A Systematic Methodology to Develop Resilient Cache Coherence Protocols

Konstantinos Aispos Li-Shiuan Peh

Presented by: Helen Arefayne Milind Furia

What is cache coherence?

- Loading the correct value when the same data is stored in multiple caches

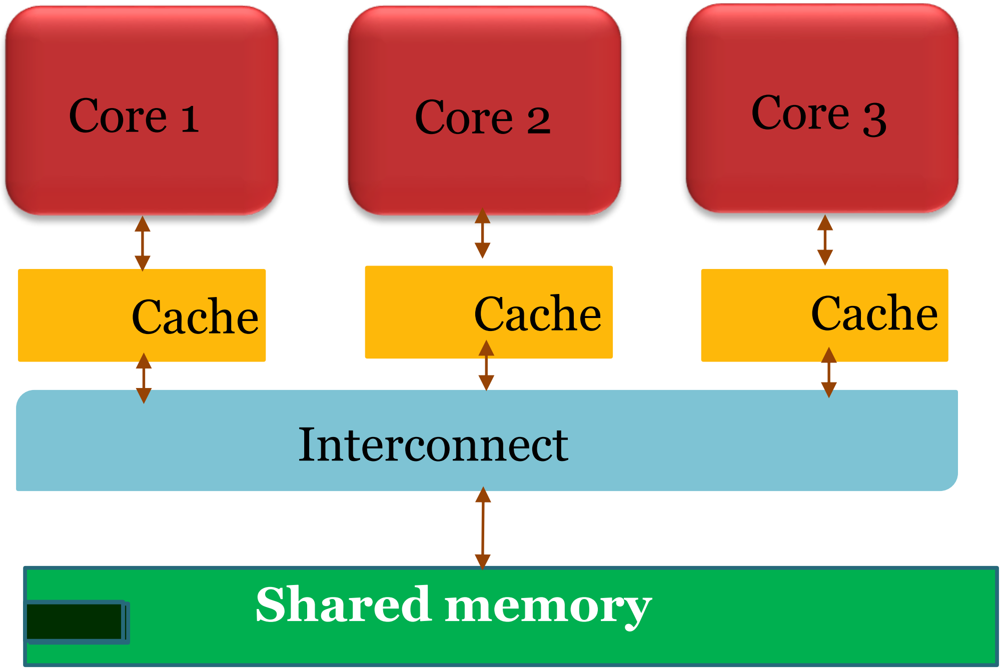

Incoherent state!

Types of cache coherence

- Directory-based
- Broadcast-based
- Snooping

Problem

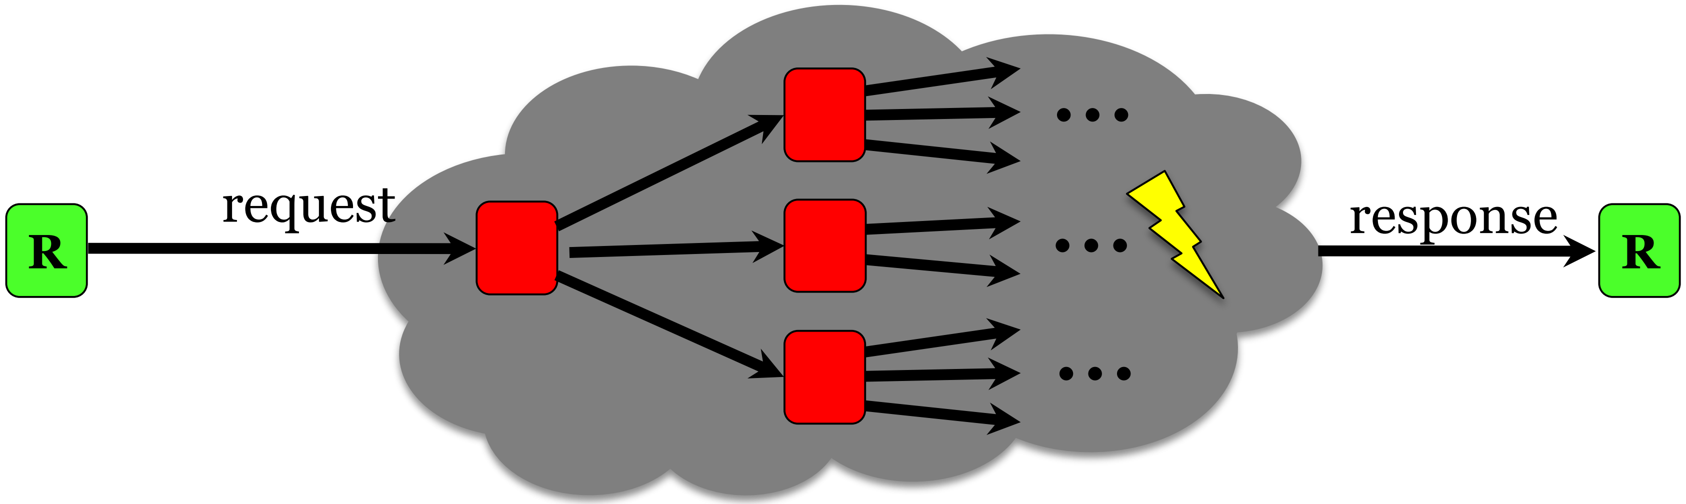

Unreliable interconnect → Suspended transaction

Cause

- Transient faults

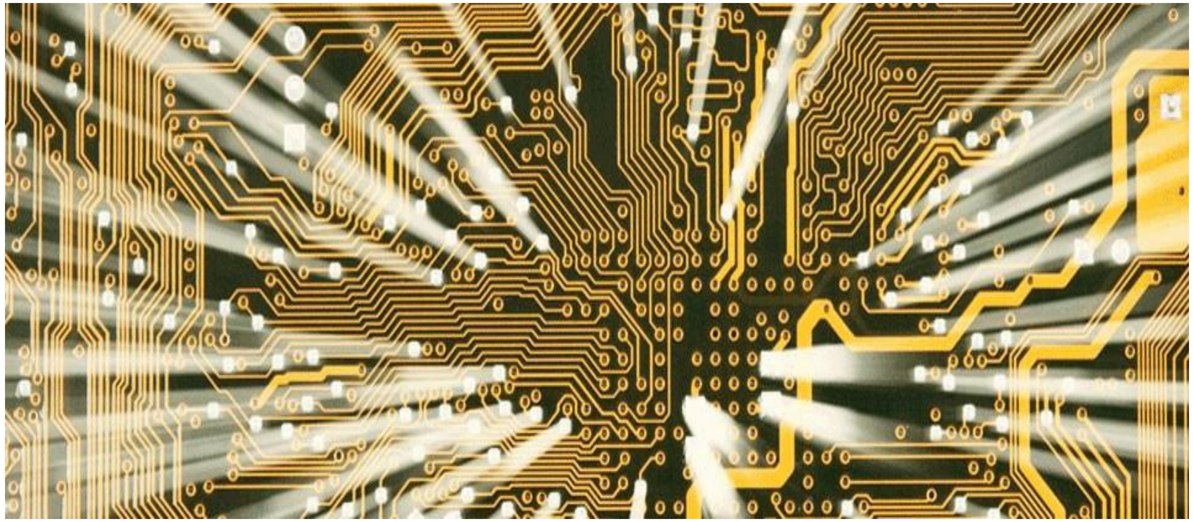

Solution

- Extend coherence protocols for resilience
  - Detect deadlocks
  - Retransmit lost messages

