ETH zürich

Consensus in a Box
Inexpensive Coordination in Hardware

Zsolt István, David Sidler, Gustavo Alonso, Marko Vukolic*

Systems Group, Department of Computer Science, ETH Zurich

*IBM Research, Zurich

Systems@ETH zürich

Motivation: Cost of consensus

- Consensus is an essential function in datacenters
- How can consensus be made inexpensive?

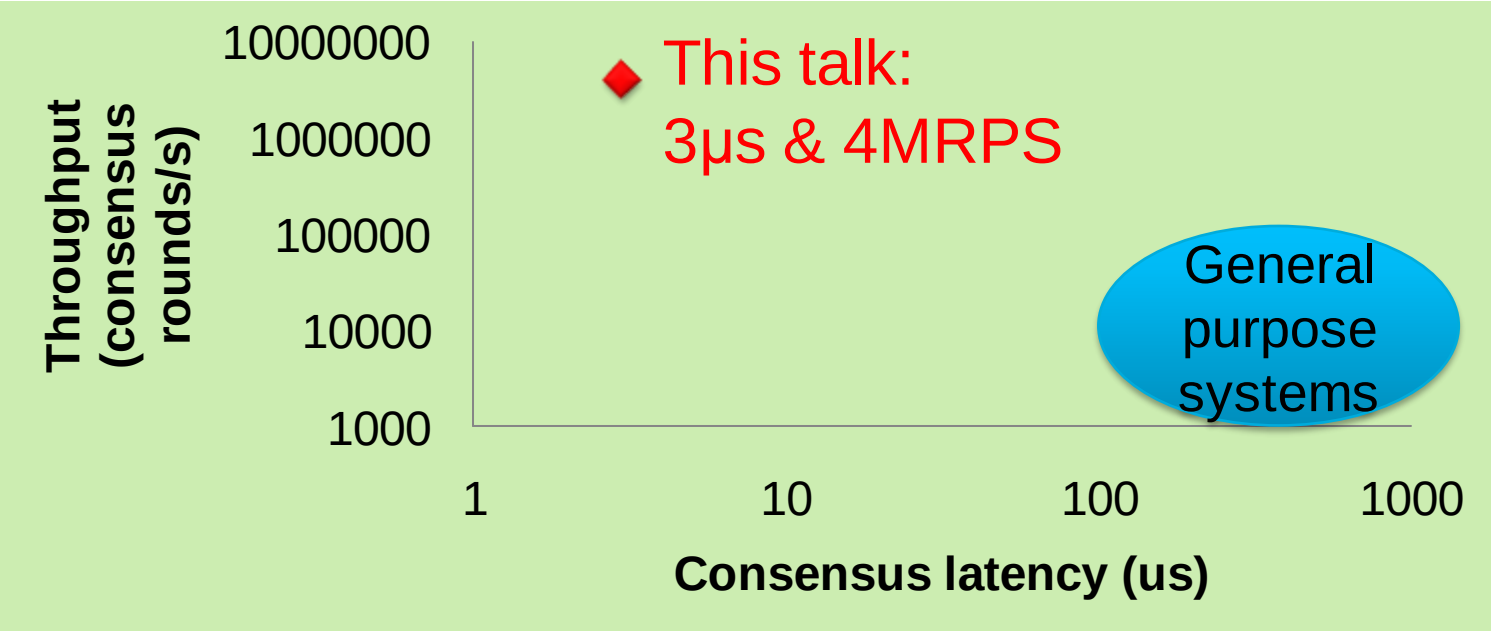

Related work

- Speeding up consensus is an important problem
  - Related work in networking, systems, HPC, etc.

