

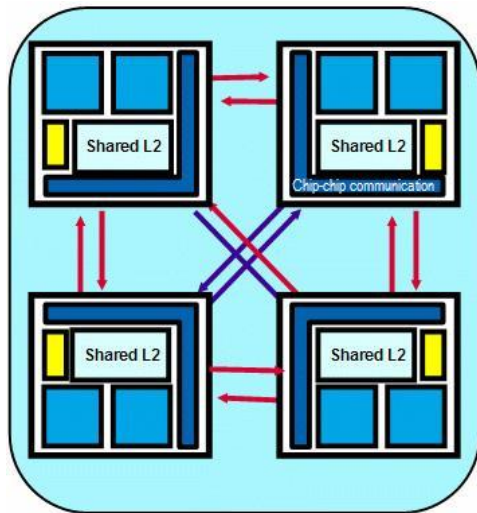
NUMA-Friendly Stack (using Delegation and Elimination)

Irina Calciu
Justin Gottschlich
Maurice Herlihy


HotPar '13



Trends for Future Architectures




Uniform Memory Access (UMA)



Non-Uniform Memory Access (NUMA)

NUMA NODE (multiple cores, shared Last Level Cache)

NUMA NODE (multiple cores, shared Last Level Cache)



NUMA NODE (multiple cores, shared Last Level Cache)

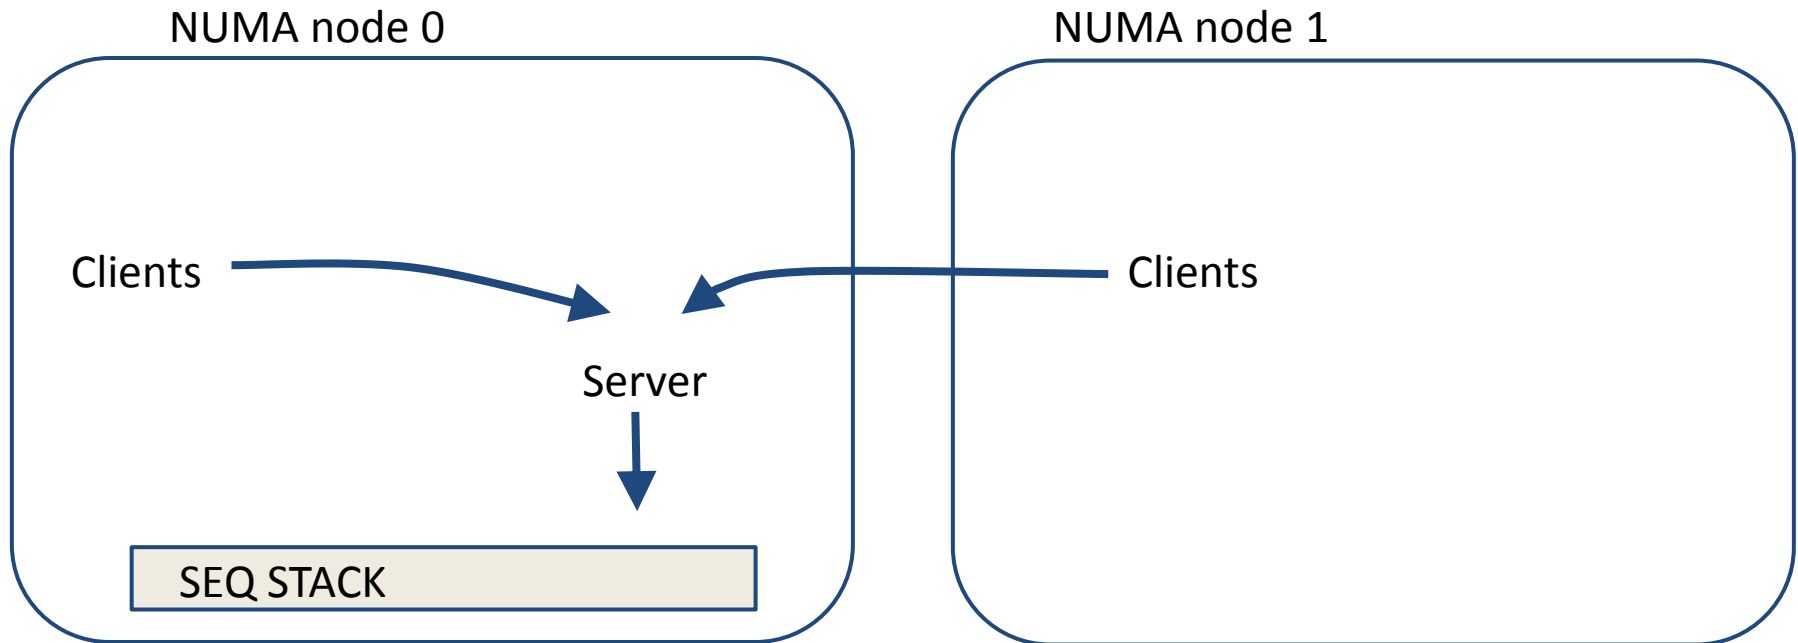
NUMA NODE (multiple cores, shared Last Level Cache)