Related work

- Checkpointing [Prvulovic *et al.*, Sorin, *et al.*]
  - Pro-active
- FTDirCMP [Pascual *et al.*]
  - Protocol-specific

Characters of a resilient protocol

- Property 1
  - All initiators of transactions stay in transient state until all state go to stable state
- Property 2
  - Previously transmitted messages can be retransmitted
- Property 3
  - All nodes can tolerate duplicate messages and produce same outcome

Directory based coherence - A

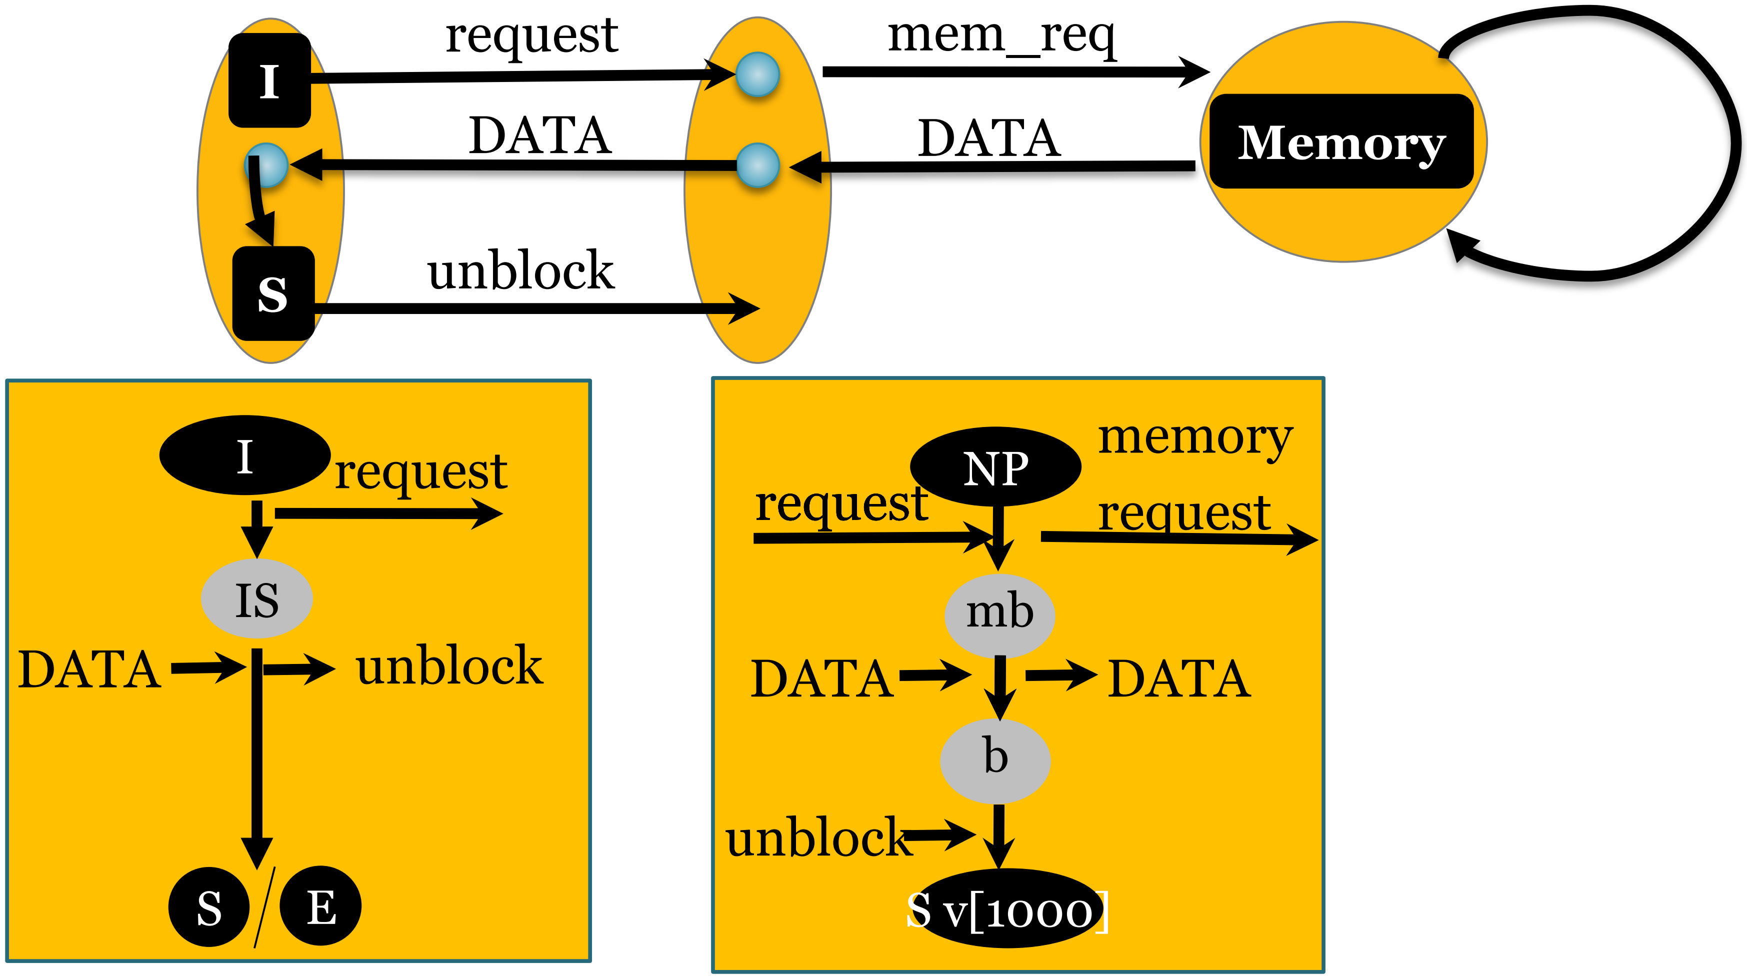

requestor

directory