- Specialized hardware can remove traditional limitations

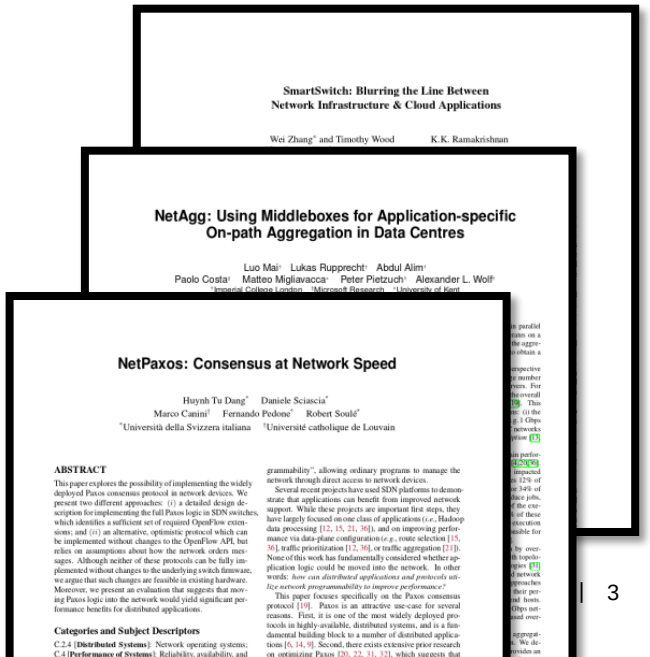

[1] Zhang et al. Smartswitch: Blurring the line between network infrastructure & cloud applications. In HotCloud'14
[2] Mai et al. NetAgg: Using Middleboxes for Application-Specific On-path Aggregation in Data Centres. In CoNEXT'14
[3] Dang et al. NetPaxos: Consensus at Network Speed. In SOSR'15
[4] Poke et al. DARE: High-Performance State Machine Replication on RDMA Networks. In HPDC'15.

Consensus in a Box

- **Why?** Consensus is expensive, but desired
- **What?** Atomic broadcast – Zookeeper's ZAB protocol
- **How?** Specialized processor on FPGA
  - Tight integration with 10Gbps networking + deep pipelining
- **Evaluation?** Drop-in replacement for Memcached with Zookeeper's replication