Cache coherency maintained between caches on different NUMA nodes

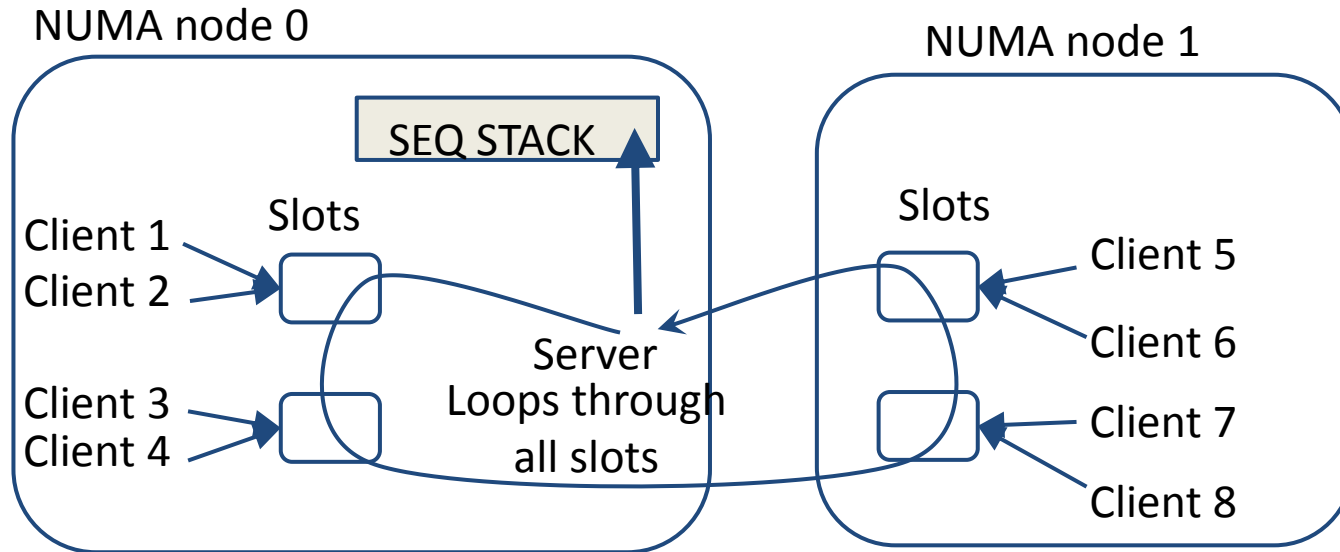
Overview

- Motivation
- Algorithms
- Results
- Conclusions

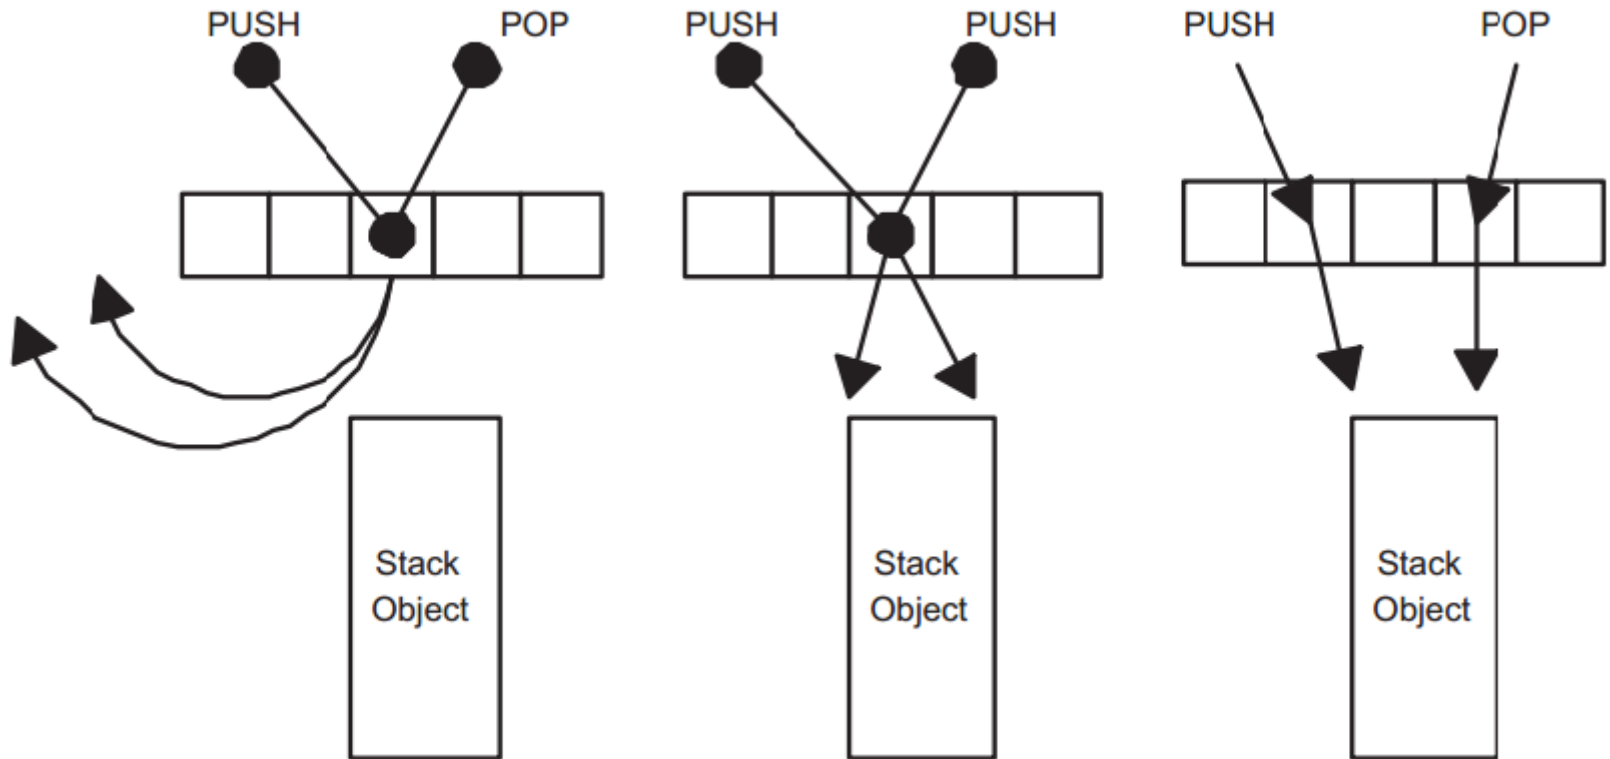
Delegation



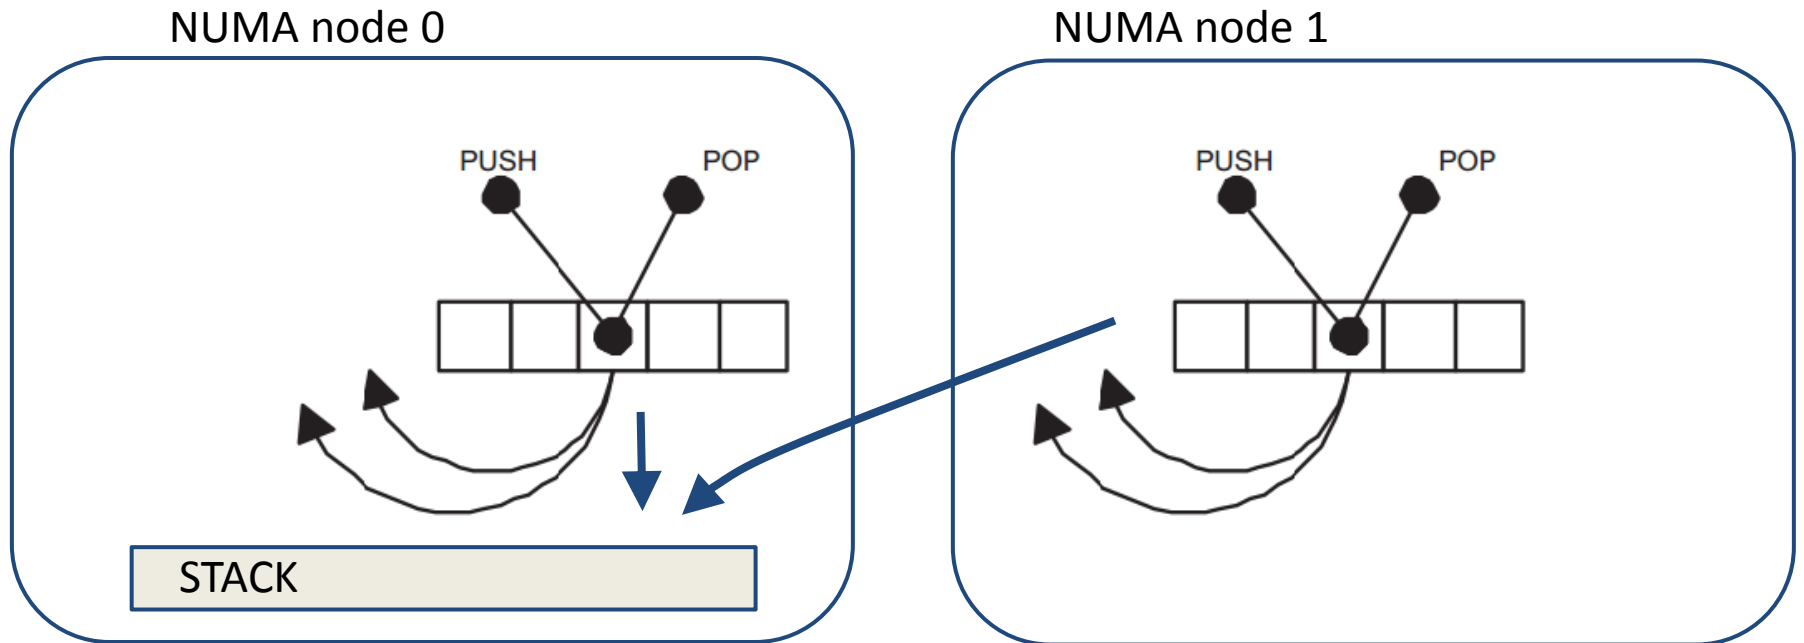
Delegation