a) data: NP/S, request: S

Enforcing property 1

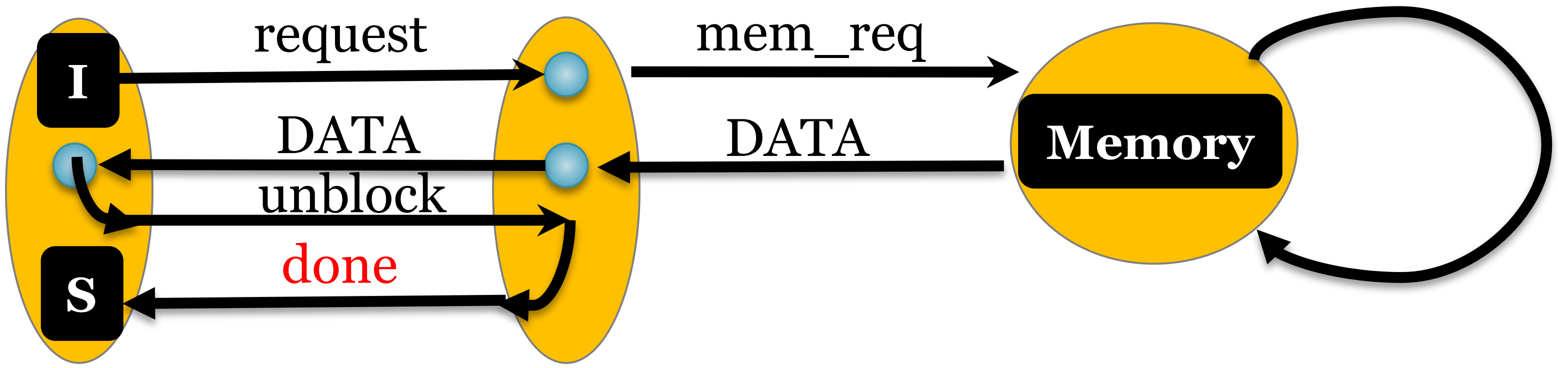

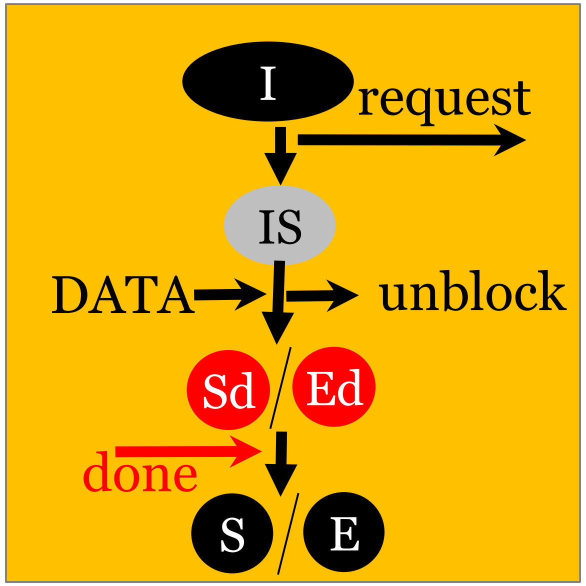

requestor

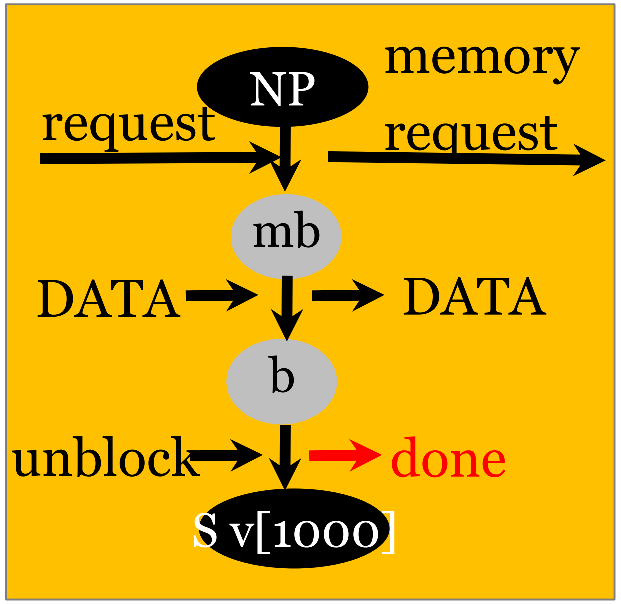

directory

a) data: NP/S, request: S

Directory based coherence - B

