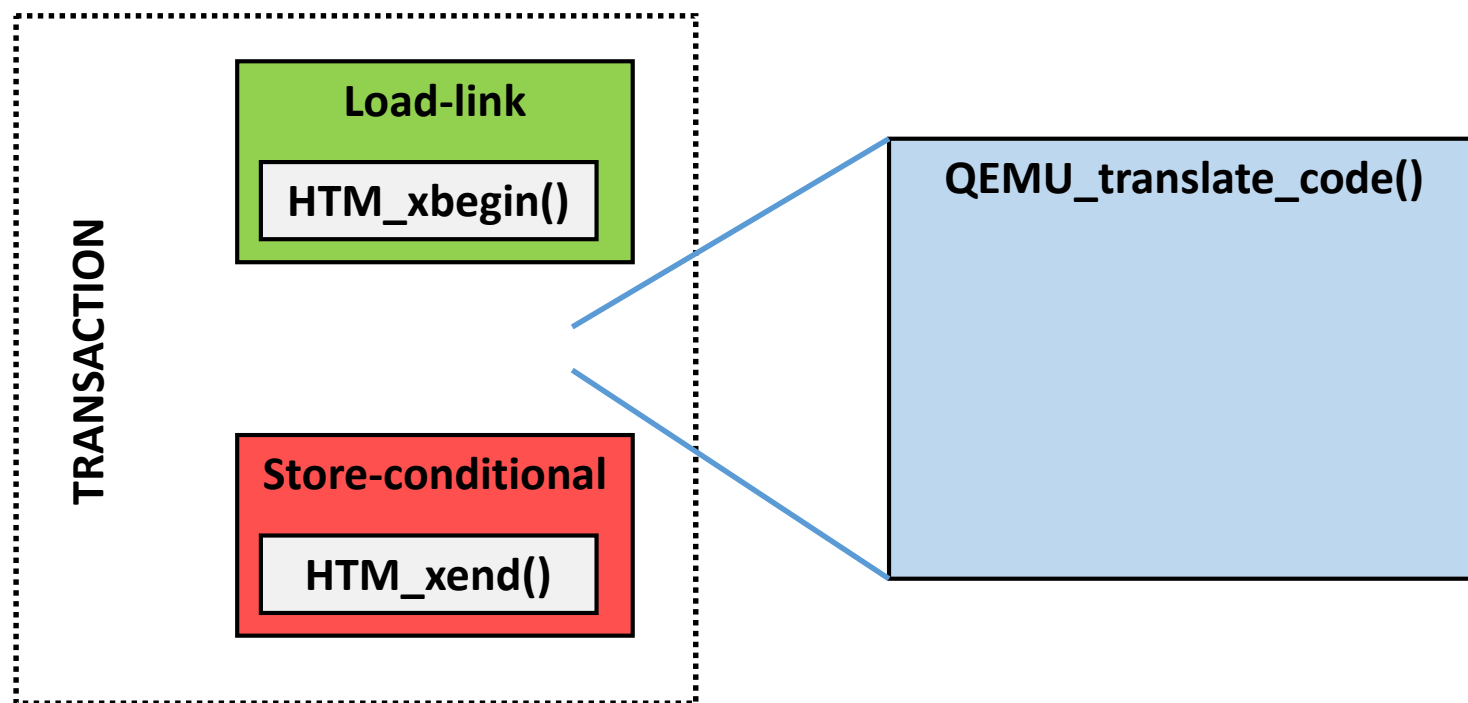
## B. PST: Page Protection Based Store Test

- Key idea: Use MMU to automatically check stores
  - No store instrumentation
- Store-conditional atomicity
  - PST: stop-the-world
  - PST-remap: remapping virtual page to achieve privilege separation



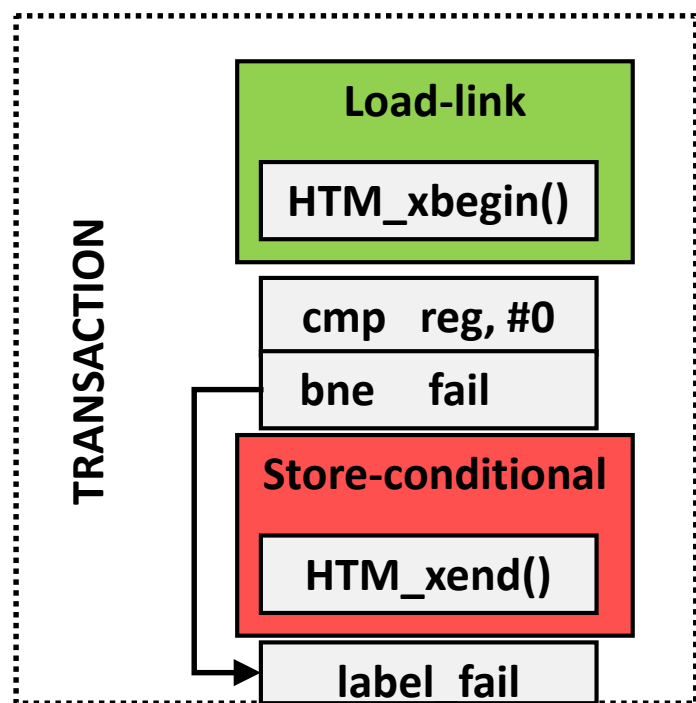
# C. HTM: Hardware Transactional Memory

- Map the LL/SC to a transaction ?
  - Code translation between emulation → large footprint!



# C. HTM: Hardware Transactional Memory

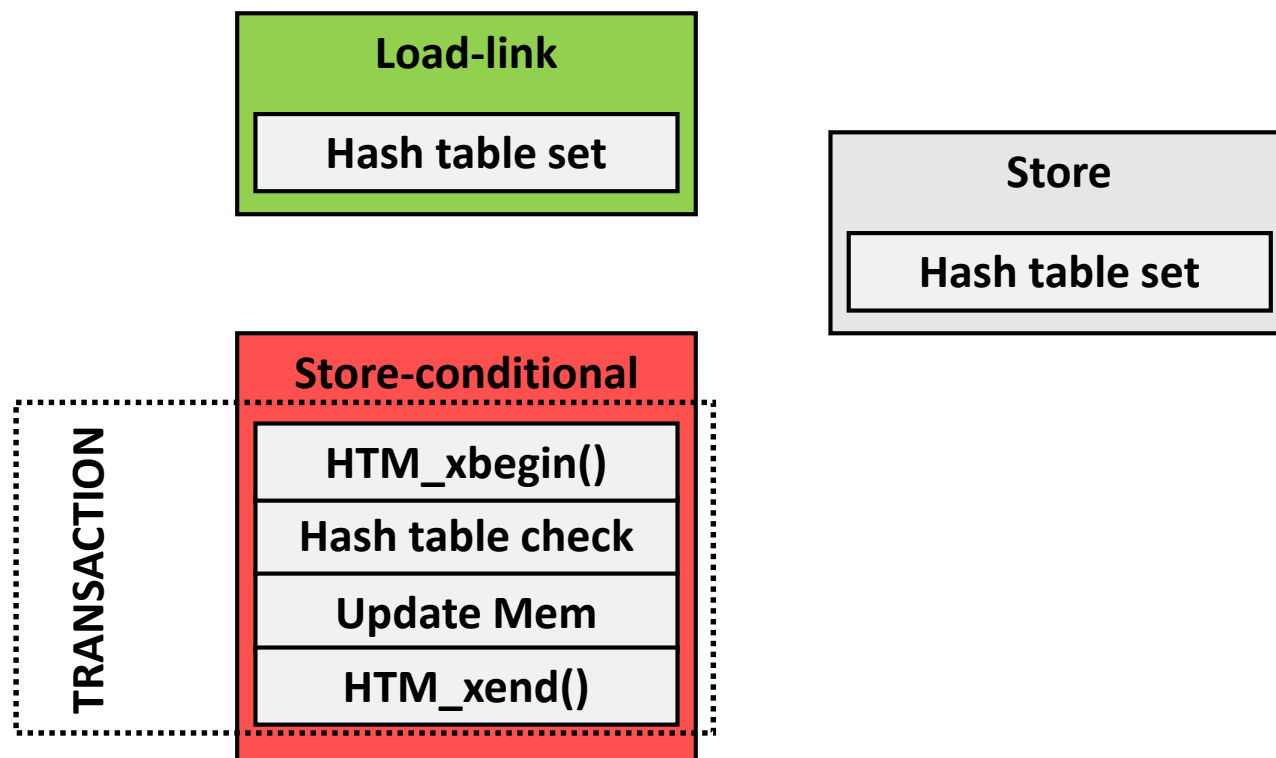
- Map the LL/SC to a transaction ?
  - Code translation between emulation → large footprint!
  - SC could be jumped over → No one ends transaction!





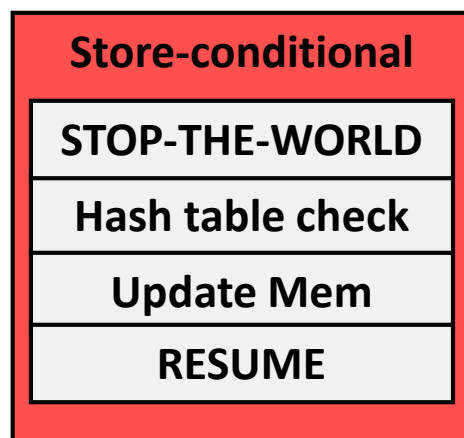

# C. HST-HTM

- Solution: Combining HST and HTM together
  - Small footprint
  - `HTM_xend()` is guaranteed