Elimination, Rendezvous



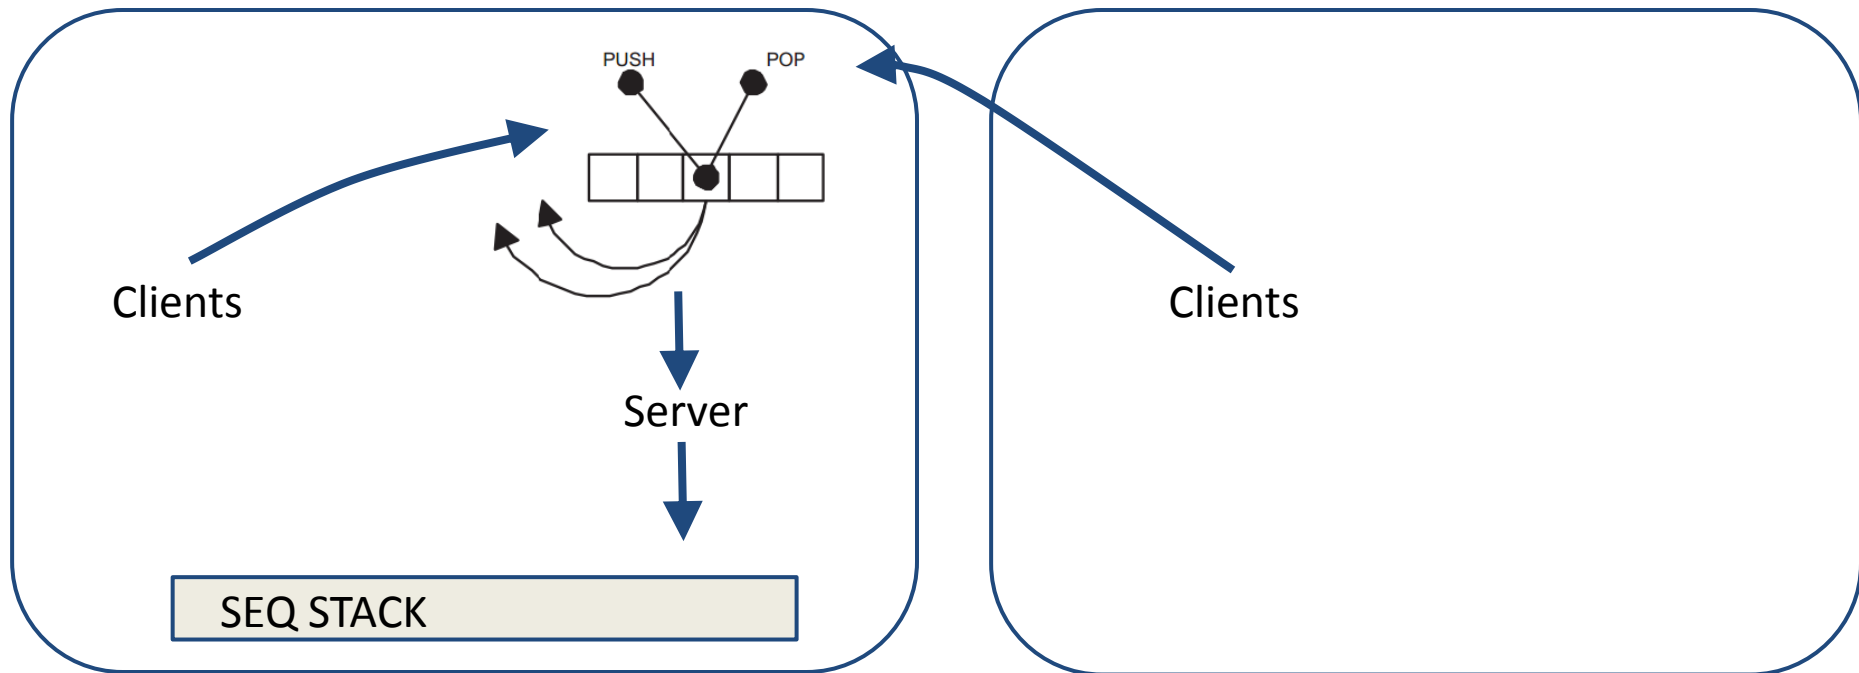
Local Rendezvous



Delegation + Elimination

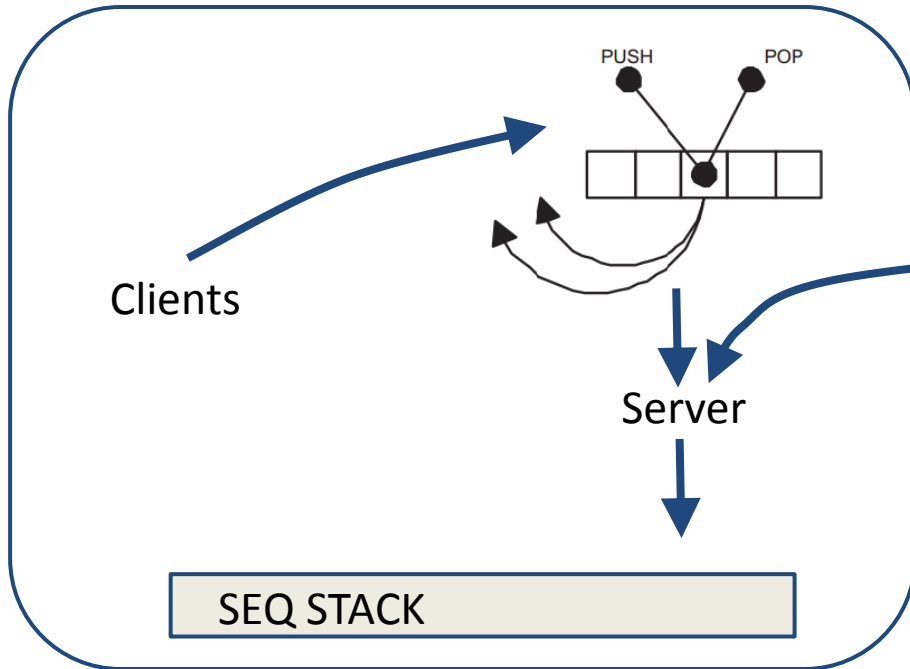
NUMA node 0

NUMA node 1

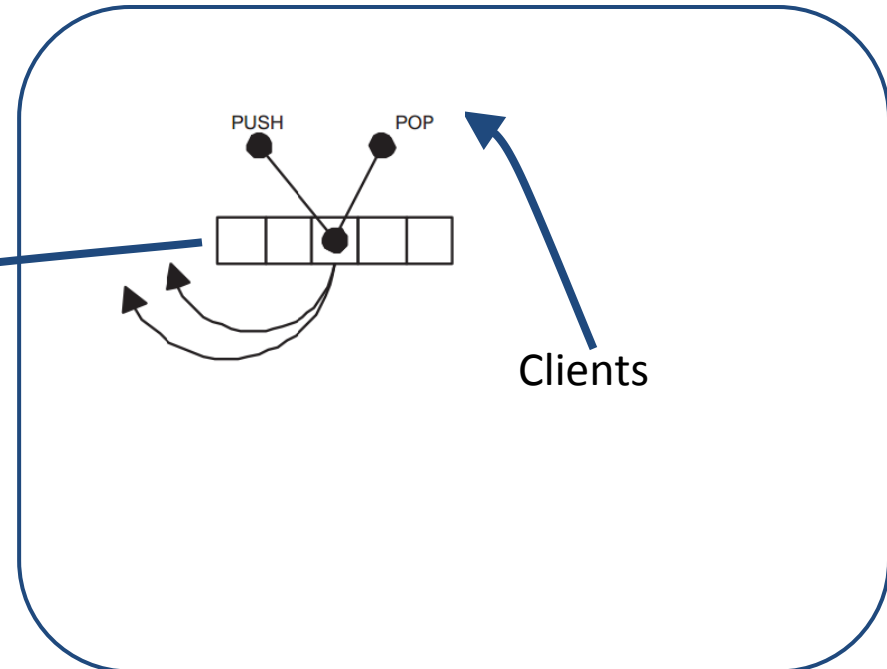


Delegation + LOCAL Elimination

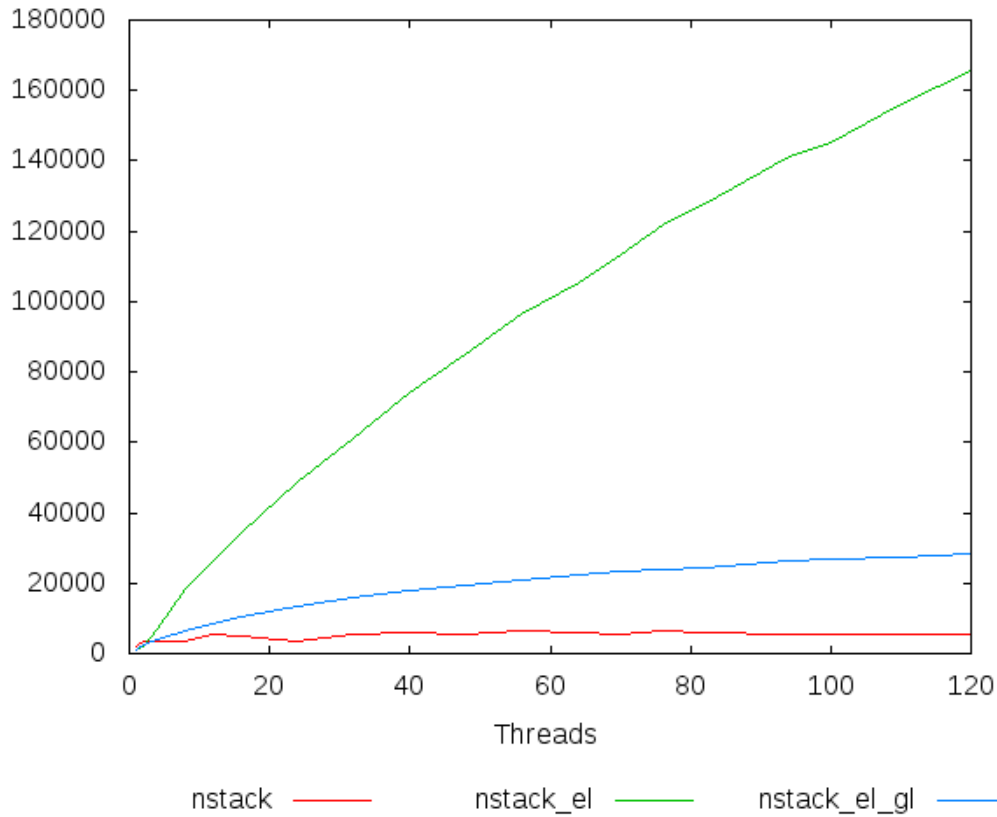