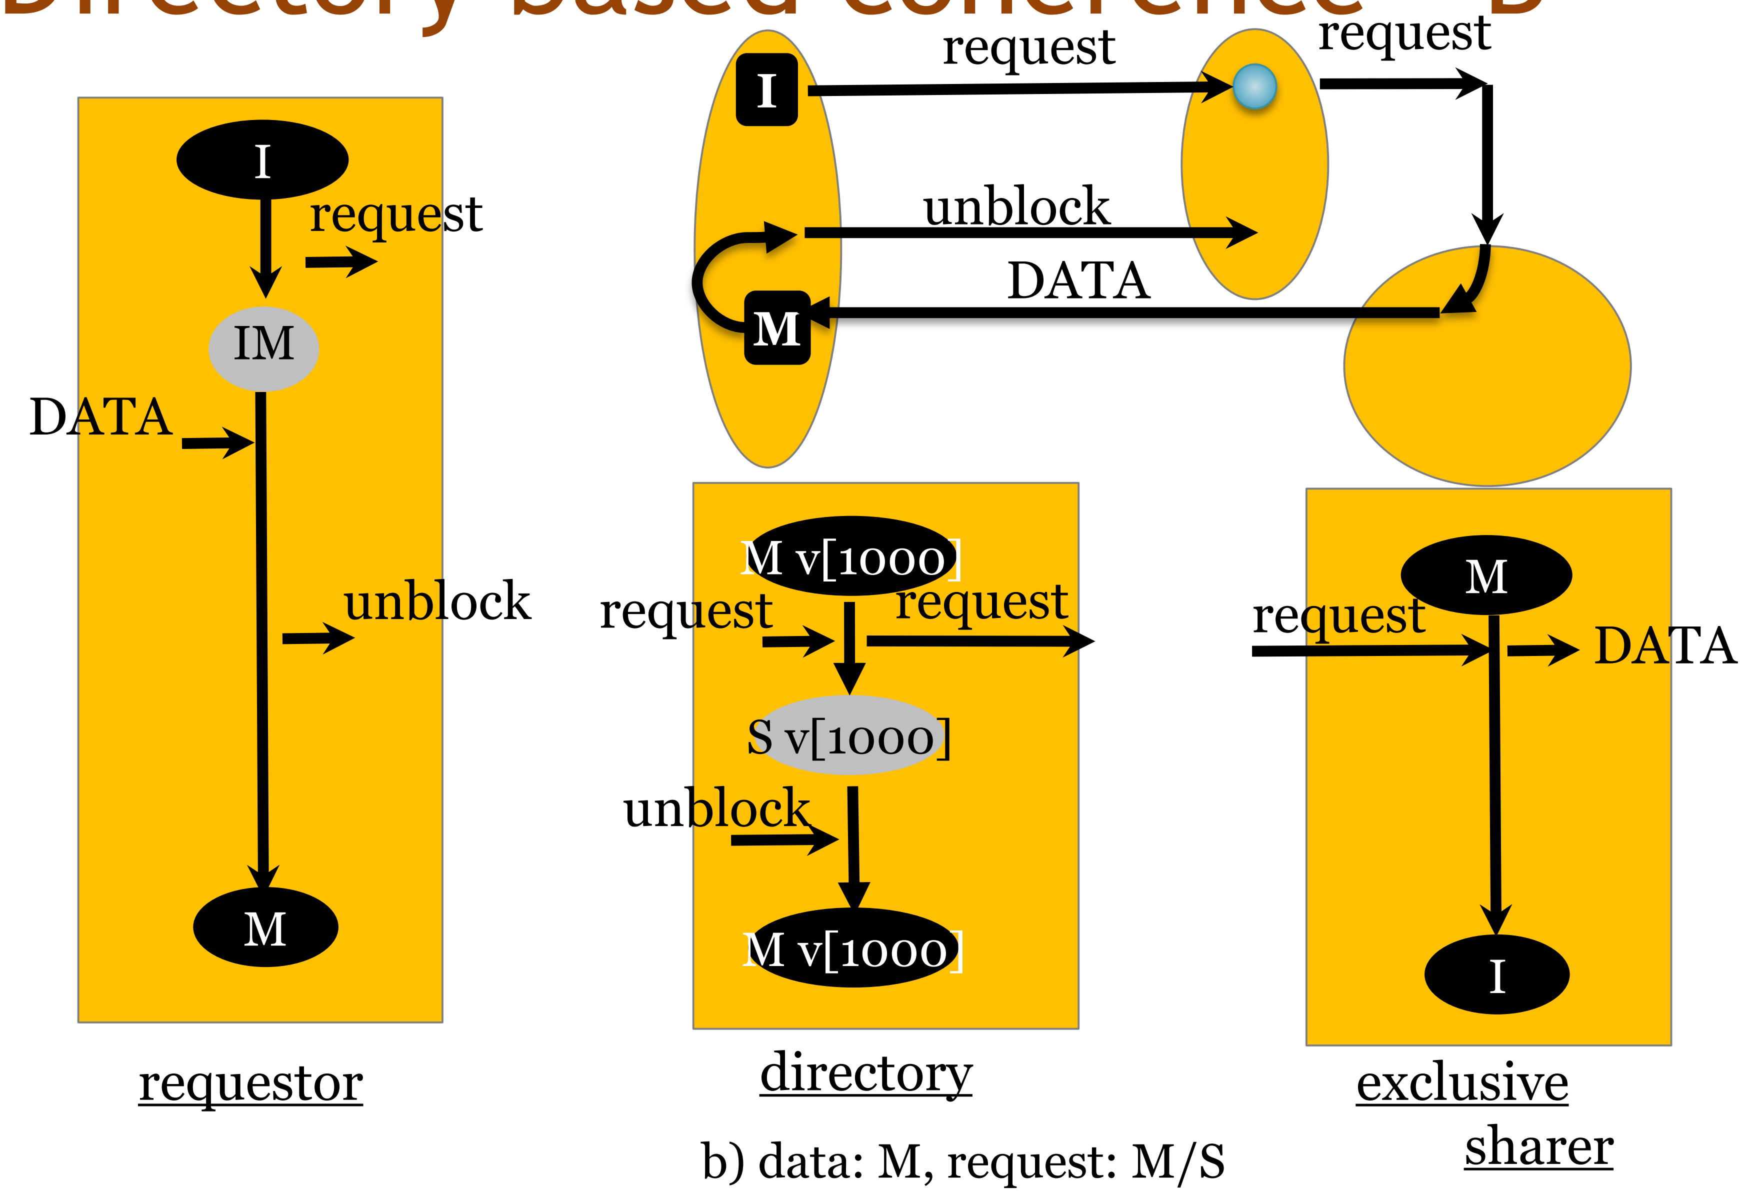

b) data: M, request: M/S

Enforcing property 2

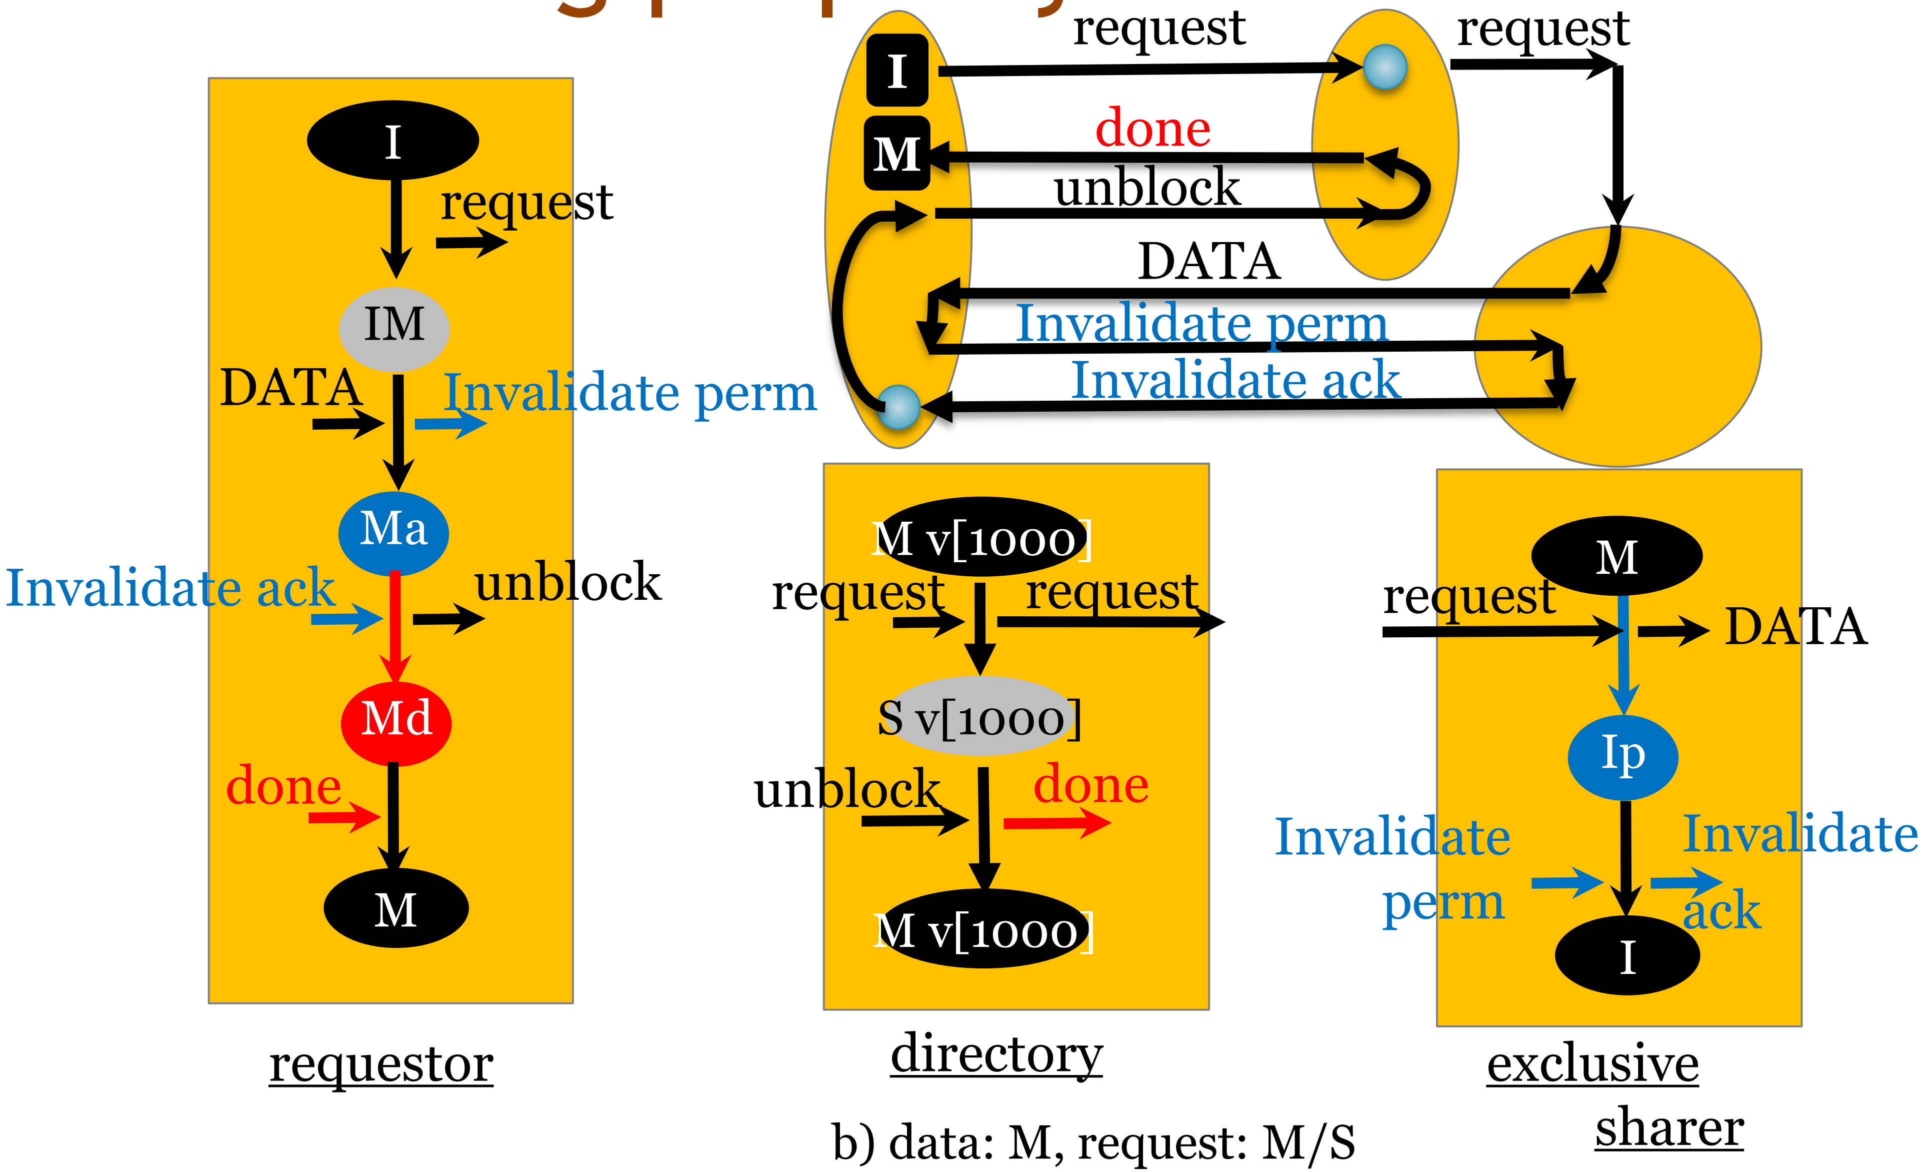

b) data: M, request: M/S

Directory based coherence - C

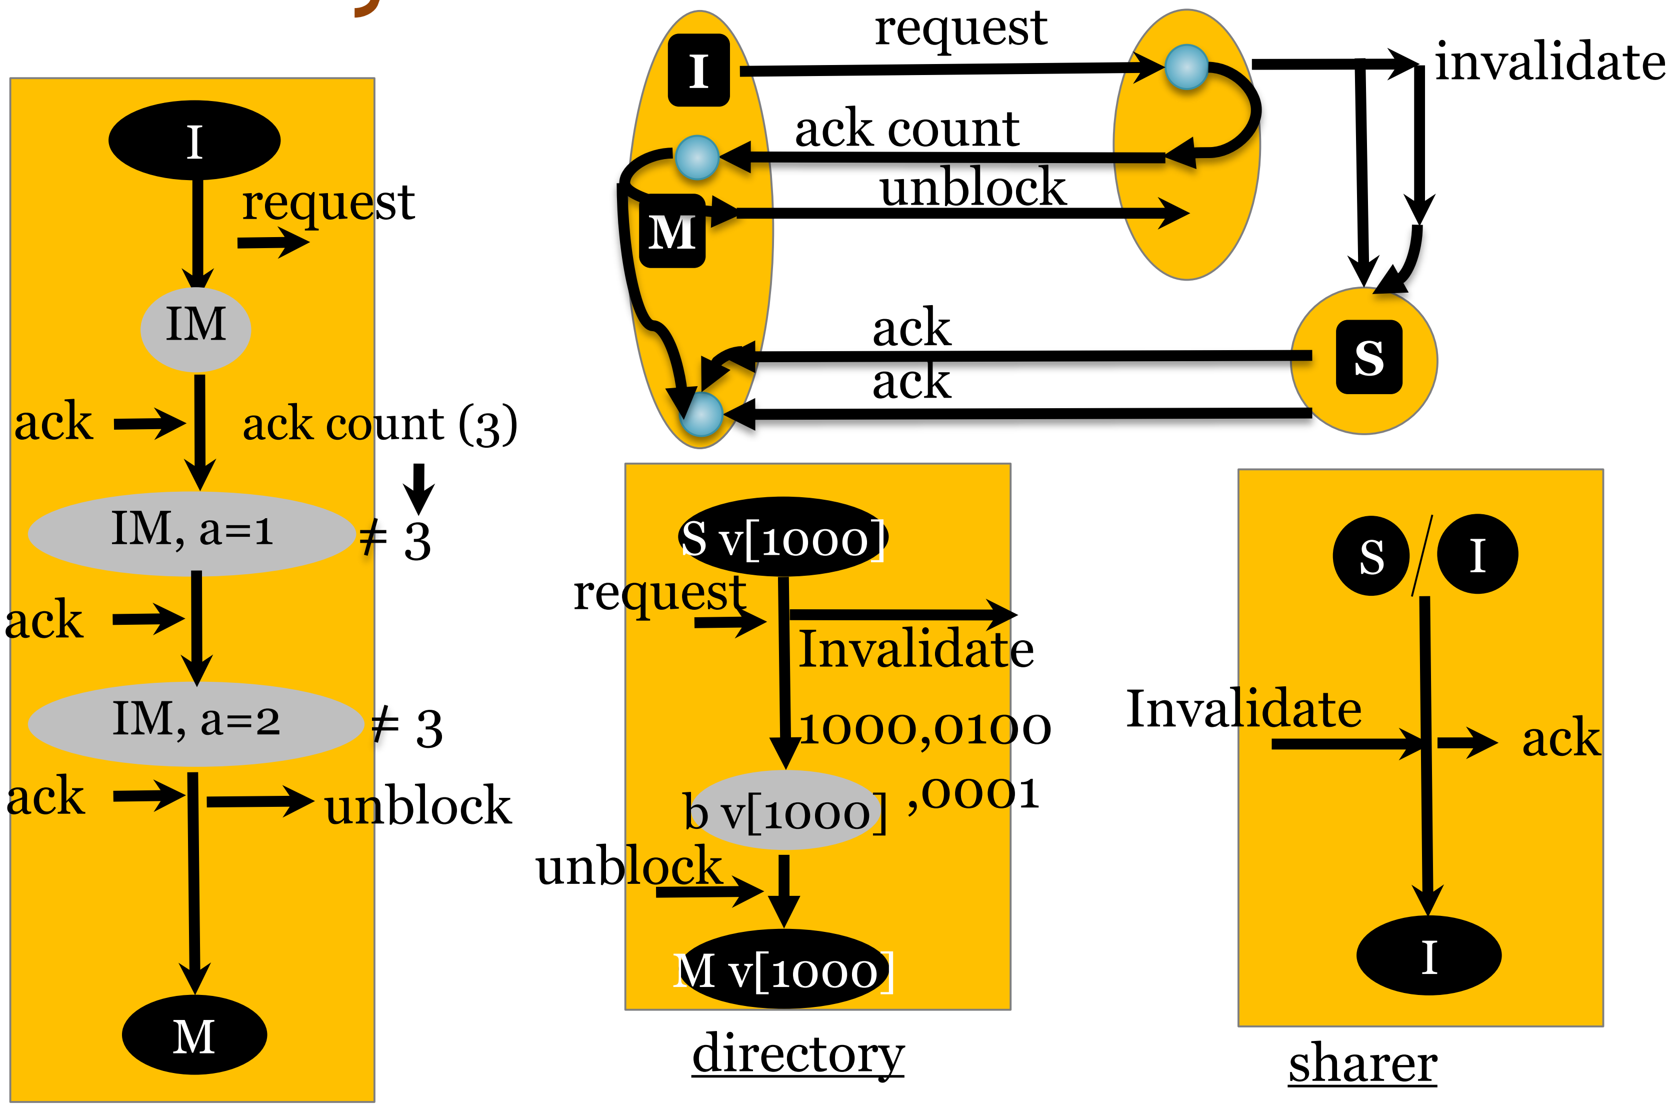

requestor

c) data: S, request: M

