

DMP

Deterministic Shared Memory Multiprocessing

sailipa



University of
Washington

Joe Devietti, Brandon Lucia, Luis Ceze, Mark Oskin

A multithreaded voting machine

2

thread 0

```
while (more_votes) {  
  load t <- votes  
  t++  
  store t -> votes  
}
```




votes == 2

thread 1

```
while (more_votes) {  
  load t <- votes  
  t++  
  store t -> votes  
}
```



votes == 2



votes == 1

A multithreaded voting machine

3

data race

thread 0

thread 1

```
while (more_votes) {  
    store t <- votes  
}  
while (more_votes) {  
    t++  
    store t > votes  
}
```

we're not trying to make these bugs go away
we're trying to make them **come back!**

**locking discipline
violation**



Why is parallel programming hard?

4

sequential bugs

concurrency bugs

We want **parallel** programs
to behave like **sequential** program

nondeterministic

memory access
interleavings

- hard to debug
- hard to test
- hard to replicate
- hard to leverage crash information

Determinism Can Help

5

Development

- no more heisenbugs!
- time-travel debugging
- **test inputs, not interleavings**

Deployment


- reproduce bugs from field
- easy to synchronize replicas
- **software behaves as tested**

**Can we remove nondeterminism
without removing performance?**

DMP from 10,000'

6

- We only care about communicating instructions
- Deterministic serialization → same communication
 - ▣ ...but I promised you performance!
- Recover parallelism from non-communicating insns