NUMA node 0



NUMA node 1



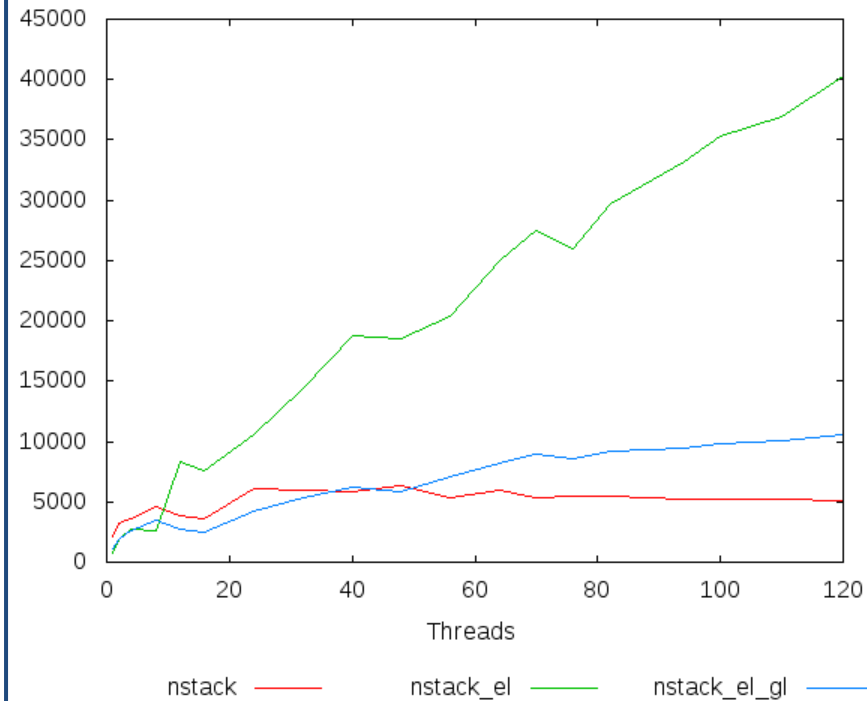
Effect of Elimination




50% push 50% pop

↑ Throughput (Better)

90% push 10% pop



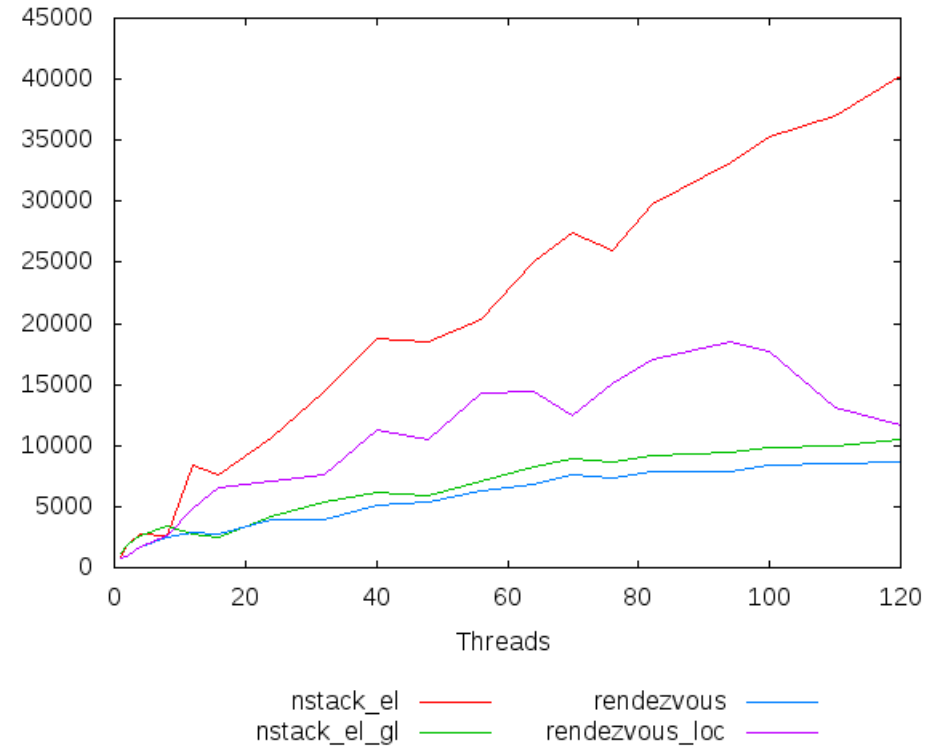
Effect of Delegation




50% push 50% pop

Throughput (Better)