Enforcing property 3

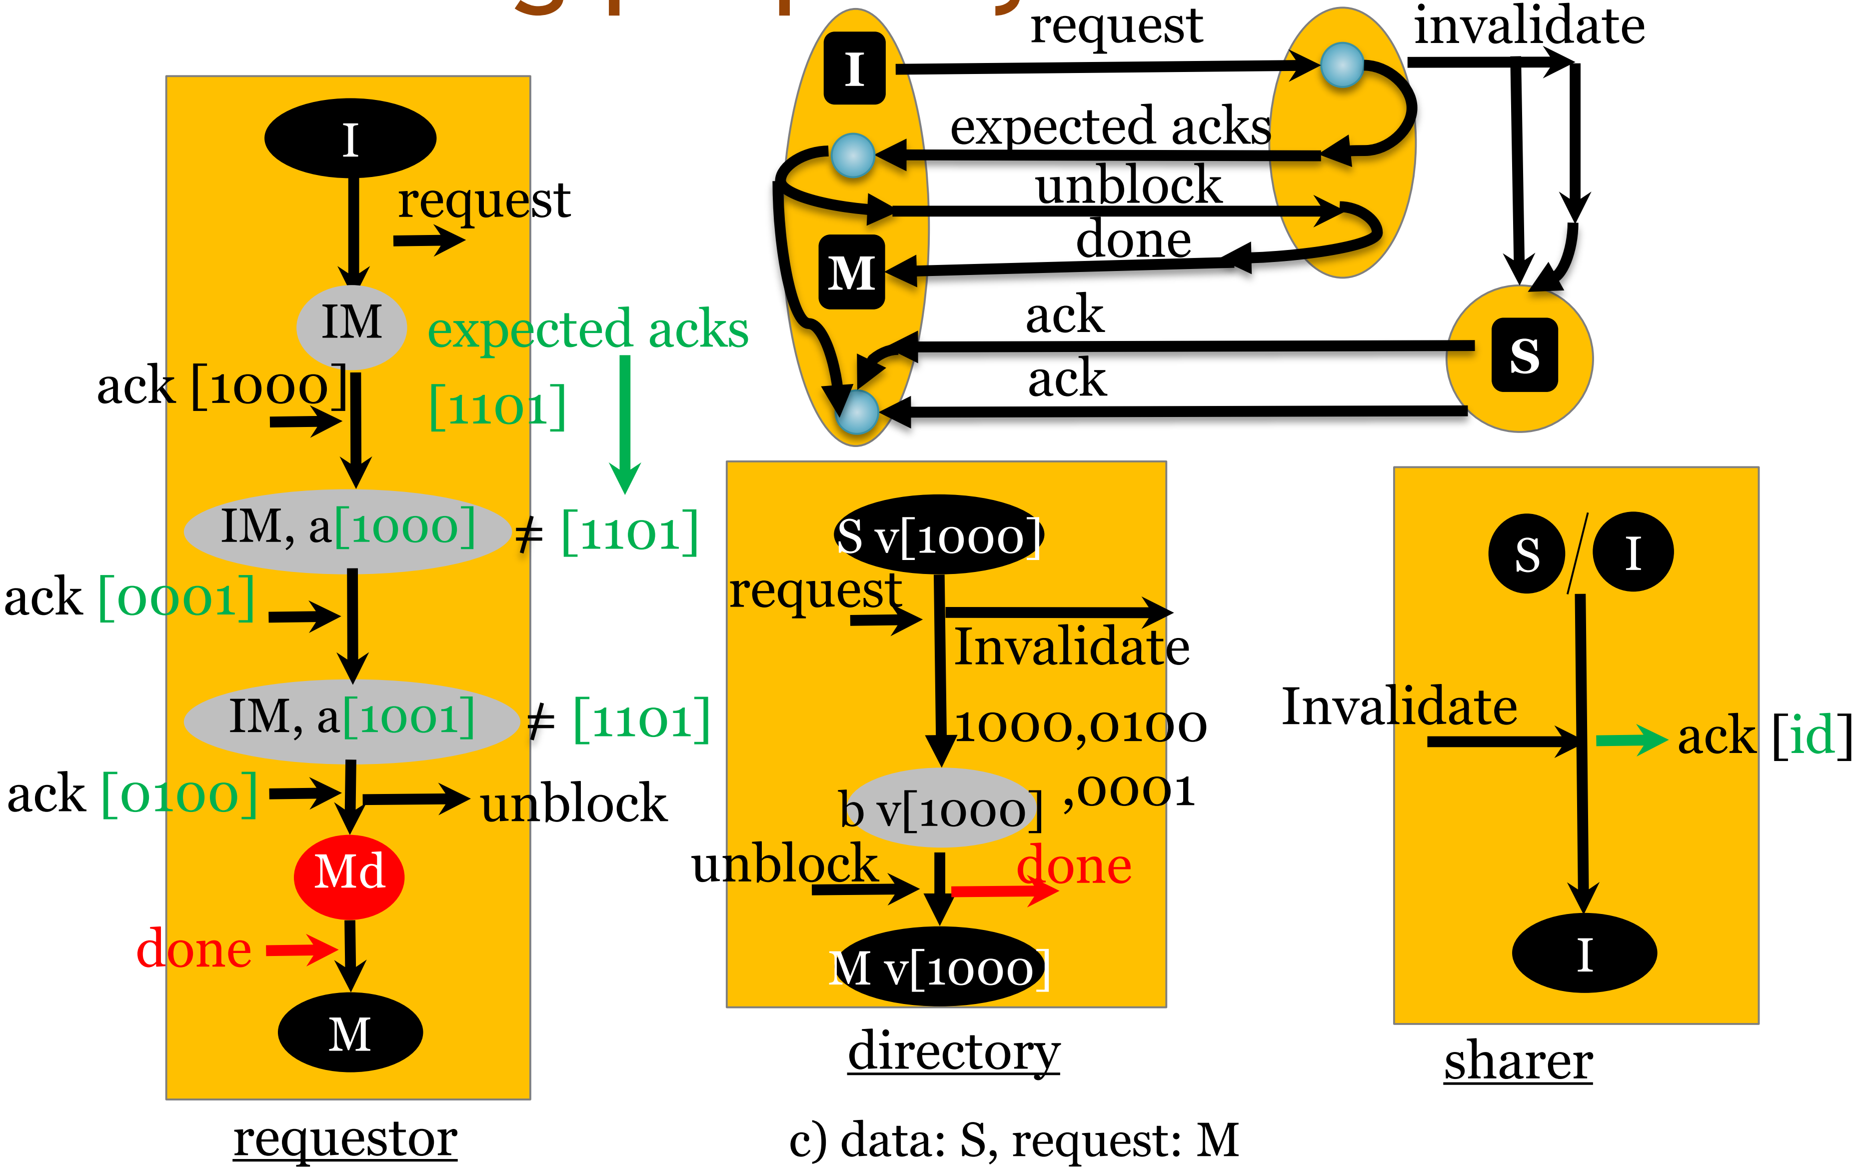

c) data: S, request: M

Experimental setup

- Wisconsin Multifacet GEMS simulator
- 64-core tiled CMP
- Private split L1 caches
- Physically distributed shared L2 cache
- Fault rates of 1 fault/ms – 1 fault/μs