Zookeeper's Atomic Broadcast

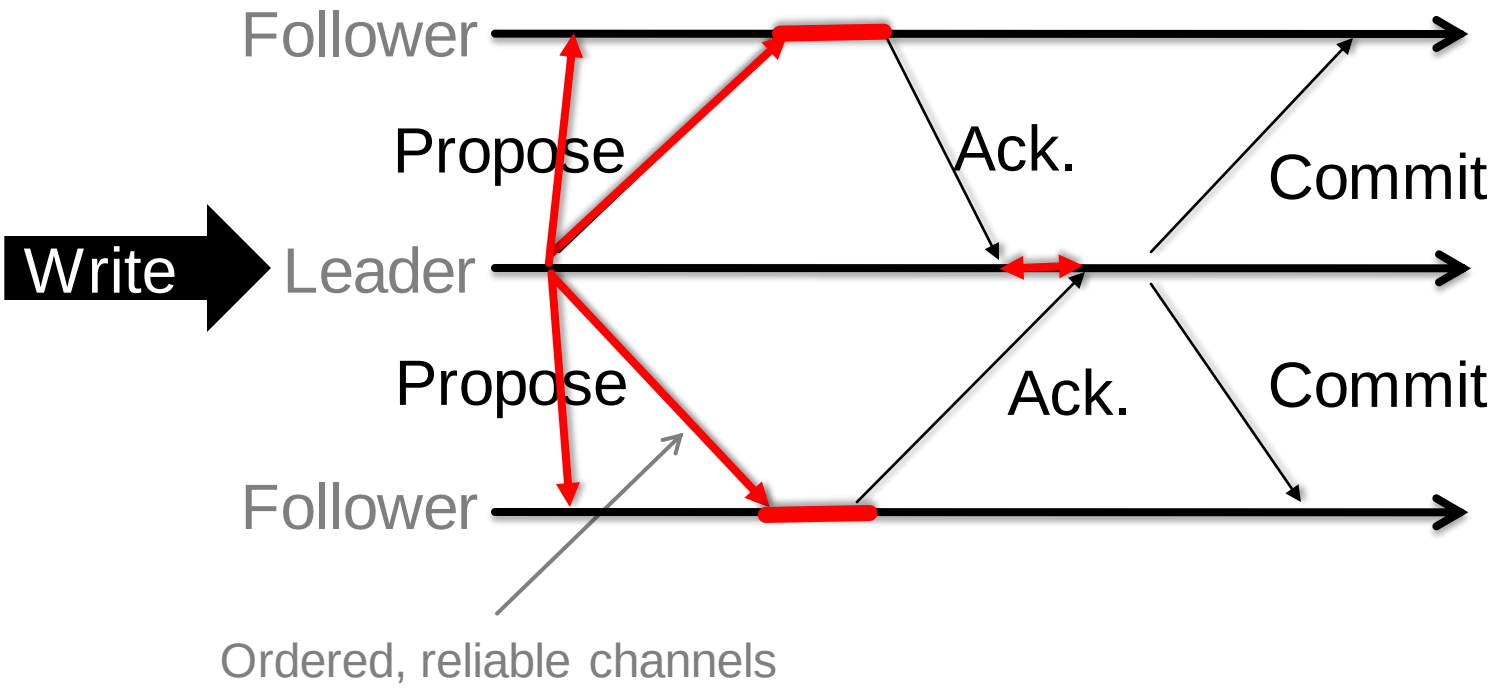

Zookeeper from 10000ft

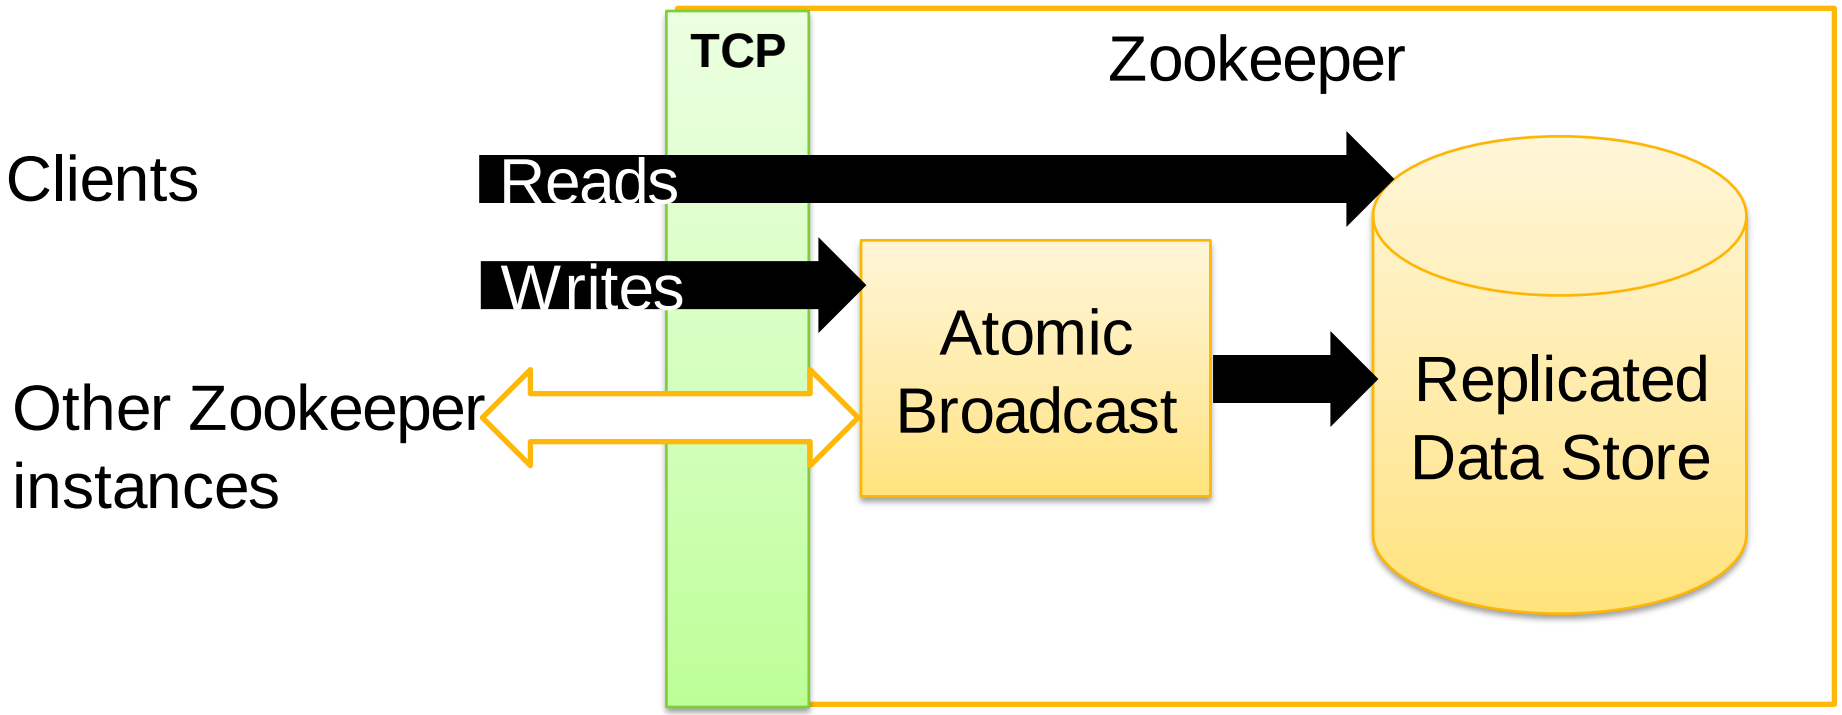

Specialized processor architecture

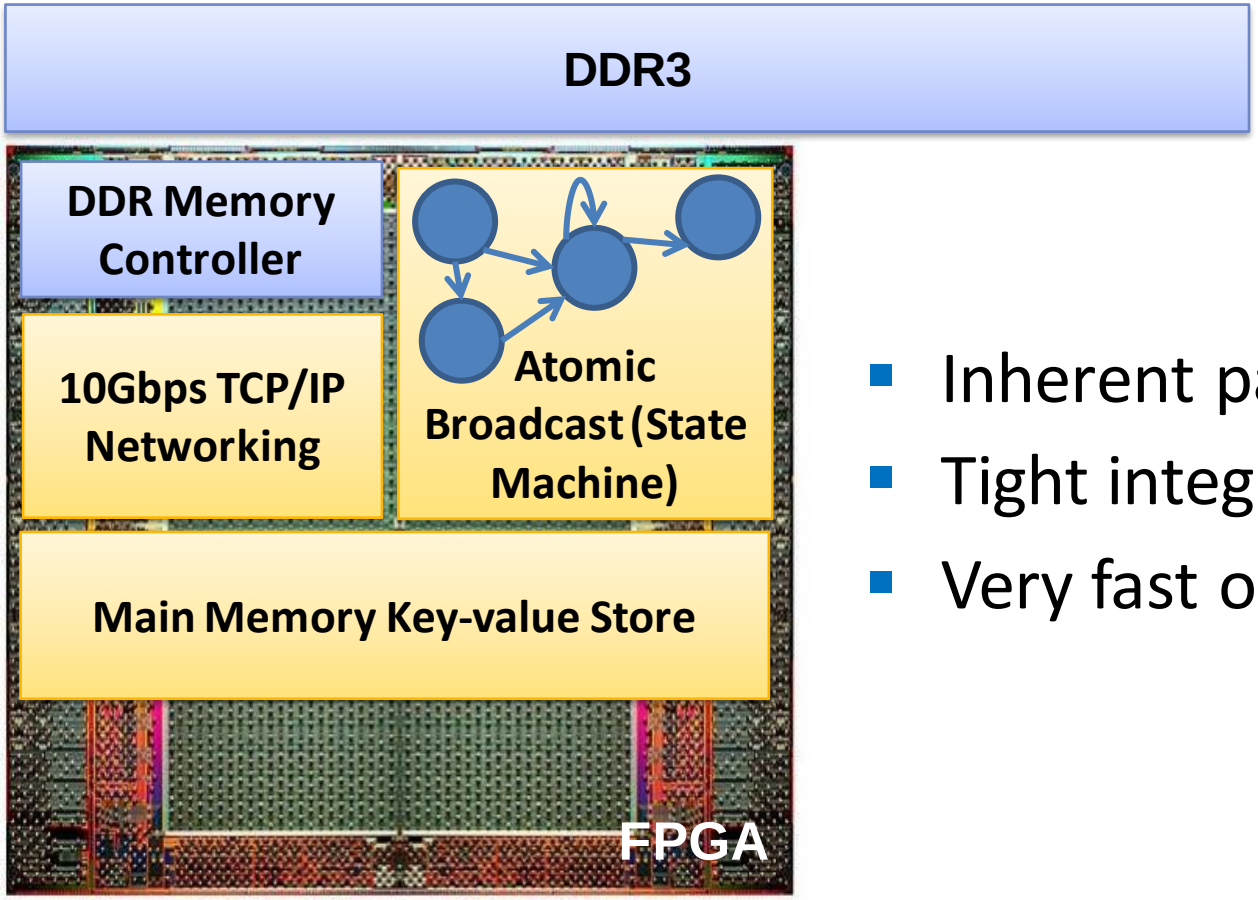

- Inherent parallelism
- Tight integration
- Very fast on-chip memory

What makes it go fast…

Latency

- **Networking optimizations**
  - Low-latency on-chip buffers for RX path
  - Datacenter and application-specific knowledge
- **Predictable behavior**
  - Fast local caches for common case behavior