90% push 10% pop



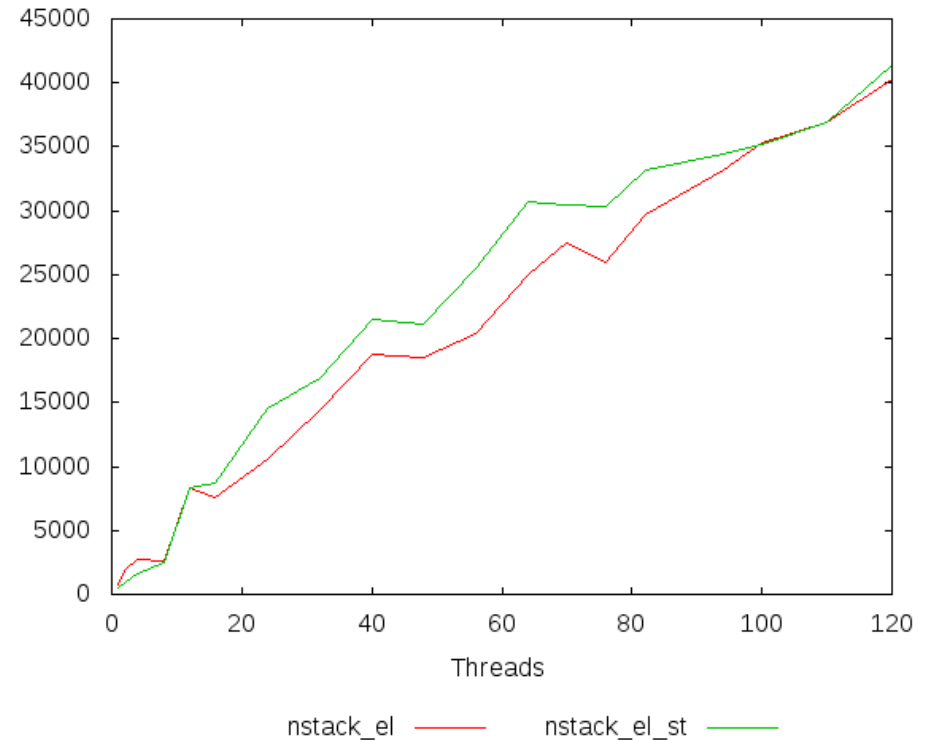
Number of Slots




50% push 50% pop

↑ Throughput (Better)

90% push 10% pop

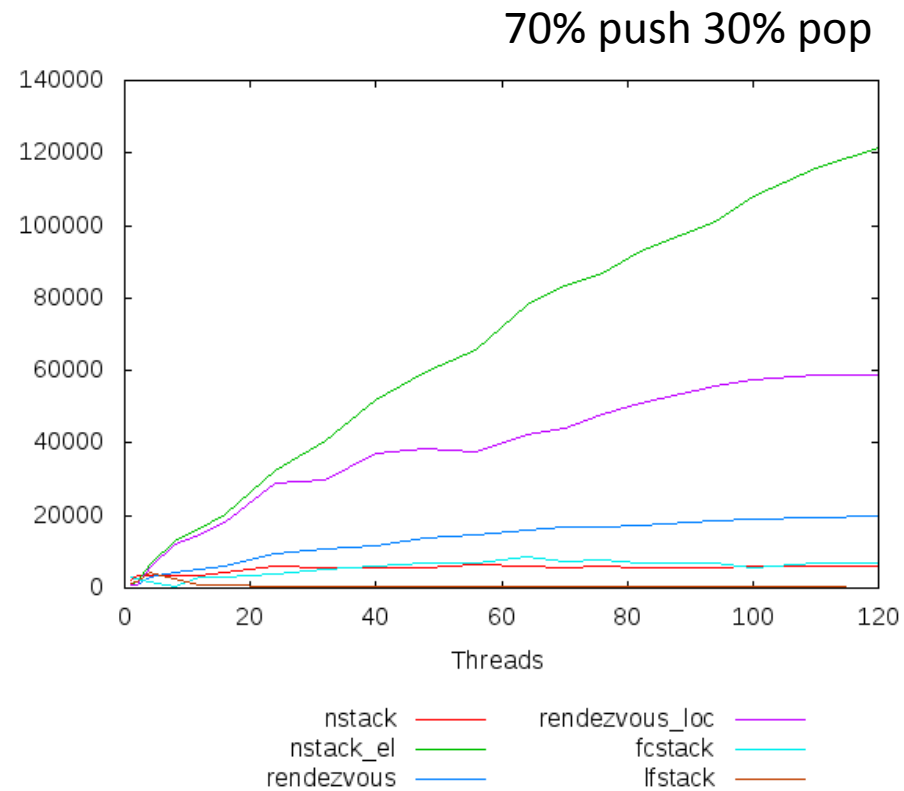


Workloads: Balanced vs. Unbalanced



50% push 50% pop

↑ Throughput (Better)



70% push 30% pop

Advantages

- Memory and cache locality
- Reduced bus traffic
- Increased parallelism through elimination

Drawbacks

- Communication cost between clients and server thread
 - Insignificant compared to the benefits
- Serializing otherwise parallel data structures
 - Parallelism through elimination
- Elimination opportunities decrease as workload more unbalanced

Open Questions

- Are there other data structures where we can use delegation and elimination?
- Are there data structures where direct access is much better?
- What can we do for those data structures?

Thank you!

Questions?