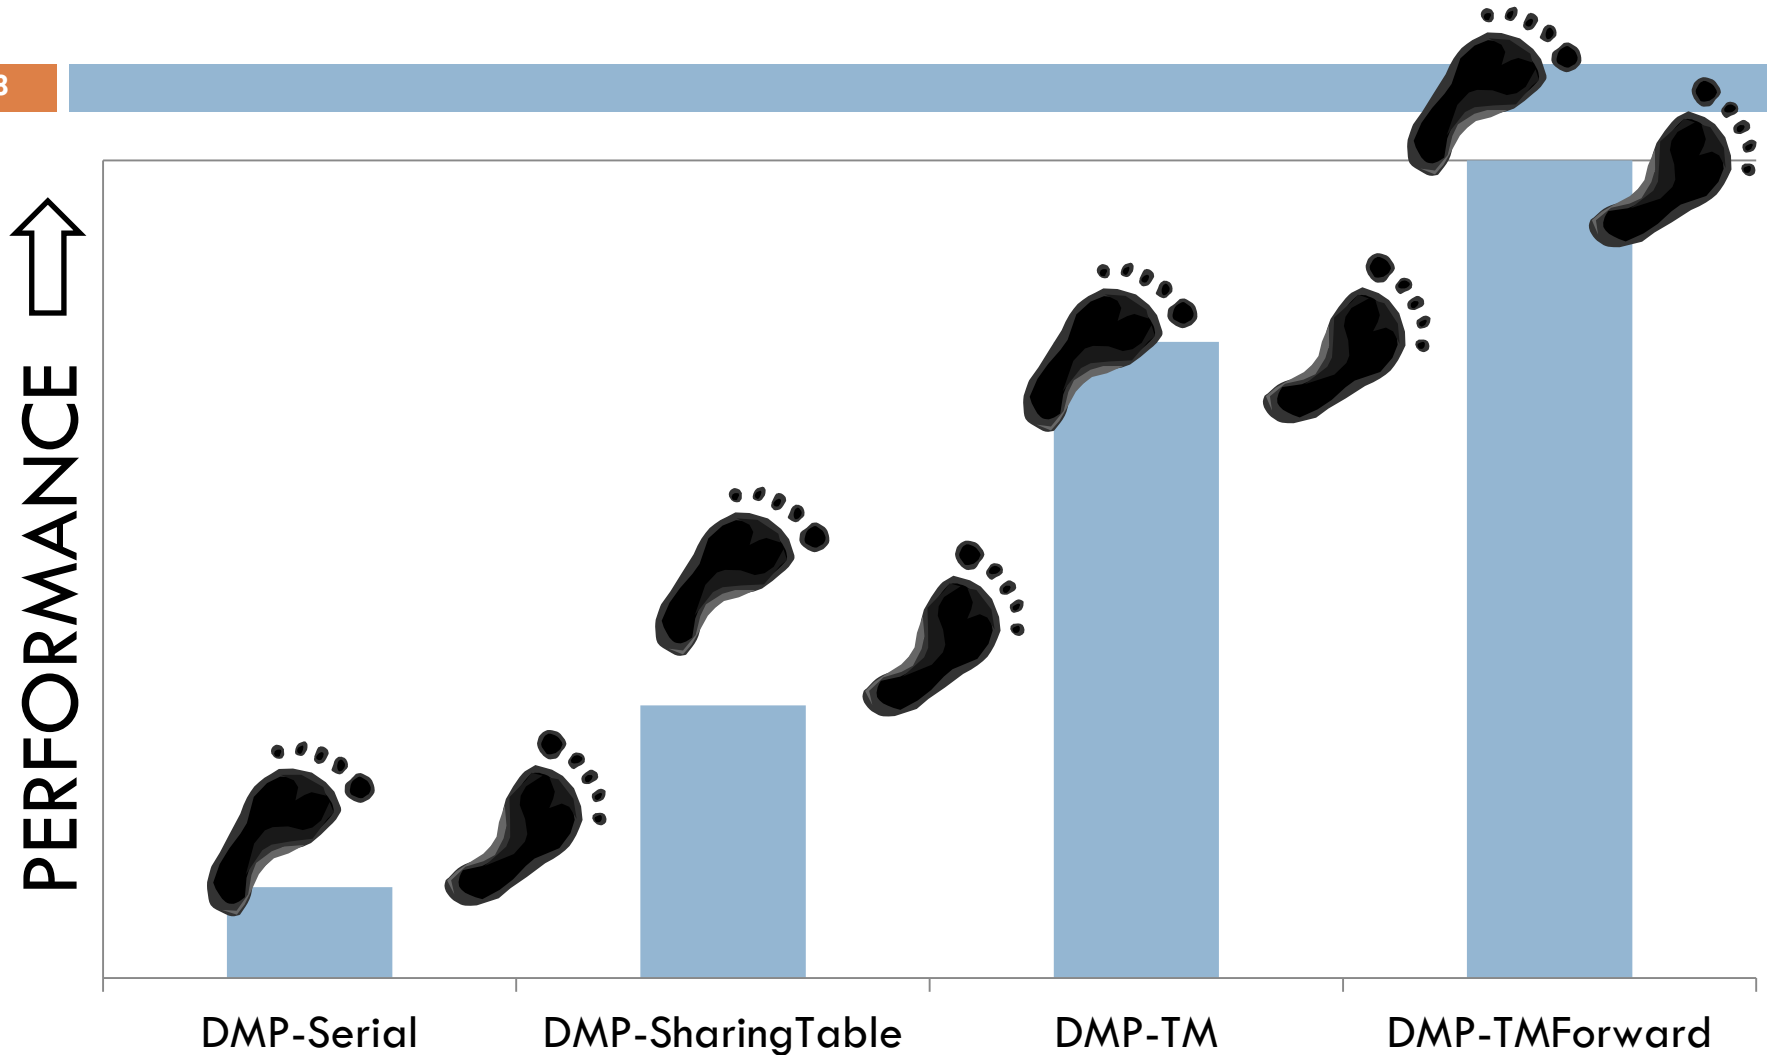
Related Work

7

Helps with...	Record + Replay
...debugging?	<input checked="" type="radio"/>
...testing?	<input type="radio"/>
...replicas?	<input type="radio"/>
...deployment?	<input type="radio"/>
Needs hw?	usually
examples:	FDR, ReRun, Capo

Talk Outline

8




COMPLEXITY →

DMP: Deterministic Shared Memory Multiprocessing - ASPLOS 2009

DMP-Serial Example


9



Can we do better?

10

- Only need to serialize communicating instructions
- Break each quantum into communication-free **parallel prefix** and communicative **serial suffix**




DMP-SharingTable

11

- Need to know when communication happens, to transition from parallel to serial mode
 - ▣ Leverage existing cache coherence protocol
 - ▣ When a line moves between processors, communication is (potentially) happening!
 - ▣ The **Sharing Table** tracks information about ownership
- State of Sharing Table must evolve deterministically
 - ▣ Only allow updates during serial suffix

DMP-SharingTable Example

12



DMP-TM:


Recovering Parallelism with Speculation

13

- DMP-SharingTable conservatively assumes that all changes in ownership are communication
 - ▣ ...but most changes in ownership are **not** communication
- Use TM support to speculate that a quantum is not involved in communication
 - ▣ If communication happens, rollback + re-execute
- Each quantum is an implicit transaction
 - ▣ **Commit quanta in-order** (need DT to commit)