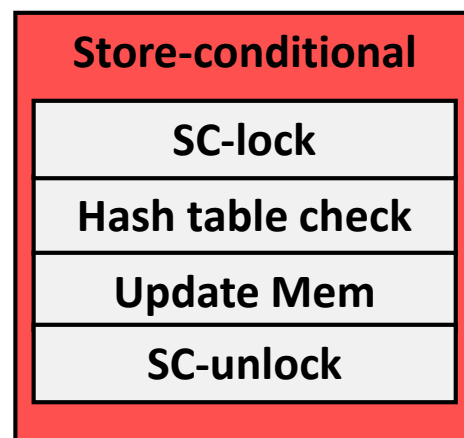



# Summary of Design

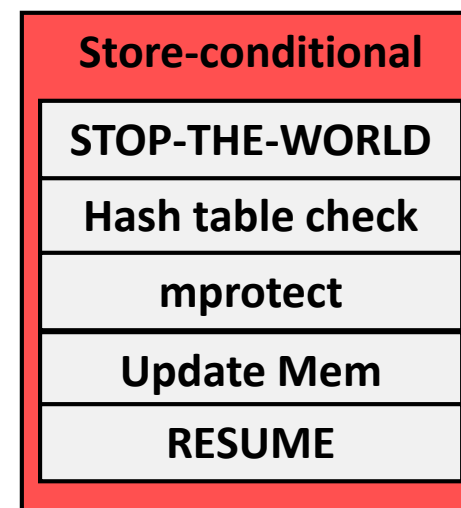


## HST



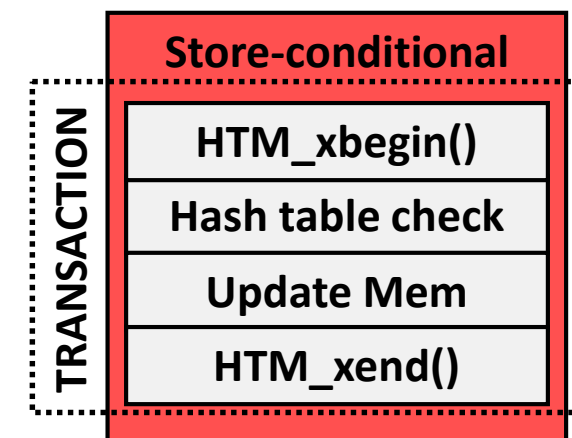


## HST-weak



## PST



## HST-HTM



# Questions to Answer for Evaluation

- Conventional assumptions
  - Save old value and check on store (QEMU) – Fast but incorrect?
  - LL/SC via MMU (PST) – Fast because not instrumenting all stores?
  - Link helpers and instrumented stores (Pico-ST, HST) - Big performance impact by having a helper run for every store?
- Are the schemes scalable in multi-threaded programs?
- What is the performance bottleneck?
- Best trade-off between correctness, speed, and portability?

# Setup

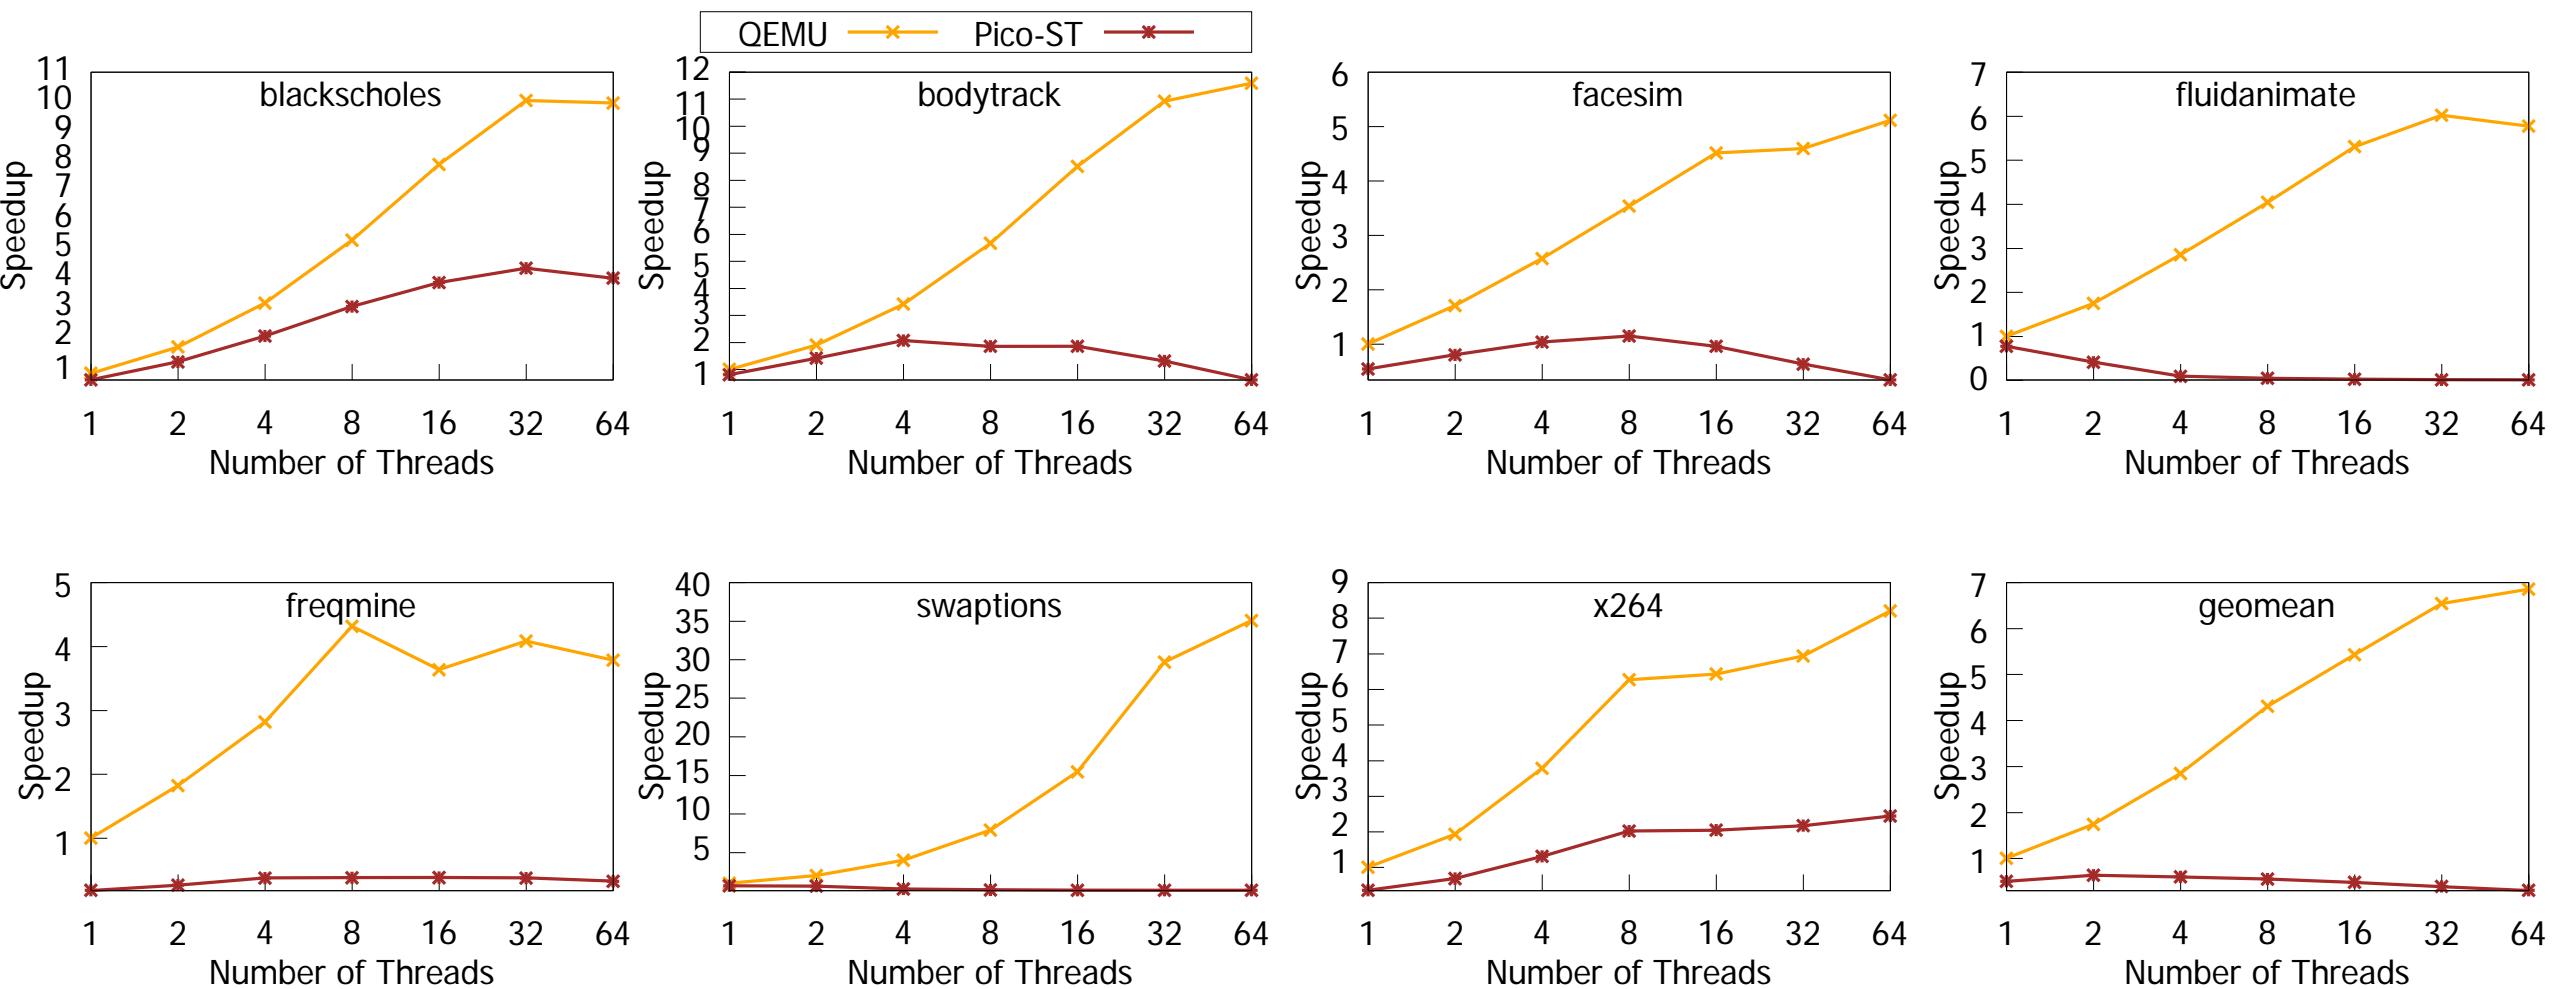
- Running ARM PARSEC benchmark suite on x86 machine
  - PARSEC version 3.0
  - Input size: *simlarge*
- Machine A
  - 52 cores + 183GB Memory
- Machine B: Intel TSX support
  - 10 cores + 64GB Memory