References

- A Scalable Lock-free Stack Algorithm

<http://www.inf.ufsc.br/~dovicchi/pos-ed/pos/artigos/p206-hendler.pdf>

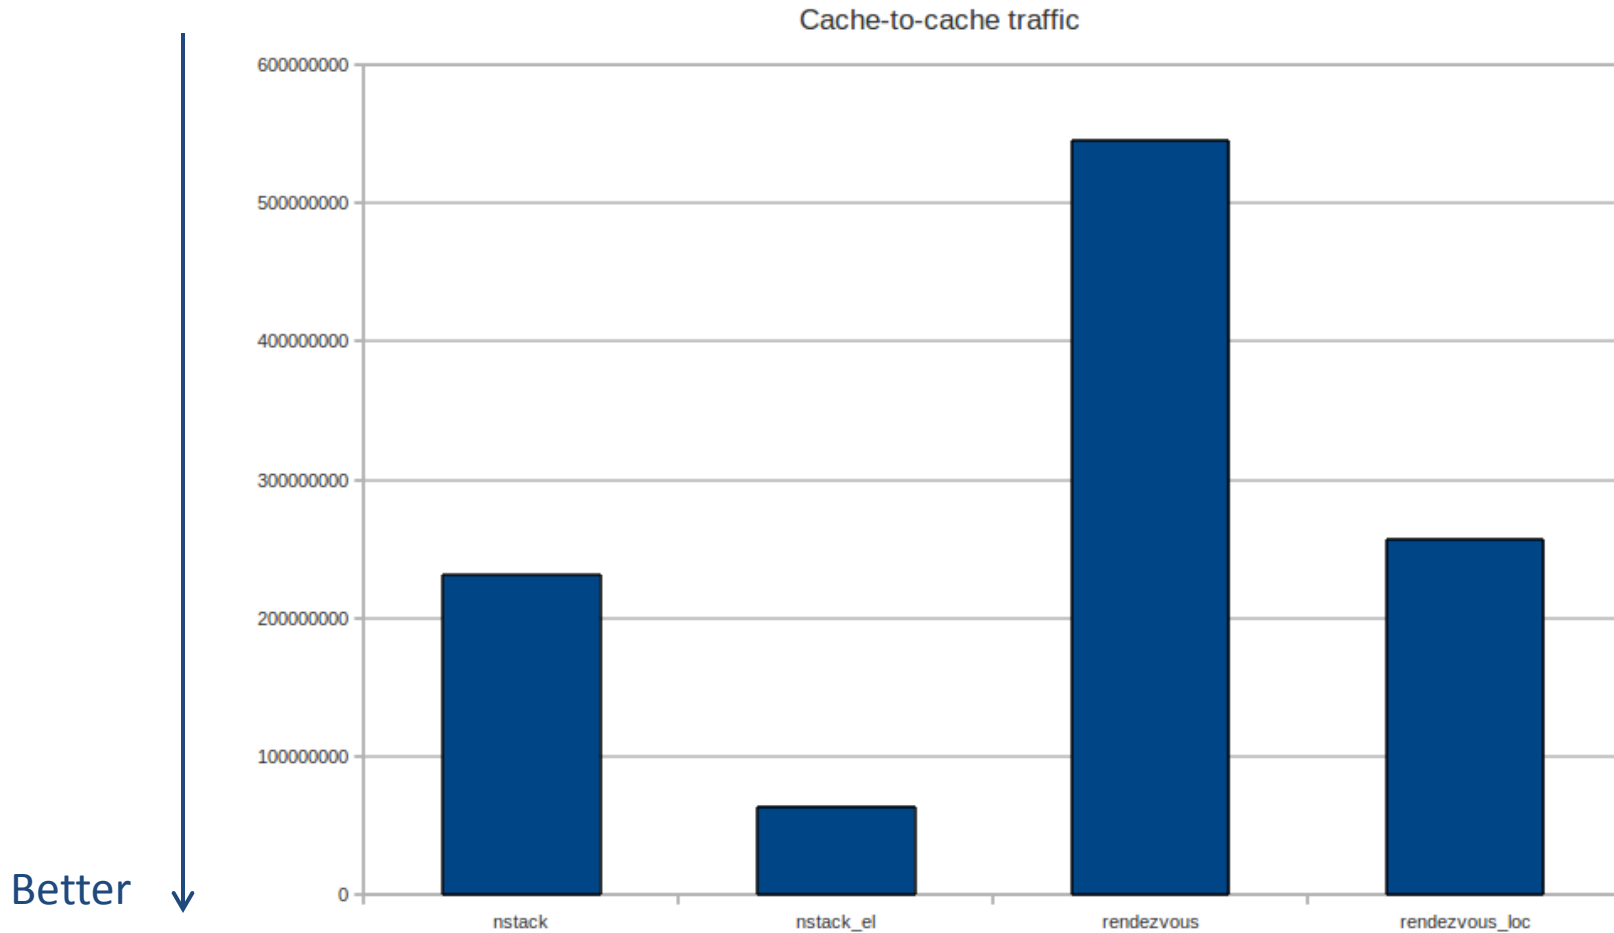
- Flat Combining and the Synchronization-Parallelism Tradeoff

<http://www.cs.bgu.ac.il/~hendlerd/papers/flat-combining.pdf>

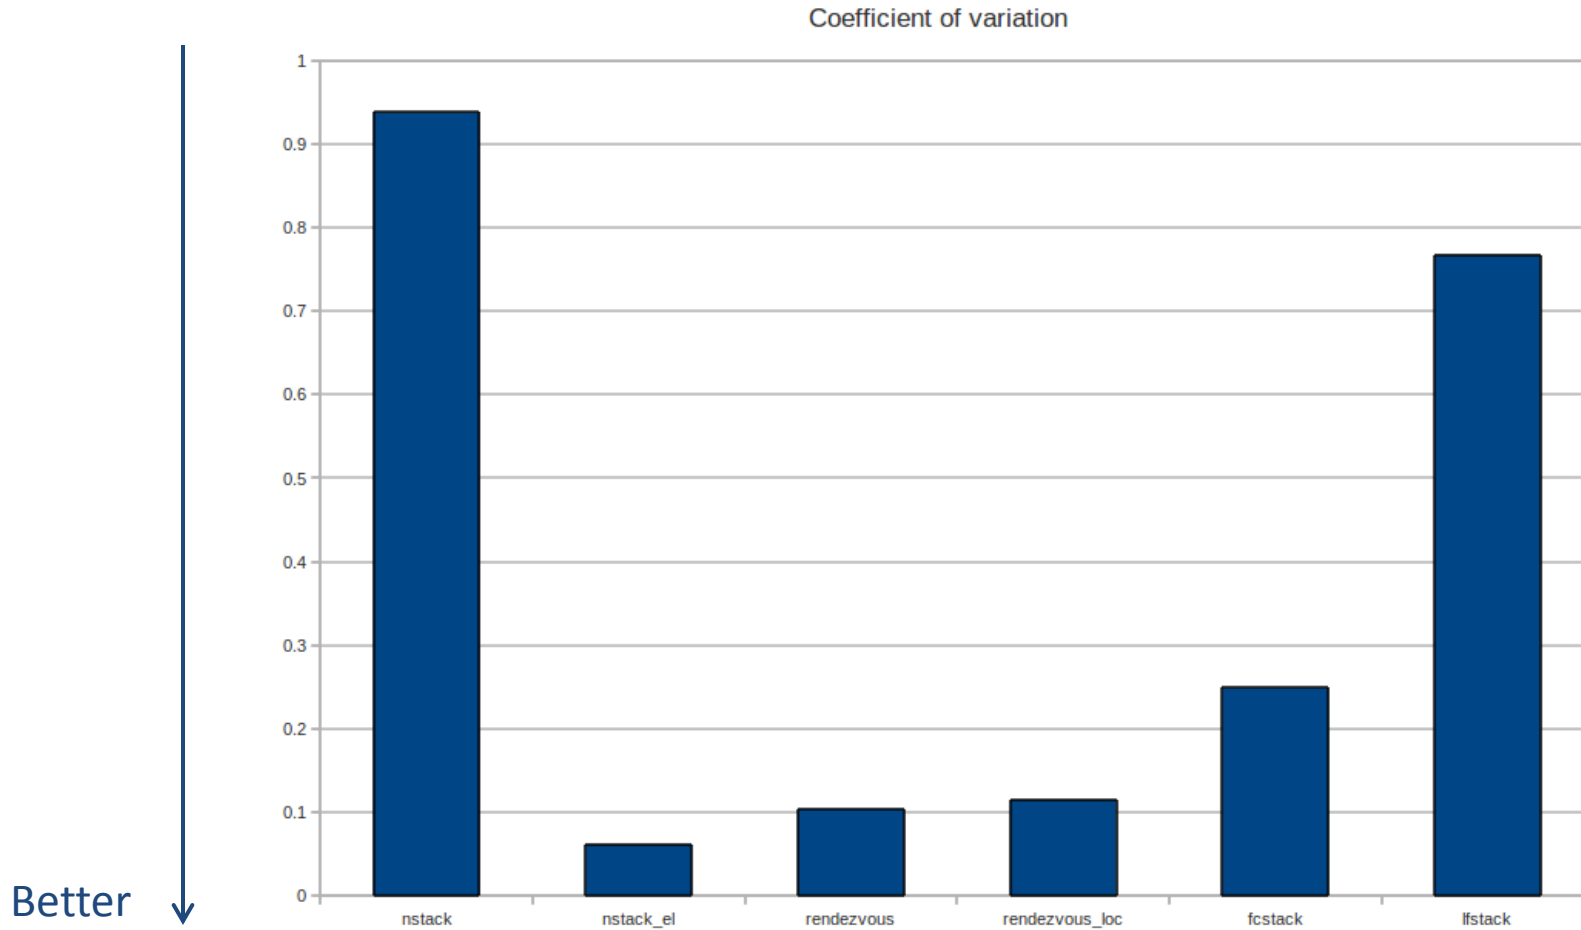
- Fast and Scalable Rendezvousing

<http://www.cs.tau.ac.il/~afek/rendezvous.pdf>

Cache to Cache Traffic




Coefficient of Variation



Flat Combining

④ infrequently, new records are CASed by threads to head of list, and old ones are removed by combiner