DMP-TM Example

14




DMP-TM-Forward: Speculative Value Forwarding

15

- DMP-TM eliminates WAW and WAR dependencies
 - ▣ but cannot speculate past true (RAW) dependences
- Idea: speculatively forward values to “future” quanta
 - ▣ coherence protocol + ordered transactions make it easy to decide *when* and *where* to forward
 - ▣ rollback if a quantum’s speculatively read data is updated before the quantum commits

DMP-TM-Forward Example


16



Better Quantum Building

17

- Any deterministic policy will work
- We want quanta that are free of communication
 - ▣ no communication → no serialization, no rollbacks




Experimental Methodology

18

- PIN-based simulator
 - ▣ Models serialization, quantum building, address conflicts and transaction rollbacks
 - ▣ Assumes constant IPC with free commits
- SPLASH2 and PARSEC benchmark suites

Results

19



- 1,000-instruction quanta
- Synchronization+burst quantum builder

Also in the paper...

20

- Software-only Sharing Table implementation
- Support for debugging
 - ▣ Adding instrumentation without affecting communication
- Making execution deterministic across machines
- Dealing with nondeterminism from I/O and the OS

Conclusions

21

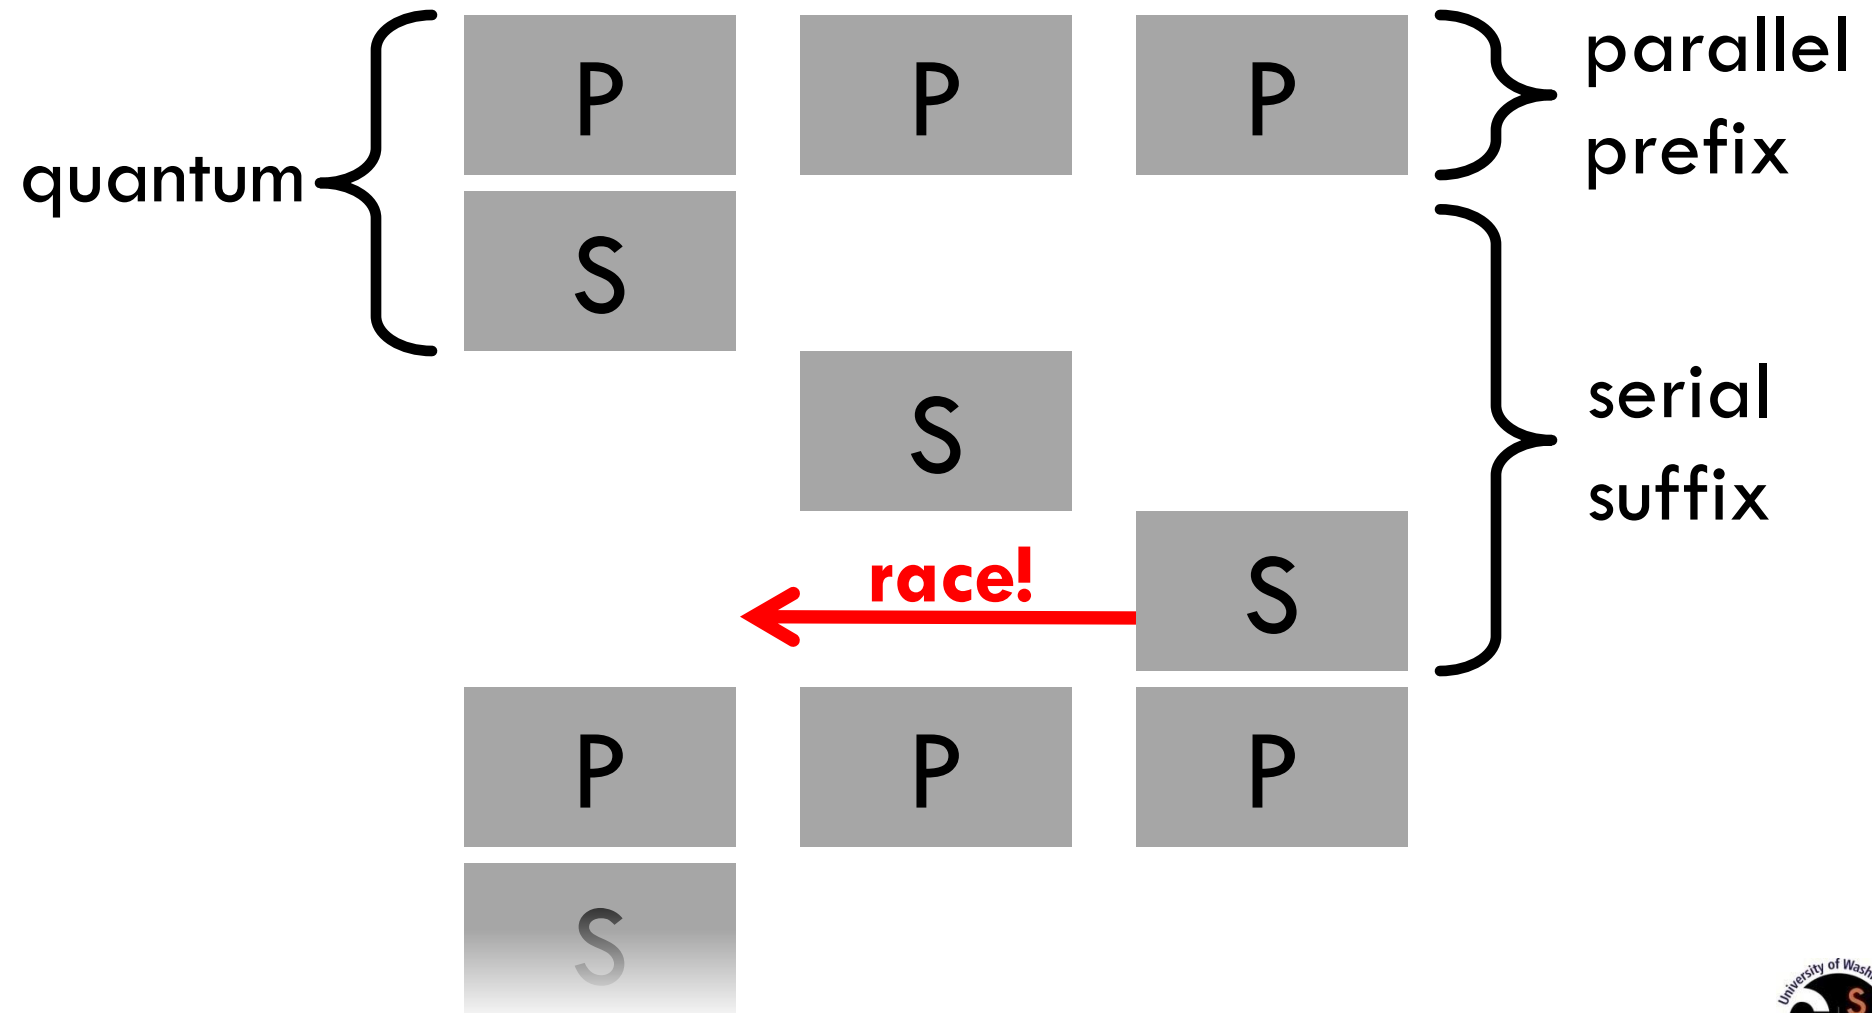
- Determinism makes parallel programming easier
 - ▣ Execution is repeatable
 - ▣ Simplifies debugging, testing, replicating and deployment
- DMP is a new multiprocessor architecture that provides determinism for arbitrary shared memory programs
 - ▣ Leverages existing architectural techniques
 - ▣ Performance very close to nondeterministic execution
- **Determinism is a worthwhile and achievable goal**

Questions?