# Correctness

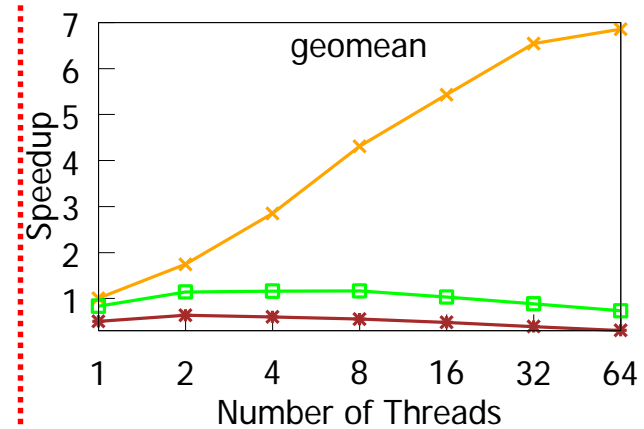







- Formula proof (details in paper)
- Experimental proof
  - A well-designed ARM lock-free stack <https://github.com/NKU-EmbeddedSystem/lock-free-stack-arm-asm>

Native ARM	QEMU 4.1	Pico-ST	HST	HST-weak	PST	PST-remap	HST-HTM	HTM
Pass	Crash	Pass	Pass	Pass	Pass	Pass	Pass	livelock