① thread writes request and spins on local record



publication list

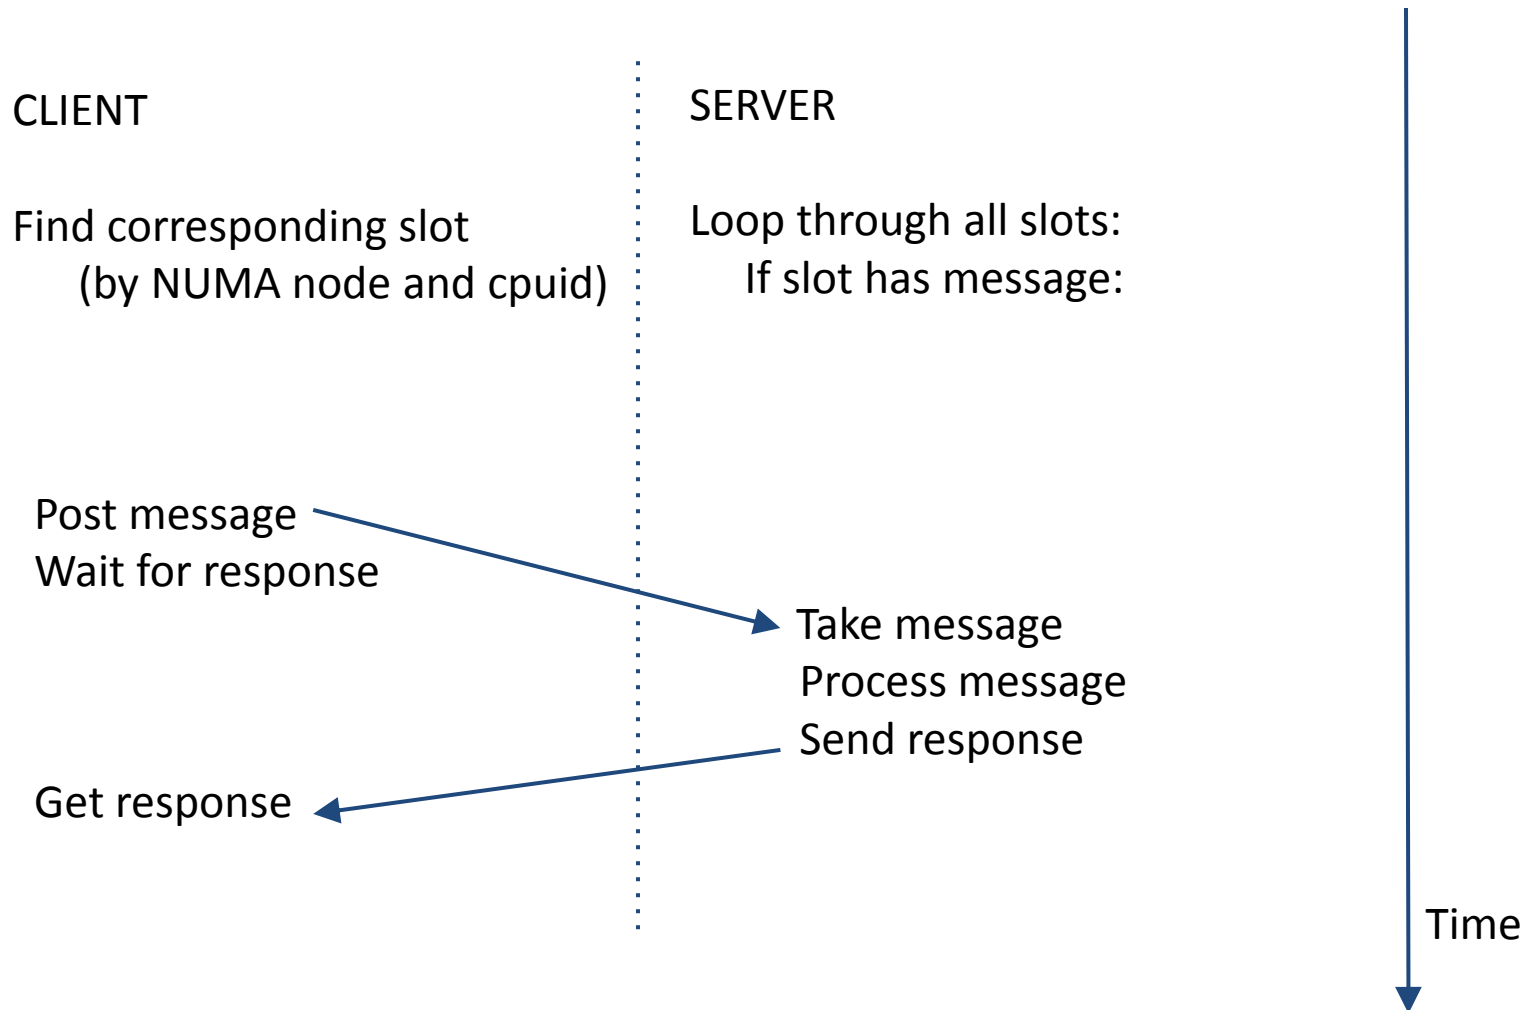
③ combiner traverses list, performs scanCombineApply()