22

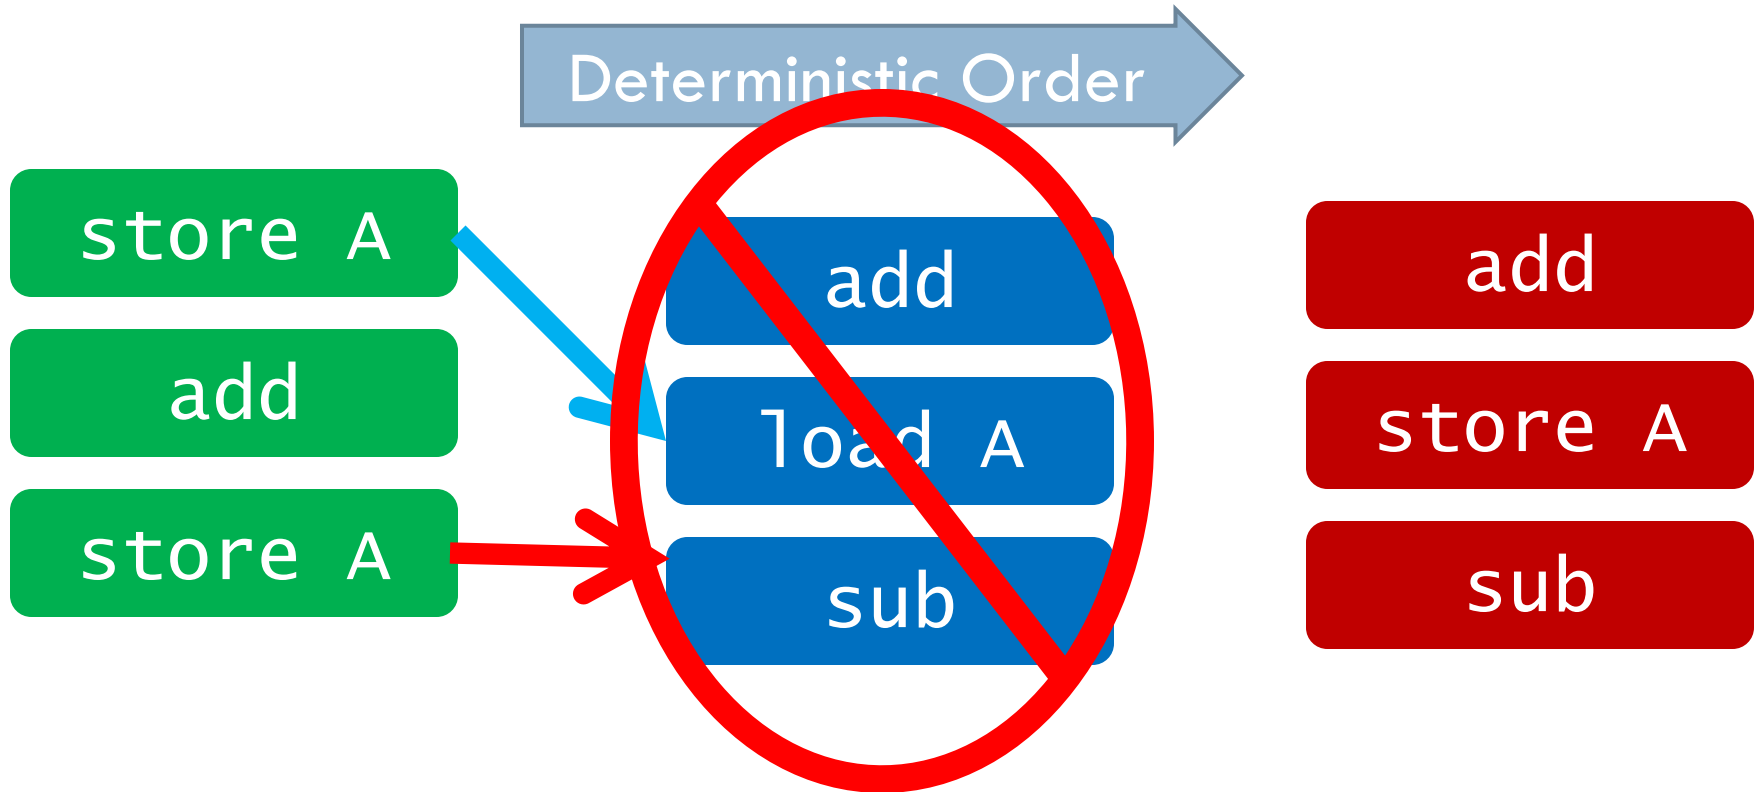
Quantum Rounds

23



TM-Forward Rollback

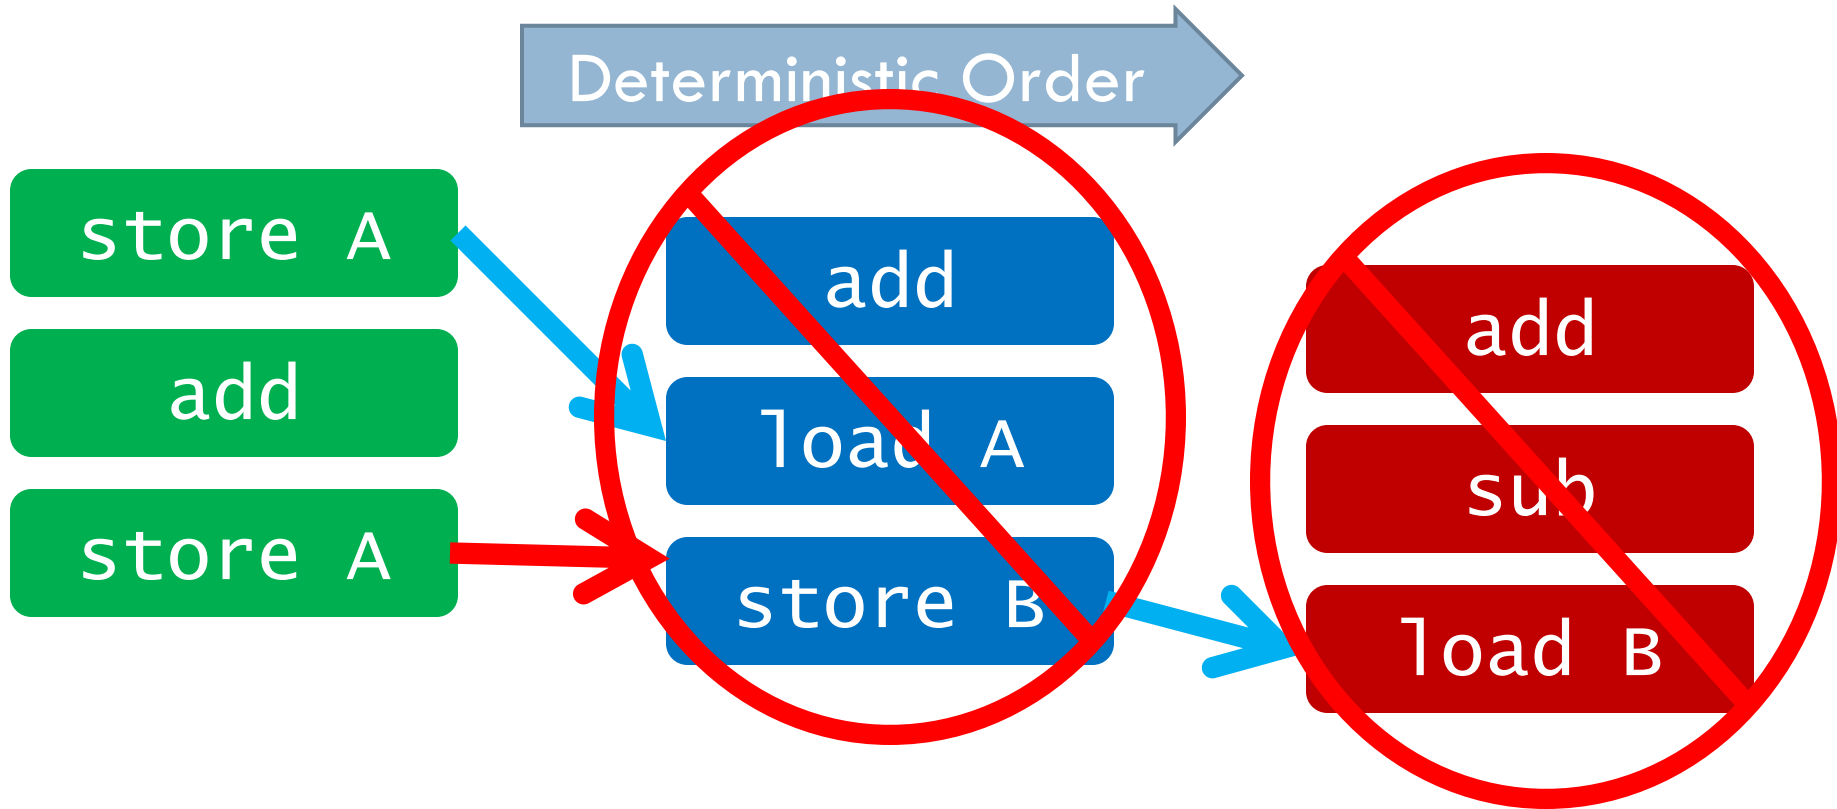
24



Need value from youngest older store


TM-Forward Recursive Rollback

25



Flash test slide

View this slide in slide show mode (shift-F5). You should see a blue box here (a Flash movie). It may take a few seconds to load. It will check Flash movie permissions and interaction with Powerpoint.



**Flash → powerpoint interaction works!
This presentation should work fine :-)**

Learn more about this presentation at

<http://linuxforlovers.wordpress.com/2009/03/15/dmp-asplos-presentation-readme/>