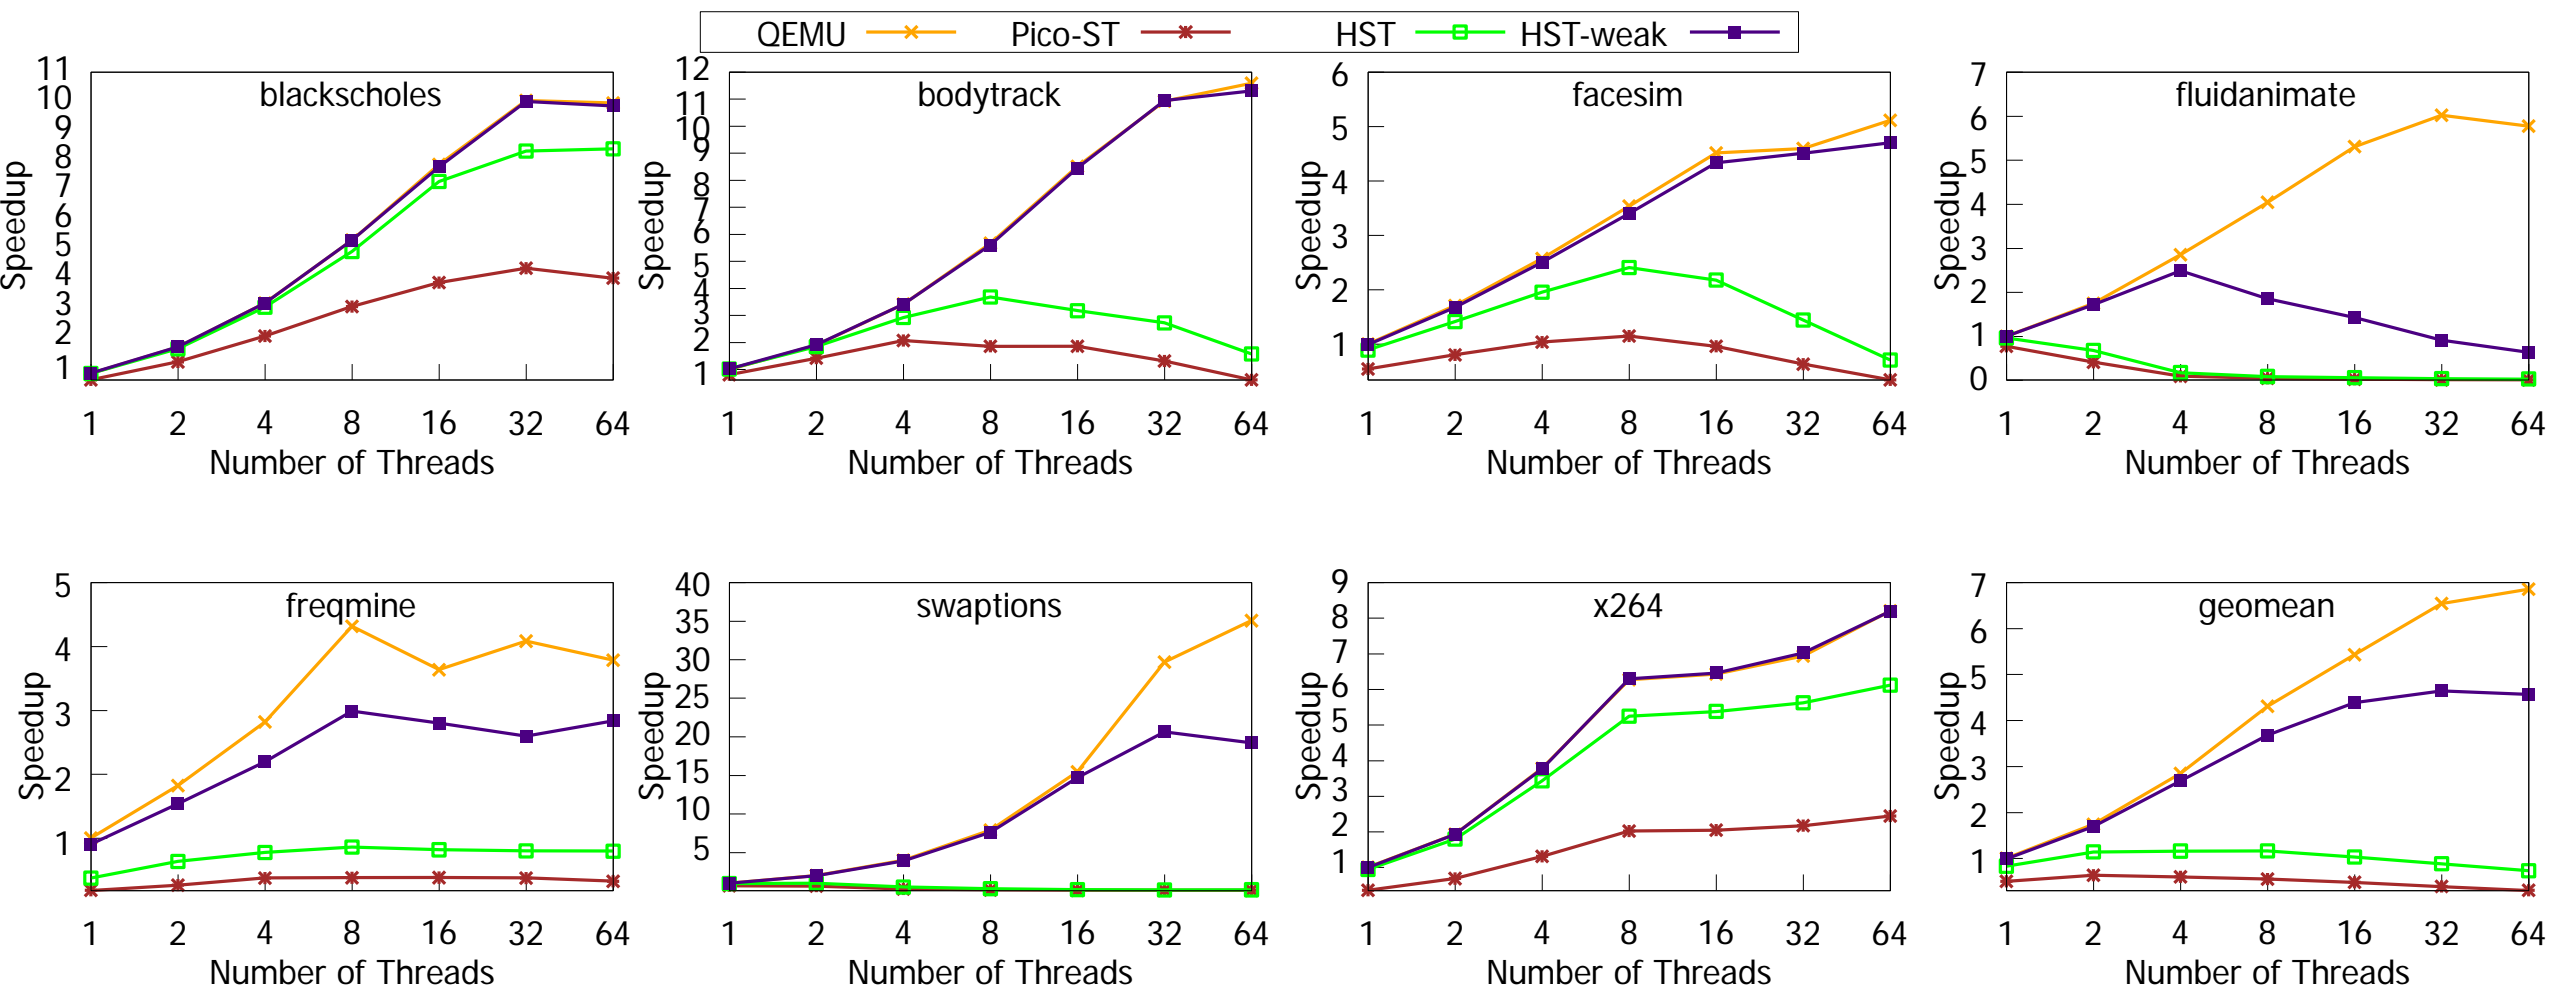
# Scalability – Portable Solutions



# Scalability – Portable Solutions

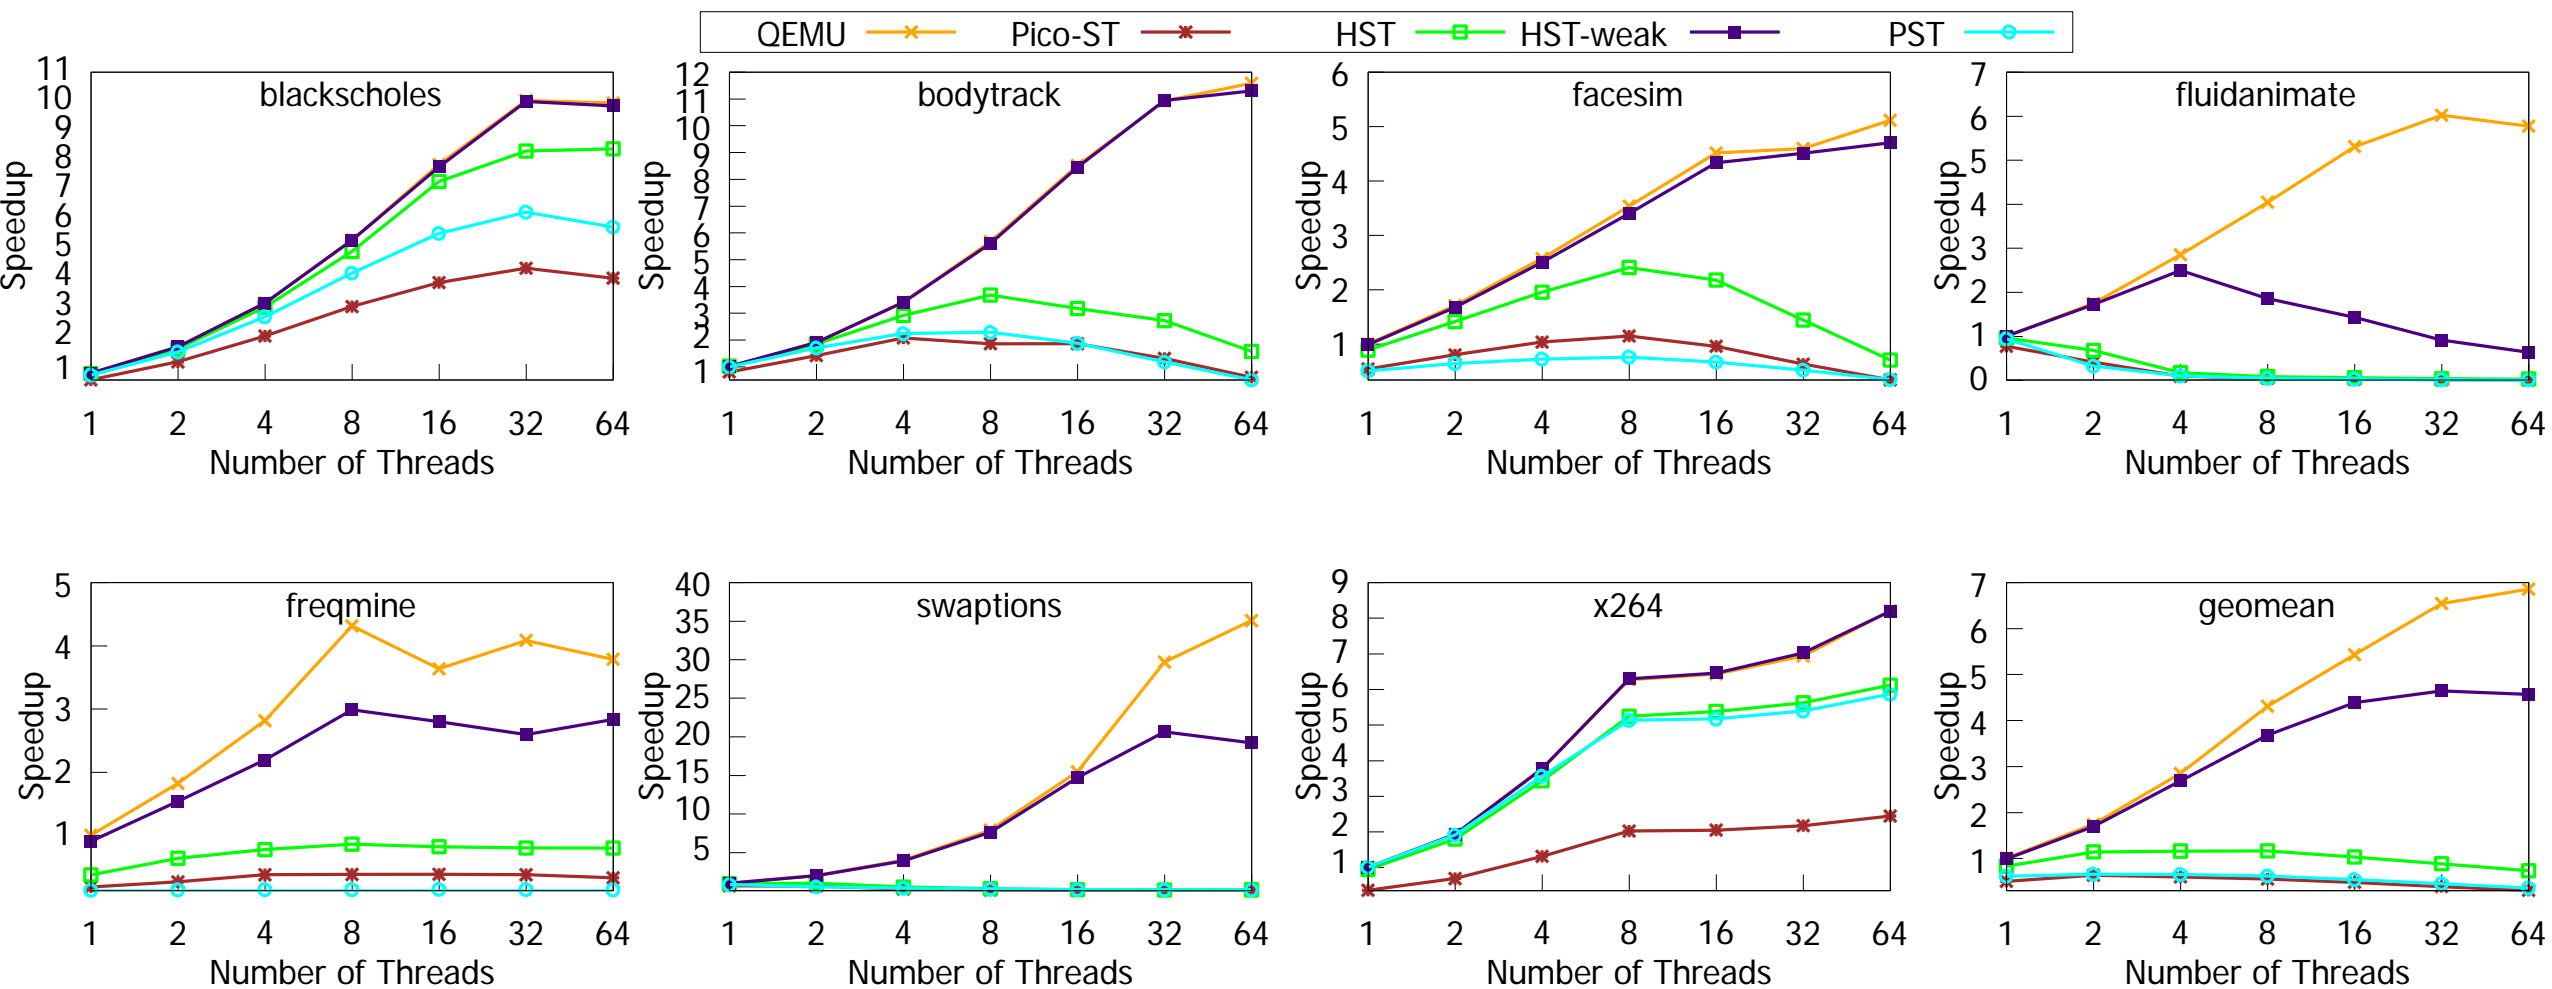


# Scalability – Portable Solutions

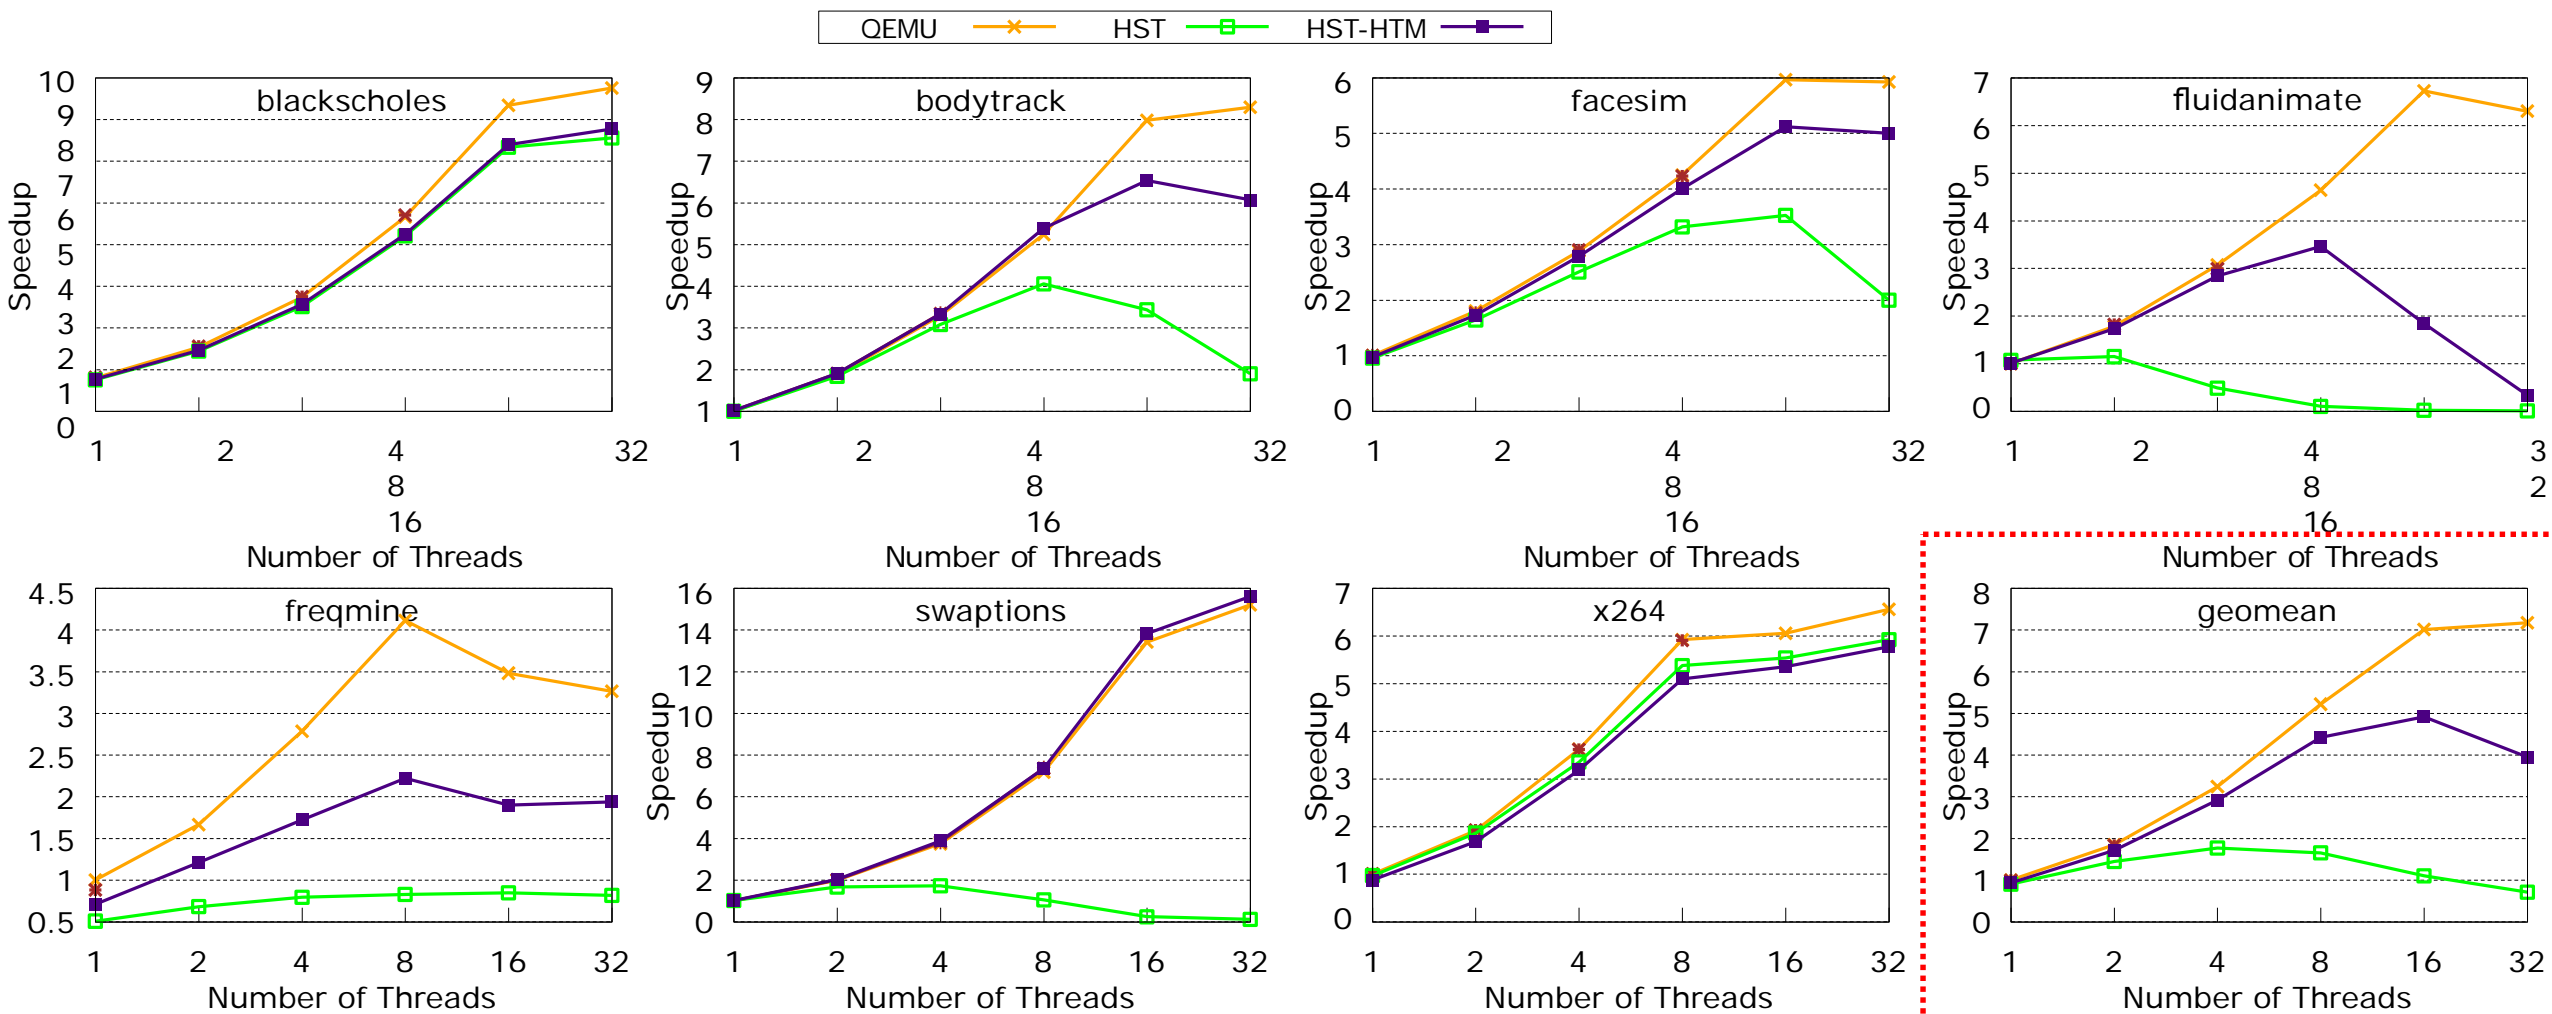




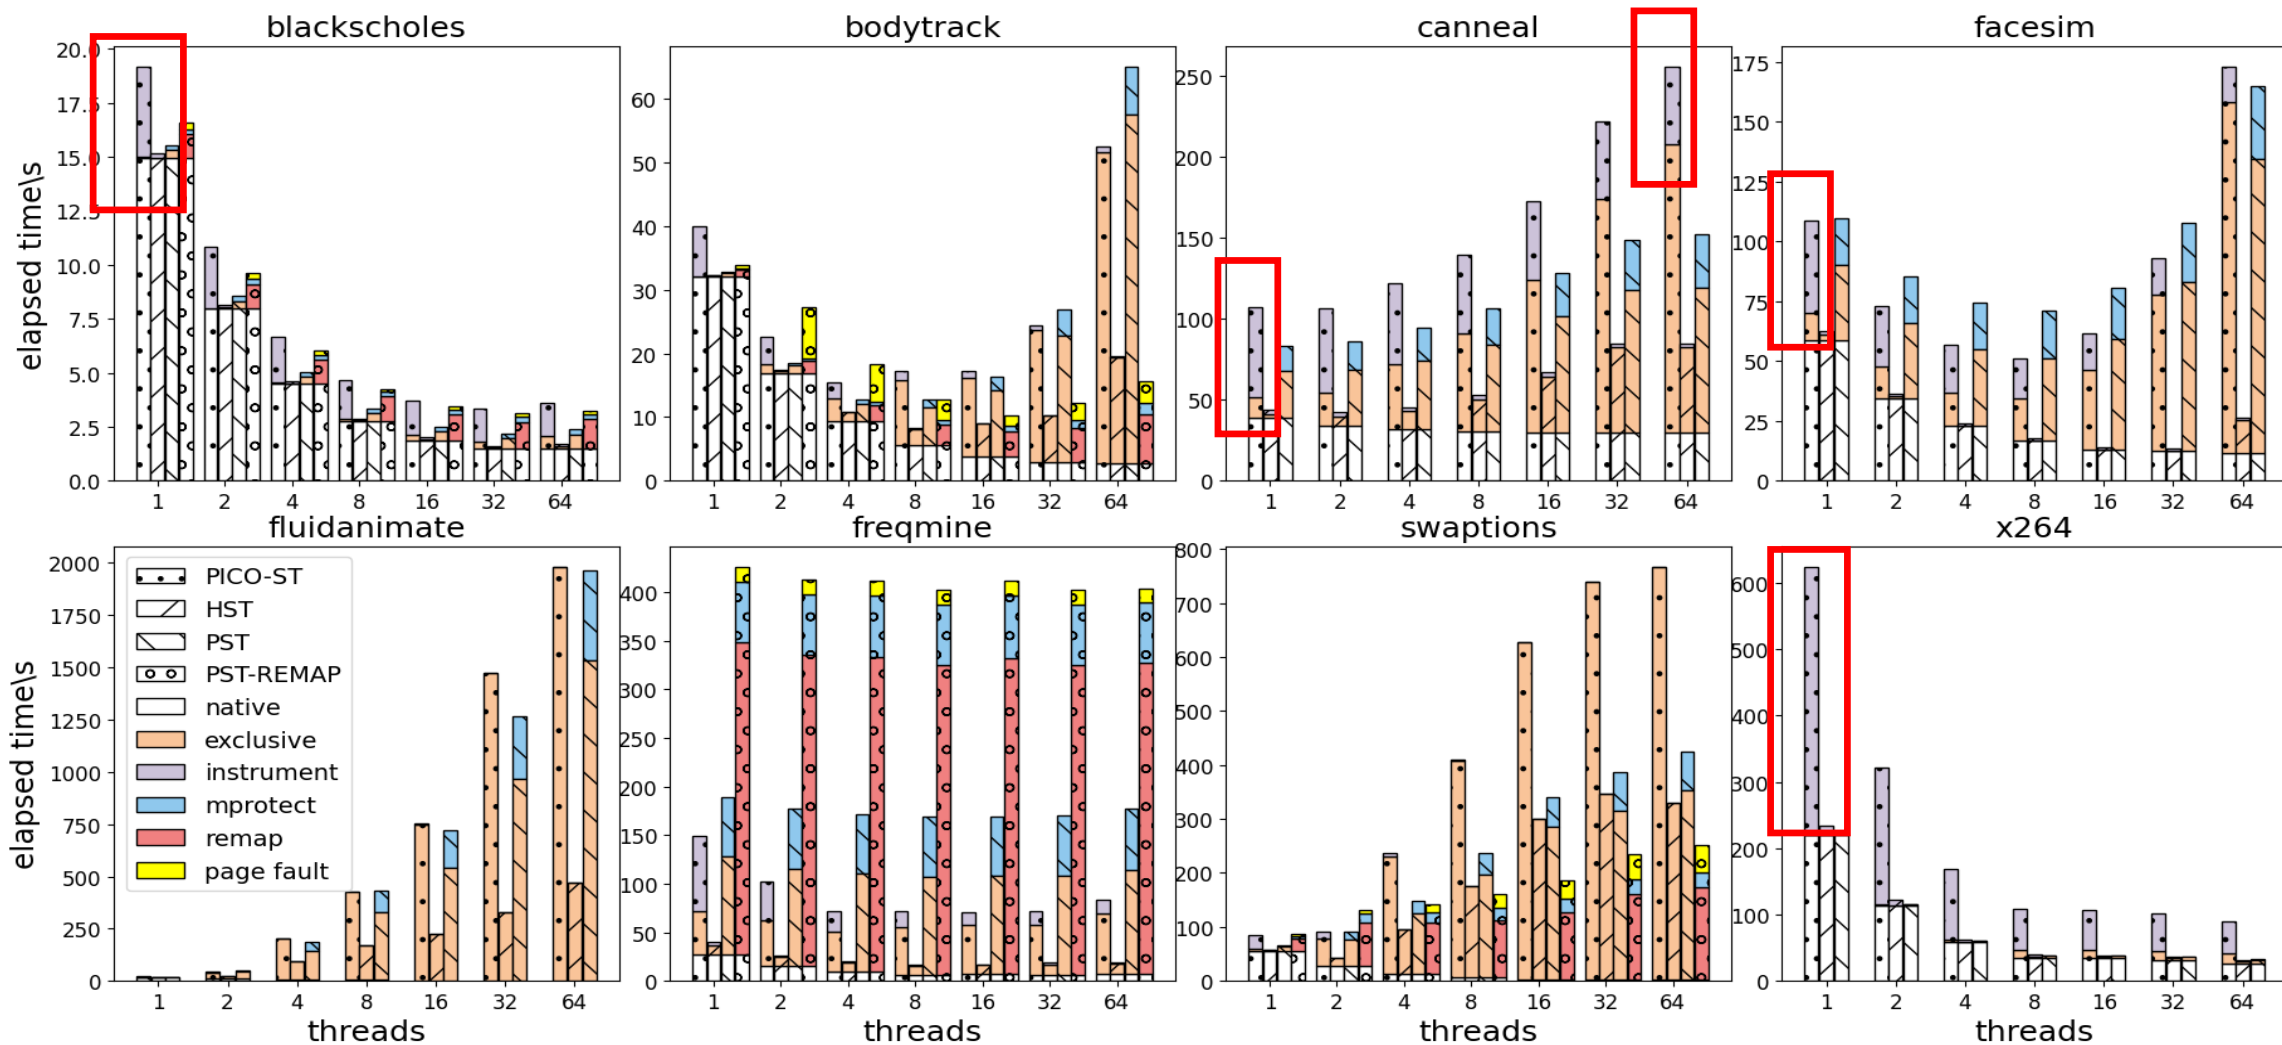
# Scalability – Portable Solutions



# Scalability – HTM based Solutions

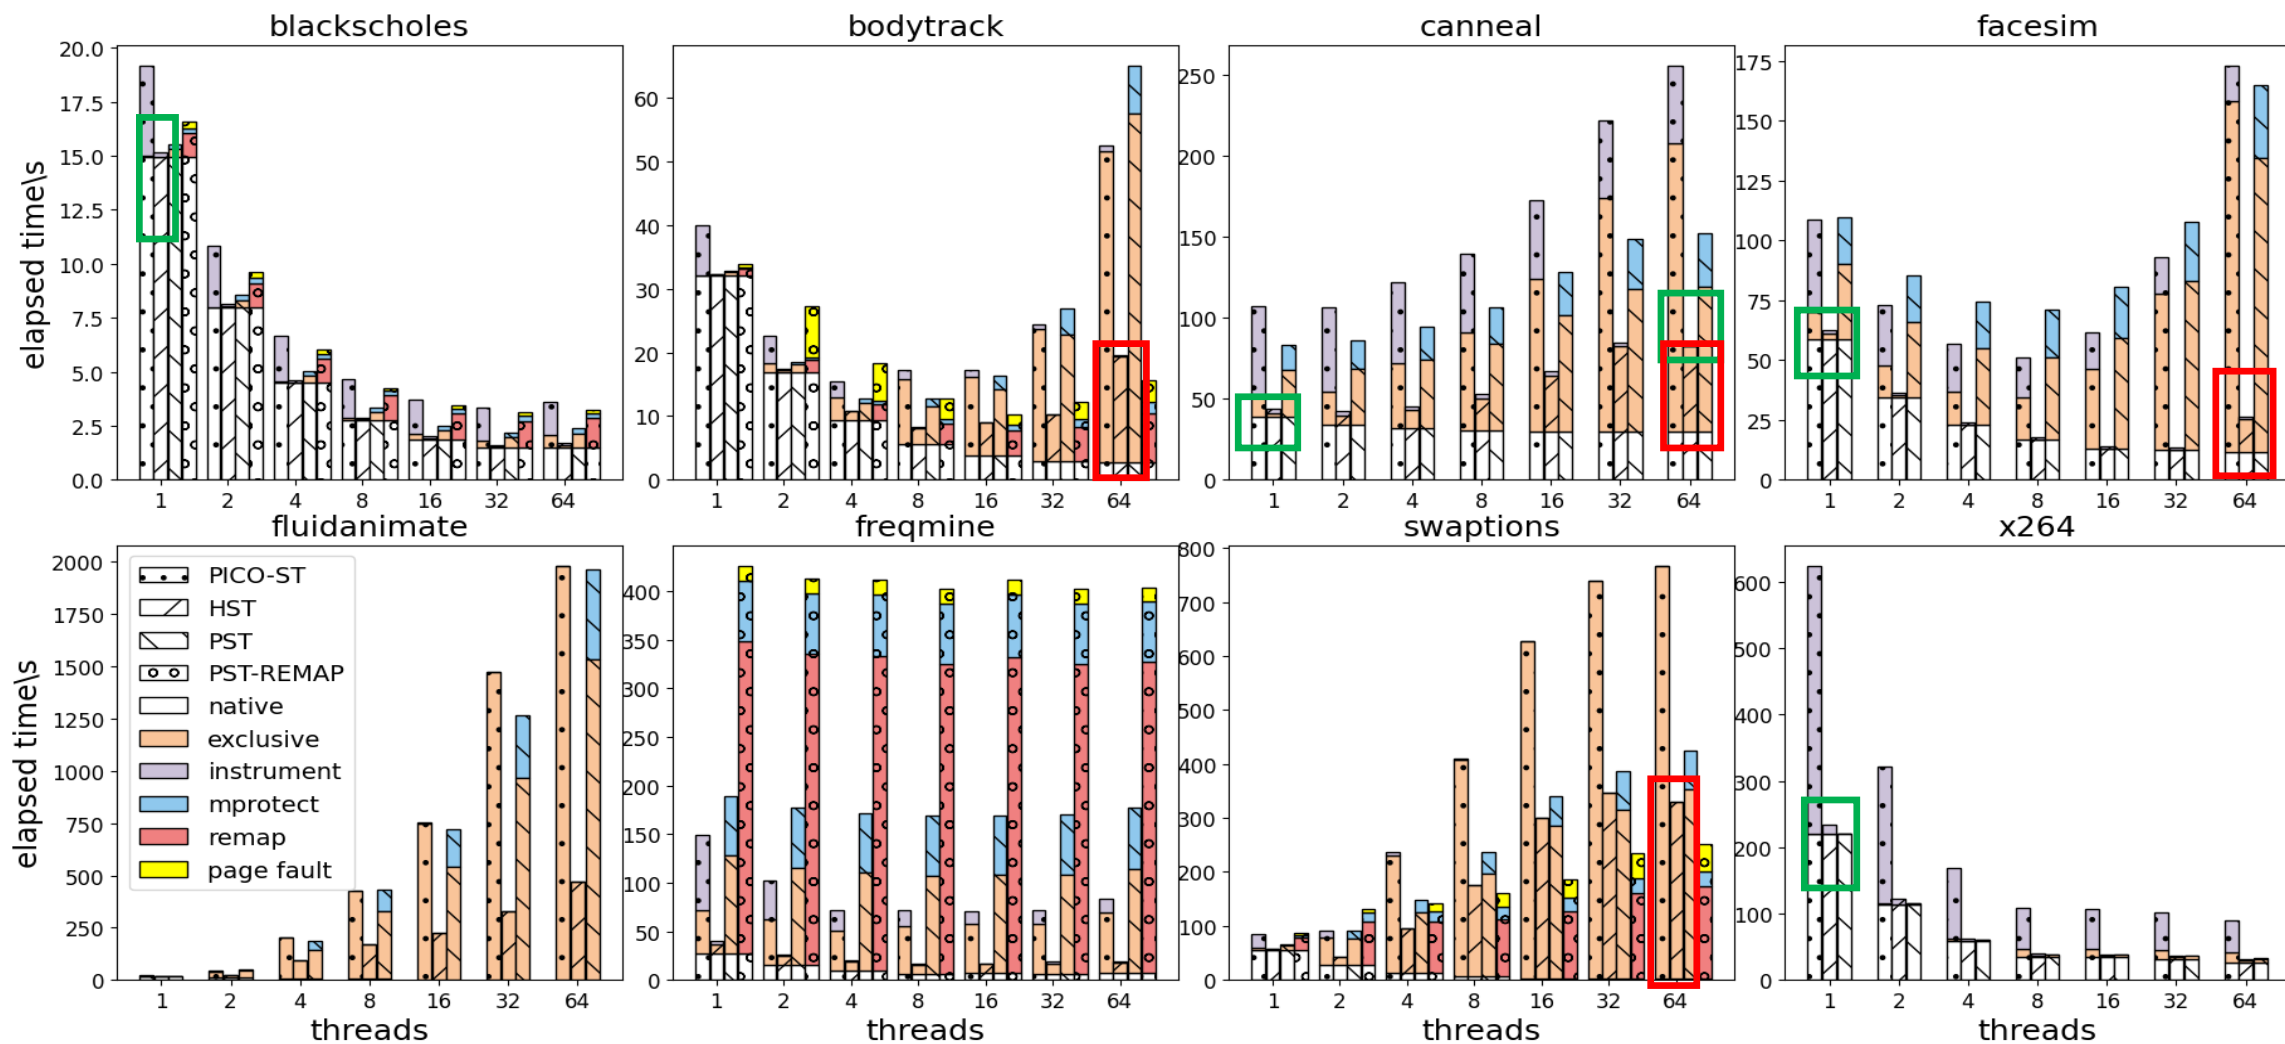