② thread acquires lock, becomes combiner, updates count

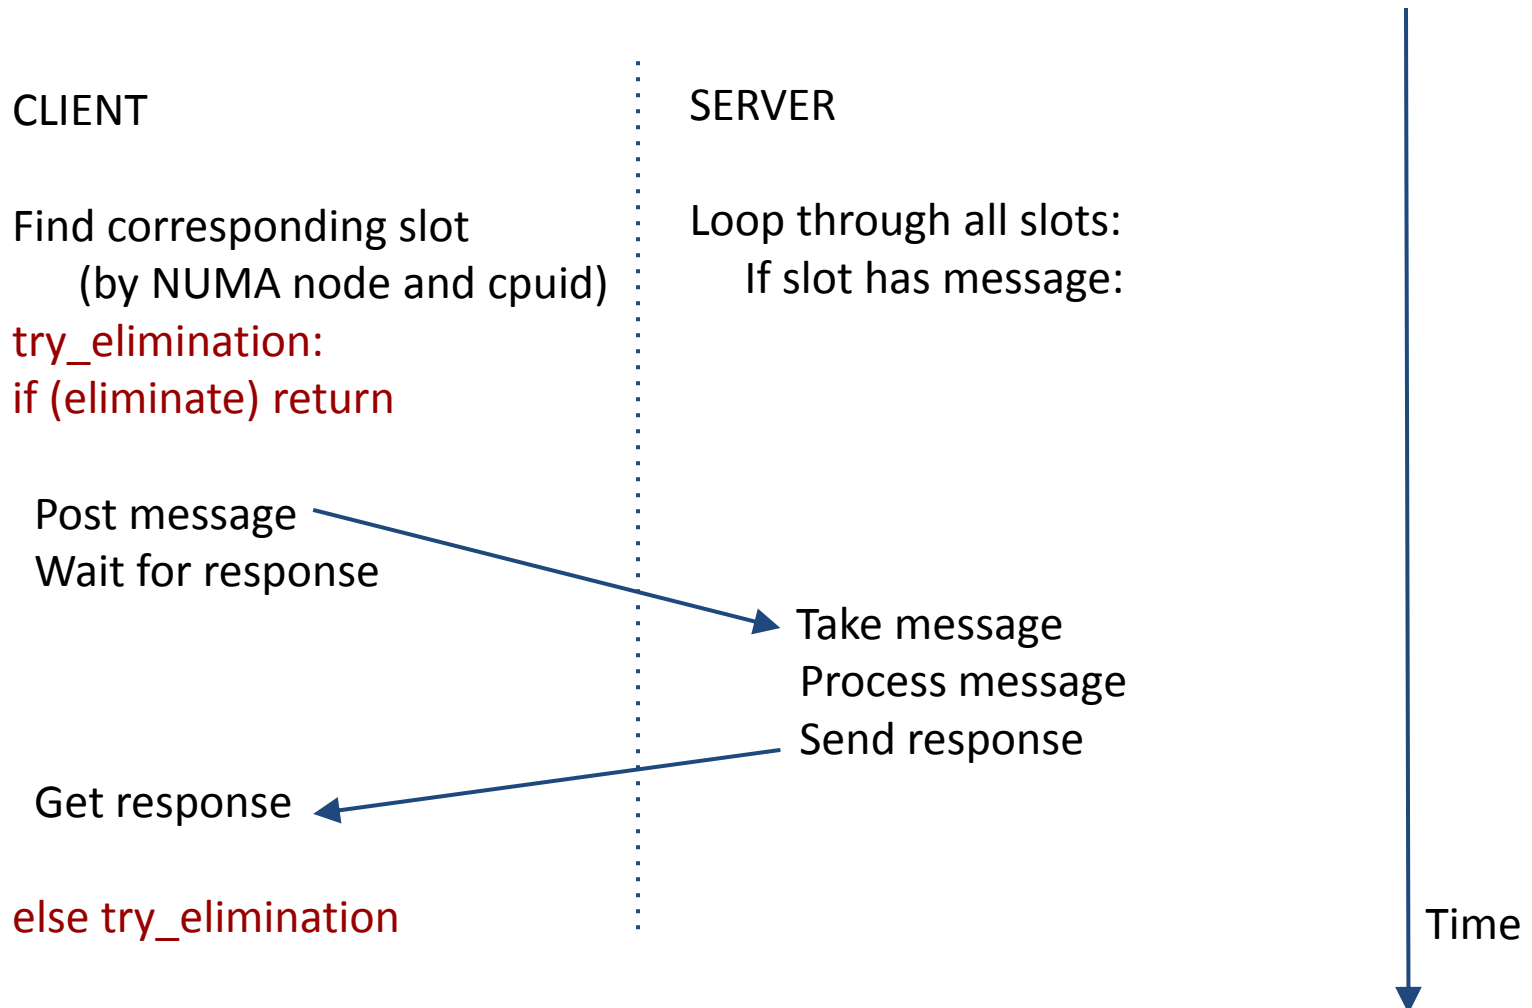


sequential data structure

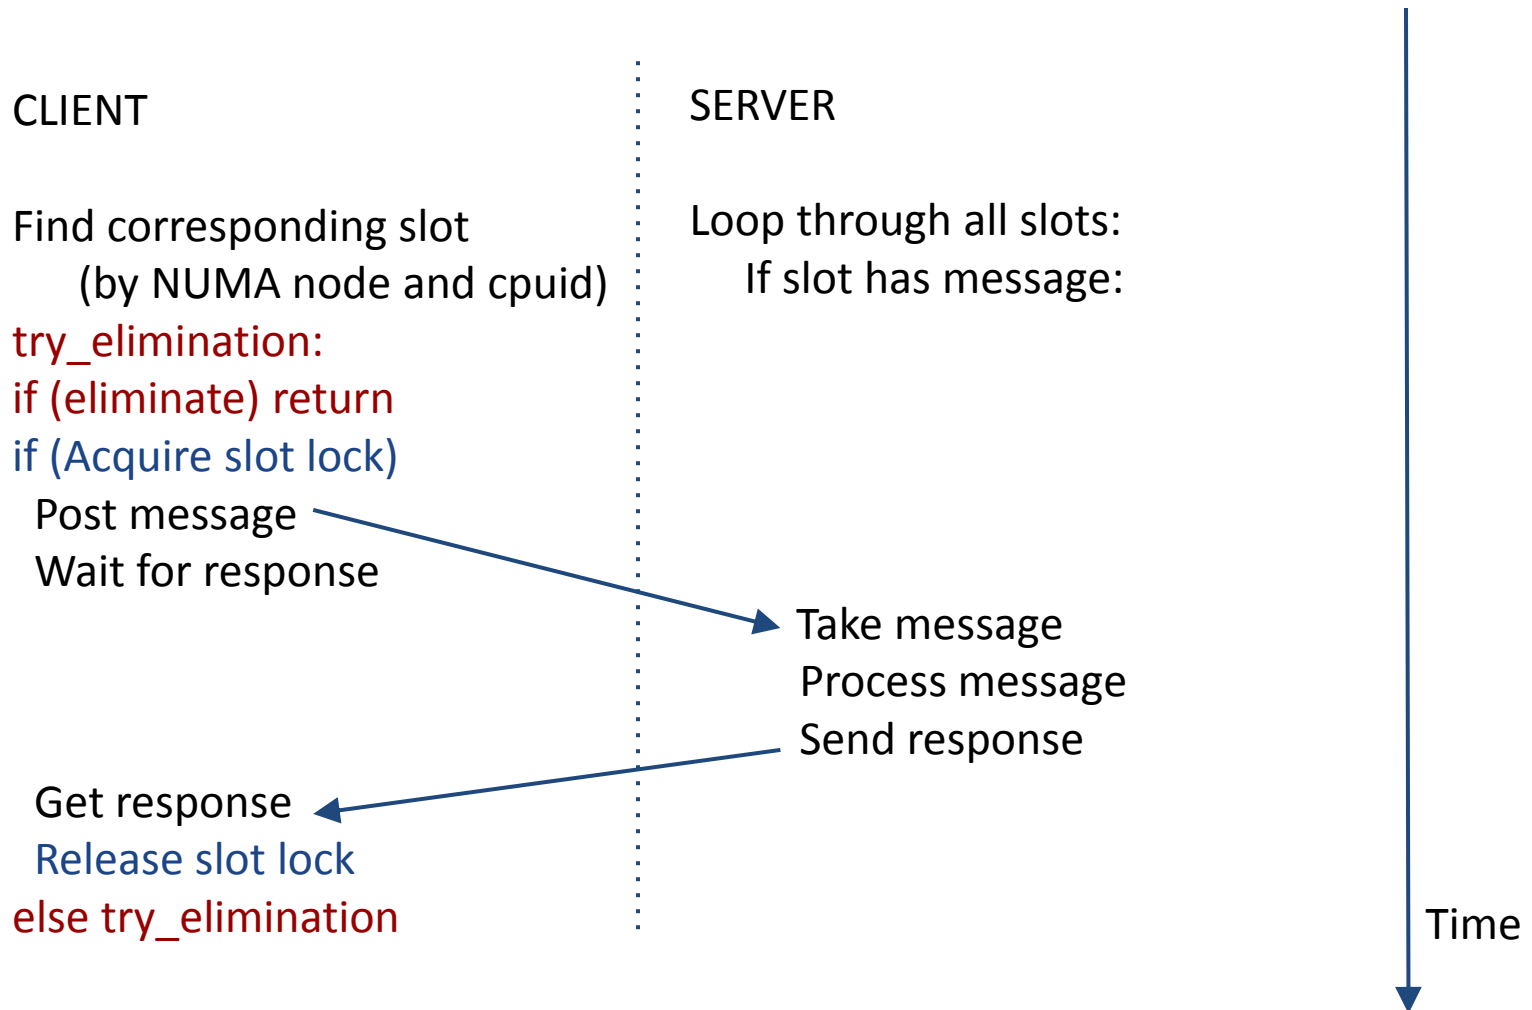
Delegation



Delegation



Delegation



Open Questions

- Performance
- Scalability
- **Power**