Performance overhead

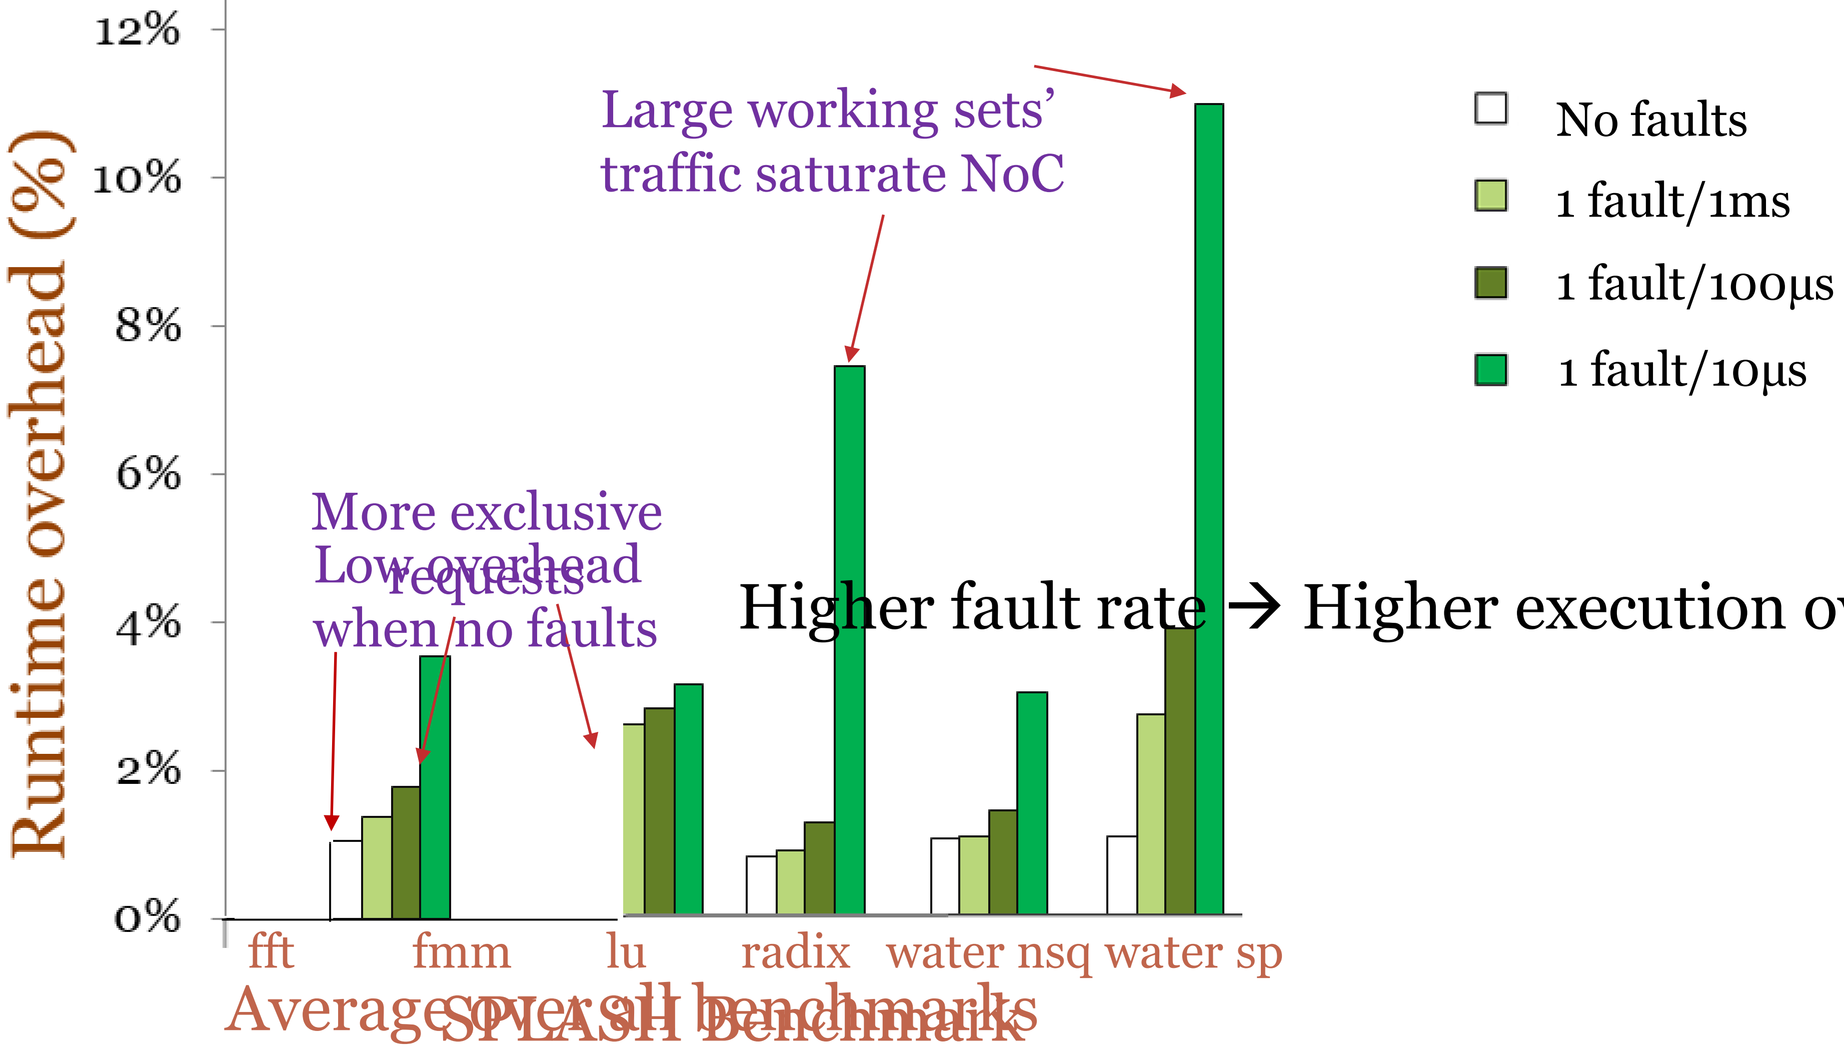

Higher fault rate → Higher execution overhead

Network congestion

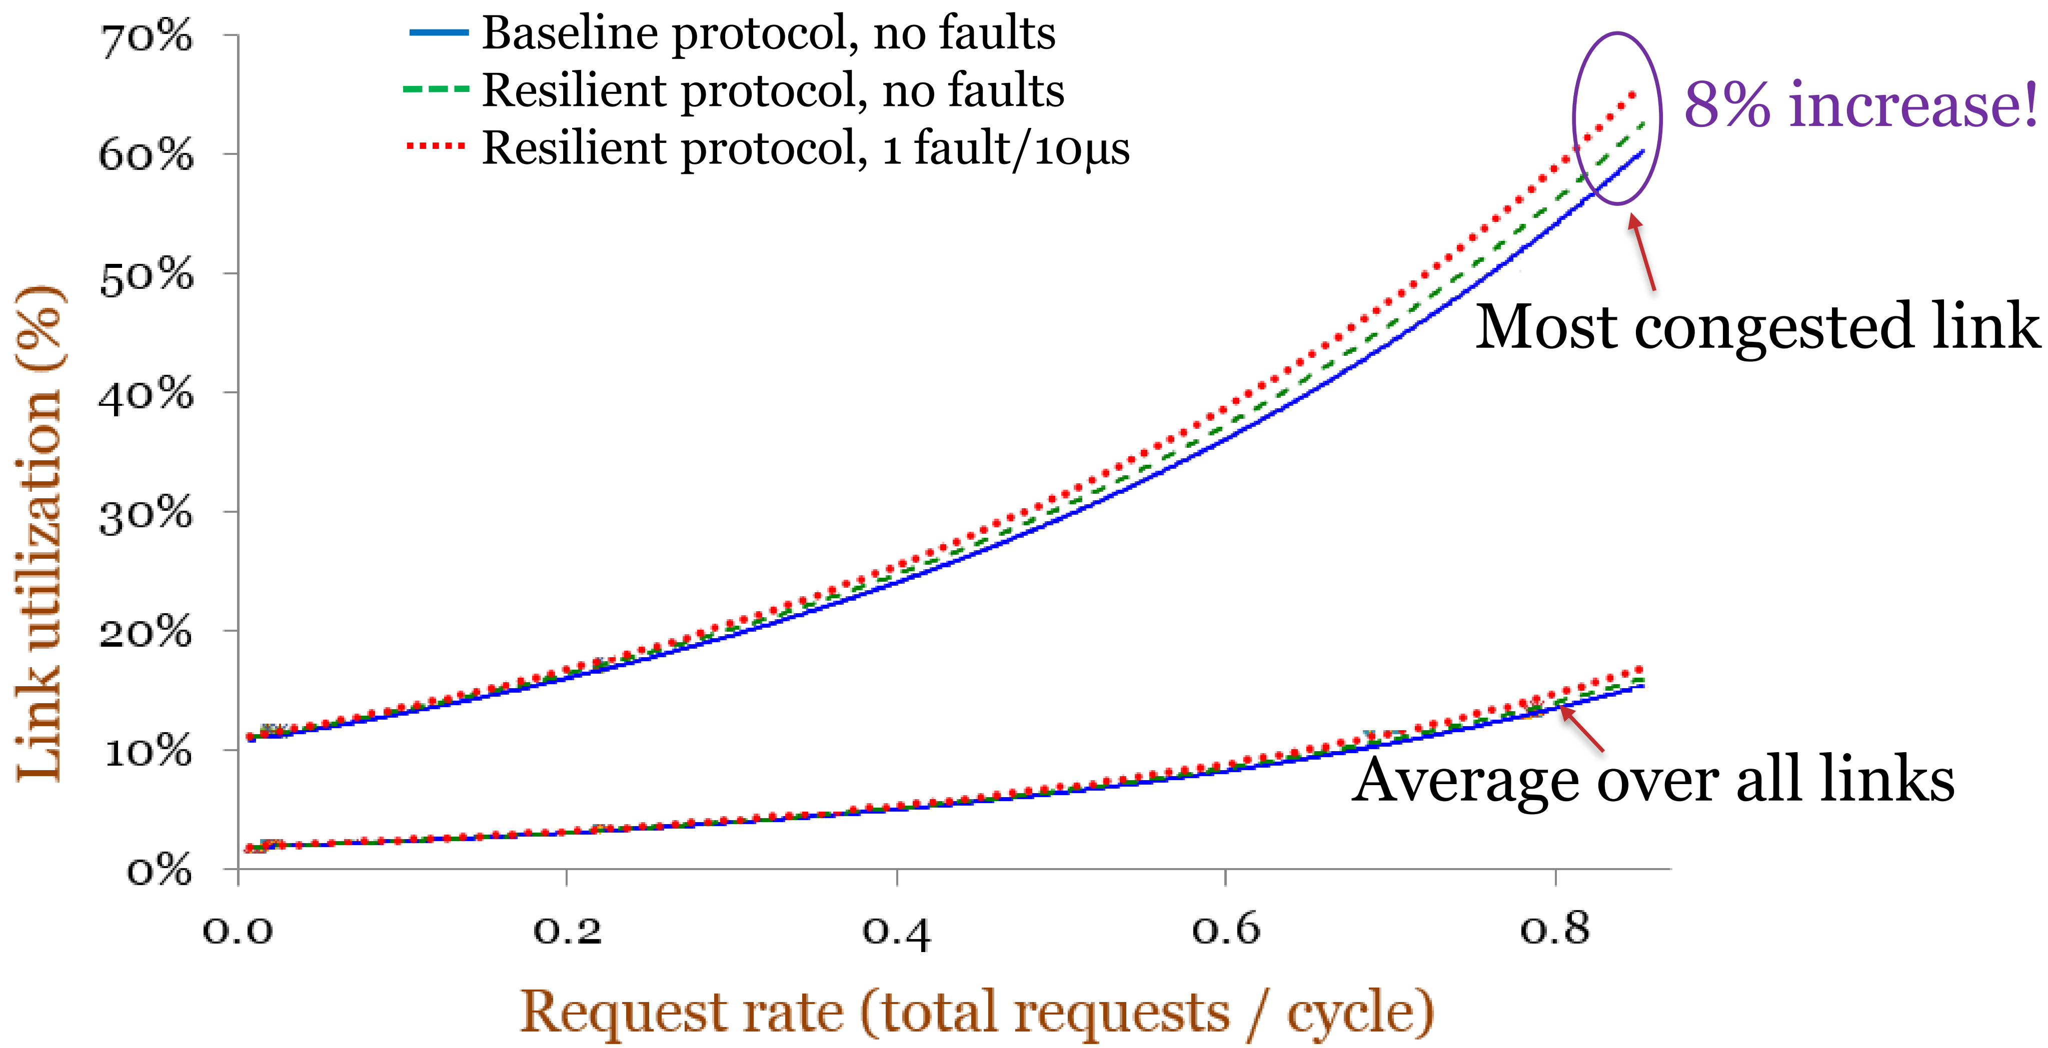

Hardware overhead

- MSHR table

	PC	Req.	….	State	Transaction ID	Sender bitvector	Timeout
0							
.							
.							
.							
31							
				←1 bit→	← 6 bits →	← 64 bits →	←13 bits →

- Router
  - Adder per buffered packet

Total= 352 bytes/node, 20 X 16 bit adder/router << core gate count

Conclusion

- Lost messages lead to suspended transactions.
- Three properties were defined that guarantee transactions will eventually complete.
- Experimental results indicated negligible hardware overhead and execution overhead of 0.8% during fault-free operation.

Questions?

Discussion

- Does addressing only transient faults guarantee sufficient resilience?
- The resilient version of the protocol is much more elaborate than the baseline. Is this worth it?

THANK YOU!