# Overhead Analysis



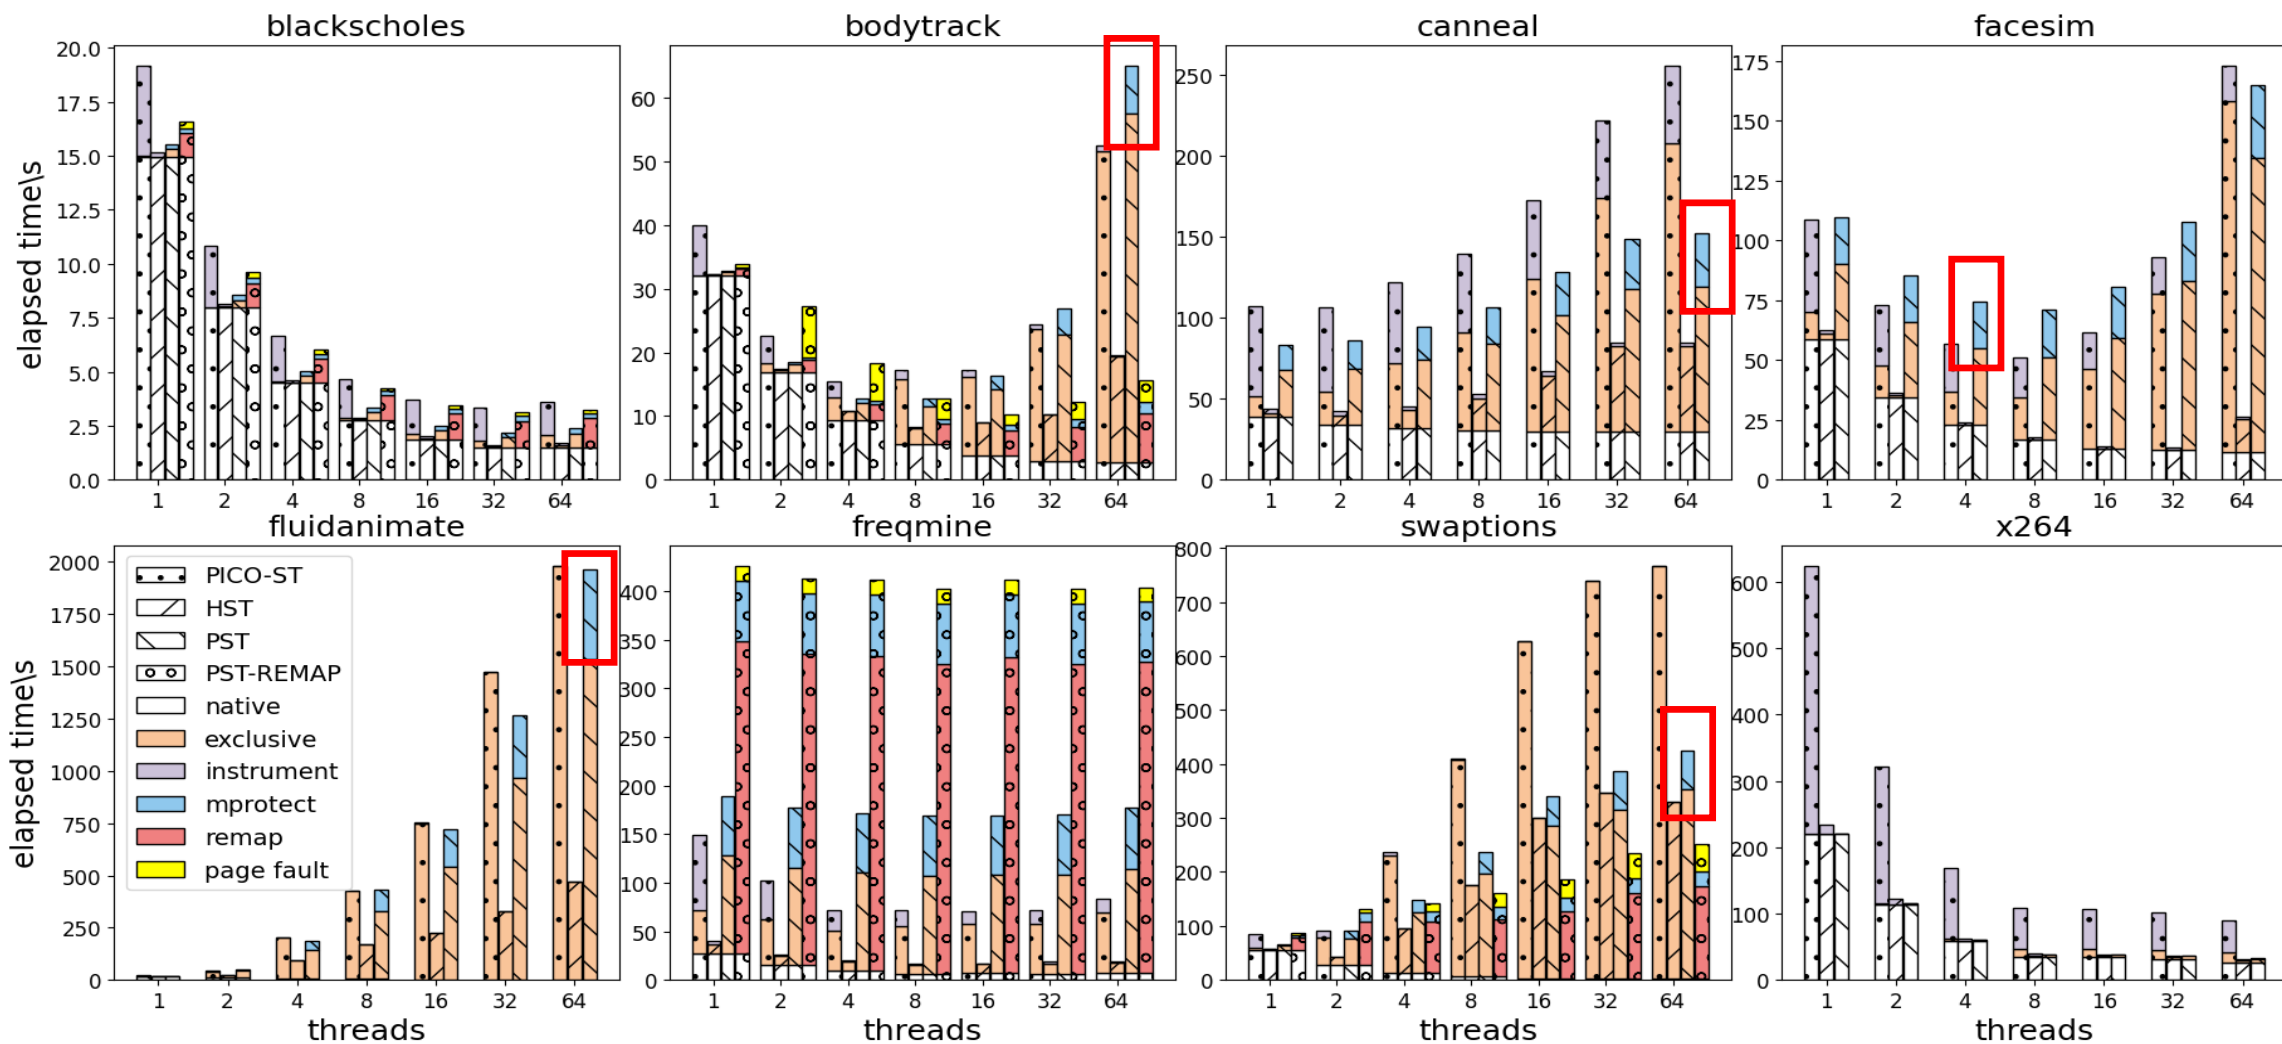
Pico-ST bottleneck: store instrumentation 20%~45%

# Overhead Analysis



**HST: store instrumentation 5%; major overhead: SC synchronization**

# Overhead Analysis



**PST: mprotect system call becomes the new overhead**

# Summary

- Trade-off between atomicity, speed, and portability

Approaches	Speed	Atomicity	Portability
HST	fast	strong	portable
HST-weak	fast	weak	portable
HST-HTM	fast	strong	HTM
PST	slow	strong	portable
PST-remap	varies	strong	portable
Pico-ST	Baseline(slow)	strong	portable
QEMU	fast	incorrect	portable
HTM	fast	incorrect	HTM

# Discussion

- Avoiding synchronization in Store-Conditional
  - Double Compare Single Swap [Timothy, 2002]
  - Intel Memory Protection Key [Park, 2019]
- Changing the translation scheme
  - Rule-based Code Translation [Jiang, 2020]
  - Adding new Intermediate Representation semantics
- Our code is available at: <https://github.com/NKU-EmbeddedSystem/ABA-LLSC>

Thank you!



# Backup Slides

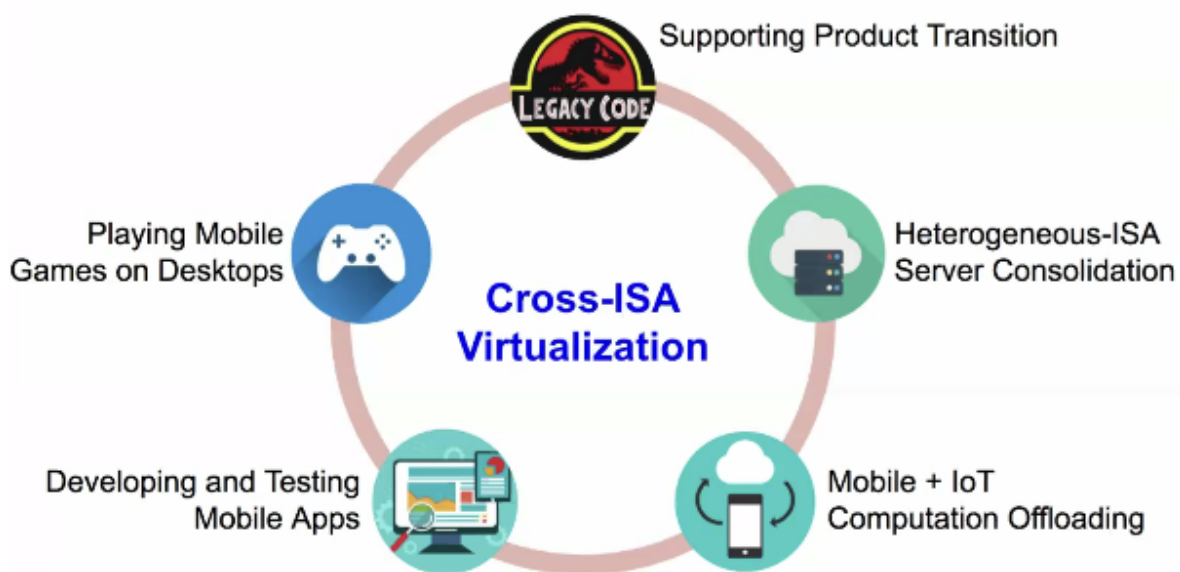
# Outline

- Motivation
  - Atomic instruction in cross-ISA emulation is important
    - Cross-ISA emulation is important: key enabling tech
    - Atomic instruction translation is important: wrong translate lead to wrong results
      - RISC-->CISC r-->w -->ABA
- Design
  - Goal: keep atomicity of LL/SC
  - Prev work:
    - a. workload on threads
      - store overhead 20%~40% (really?) -- could be fast! -- with lock-free hash table
    - b. workload on MMU
      - Seems to be promising (really?) -- could be slow!! -- large syscall overhead
    - c. workload on HTM
      - Fast & correct (really?) -- could be incorrect!! -- JIT & HTM be careful!!
  - HST, HST-weak
  - PST, PST-remap, PST-MPK
  - HST-HTM
- Evaluation
- Discussion
  - PST-MPK

# Cross-ISA Dynamic Binary Translation

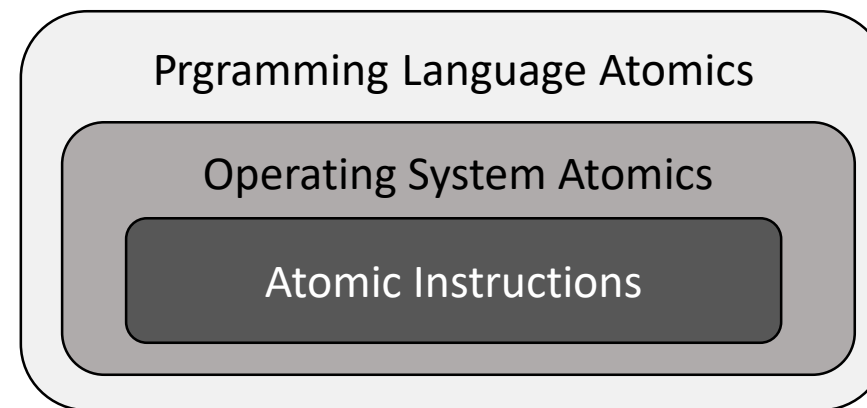
## Cross-ISA Emulation

*“A Key **Enabling** Technology”*



## Atomic Instruction


*Fundamental to **Correctness***



***Atomics are not correctly translated?***

# A. HST: Hash-table based STOre

- Challenge: store instrumentation overhead



Translated Store

```


1  find_EMR(mem)
2  lock
3  update_EMR
4  unlock
5  store  reg, mem
  
```

20%~45%  
instrumentation  
overhead

*Lesson learned: Have to reduce instrumentation code!*


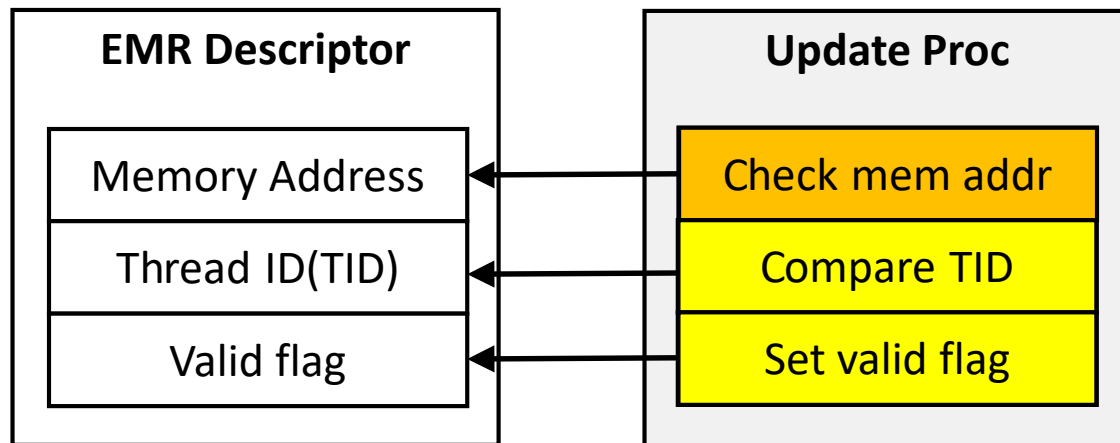
- Scalability: Avoiding lock
- Overhead: Simplifying EMR update proc

>100%  
Sync  
overhead



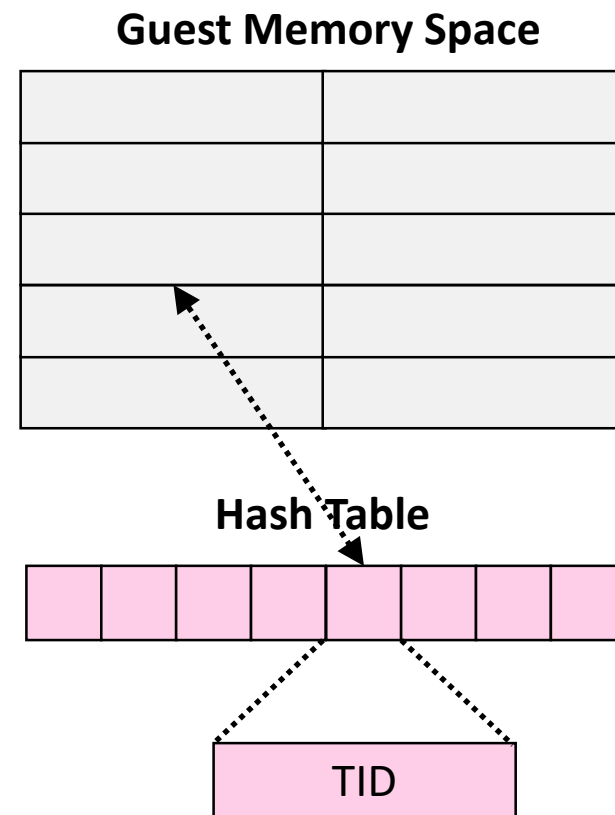
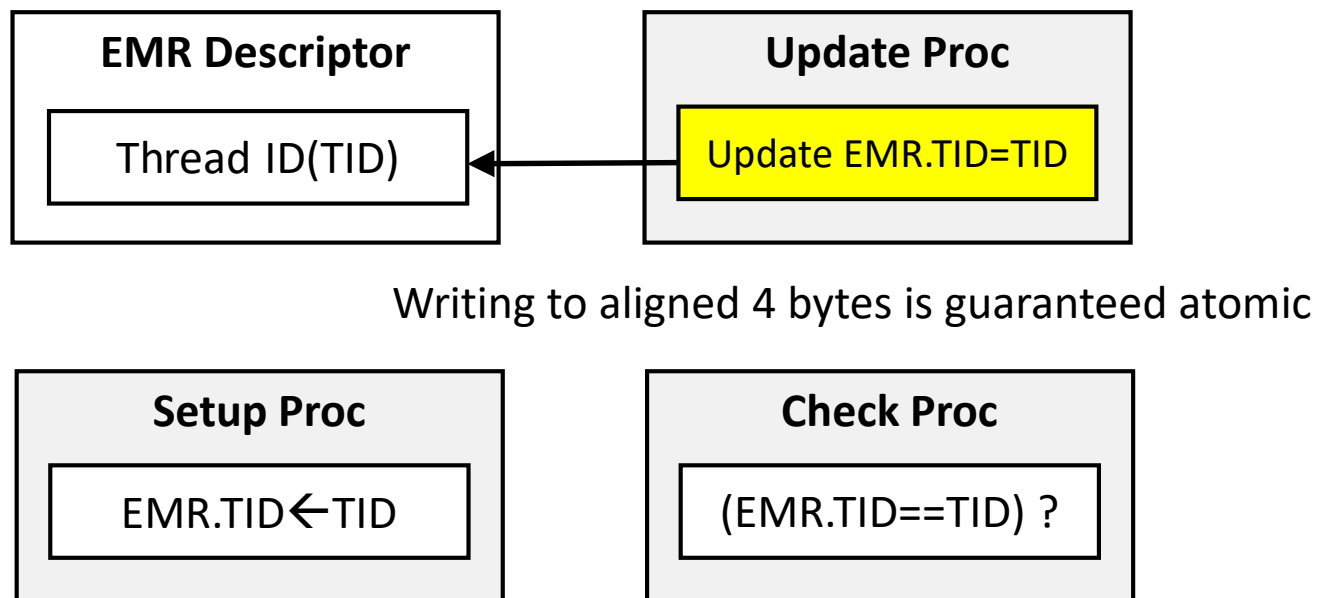
# A. HST: Hash-table based STOre

- Challenge: store instrumentation overhead
  - store instructions accounts for up to 20% instructions, highly concurrent
  - **Be Fast:** lightweight checking procedure
  - **Be Scalable:** no locking



# A. HST: Hash-table based SStore

- Challenge: store instrumentation overhead
  - How to make `update_EMR` atomic?
    - Merging memory checking
    - Merging TID and Valid Flag



# A. HST: Hash-table based SStore

- Does it provide atomicity
- Hash confliction