Throughput

- **Pipelined execution**

Deployment and Evaluation

Hardware platform

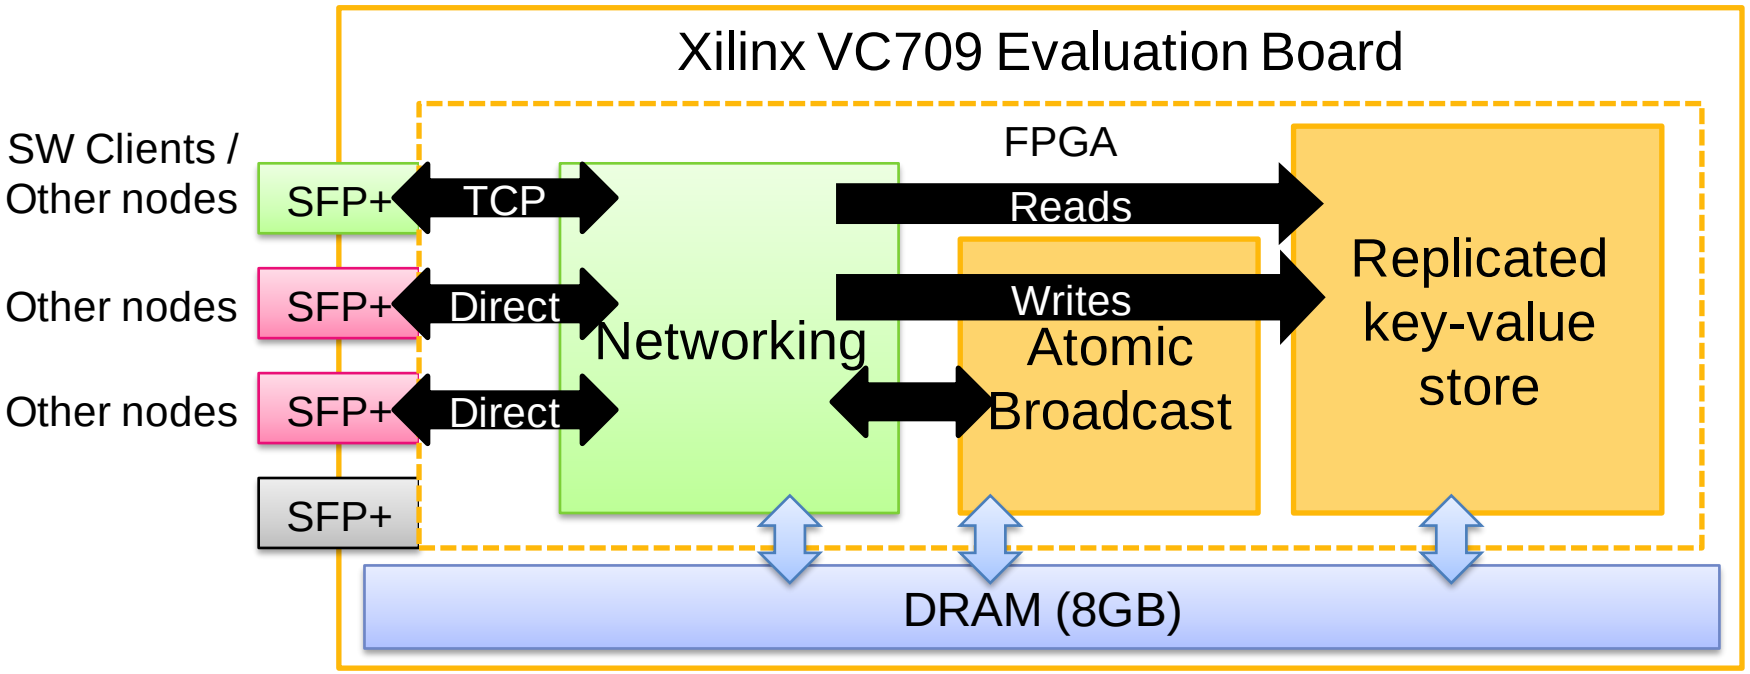

Evaluation setup

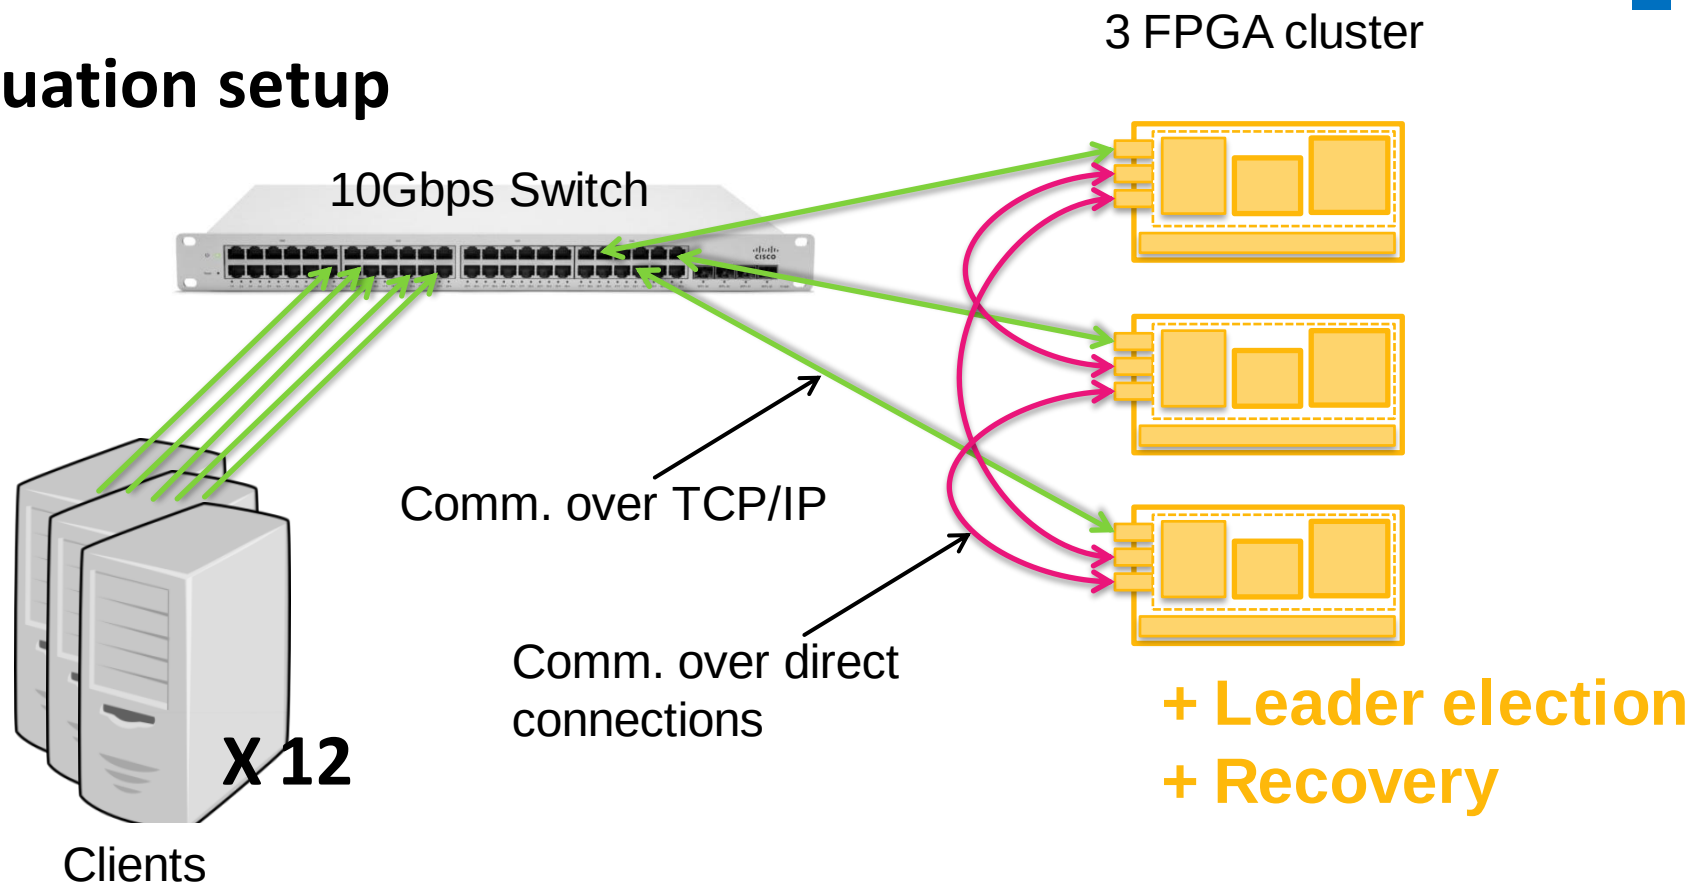

- Drop-in replacement for Memcached with Zookeeper's replication
- Standard tools for benchmarking (libmemcached)
  - Simulating 100s of clients

Latency of KVS writes (consensus)

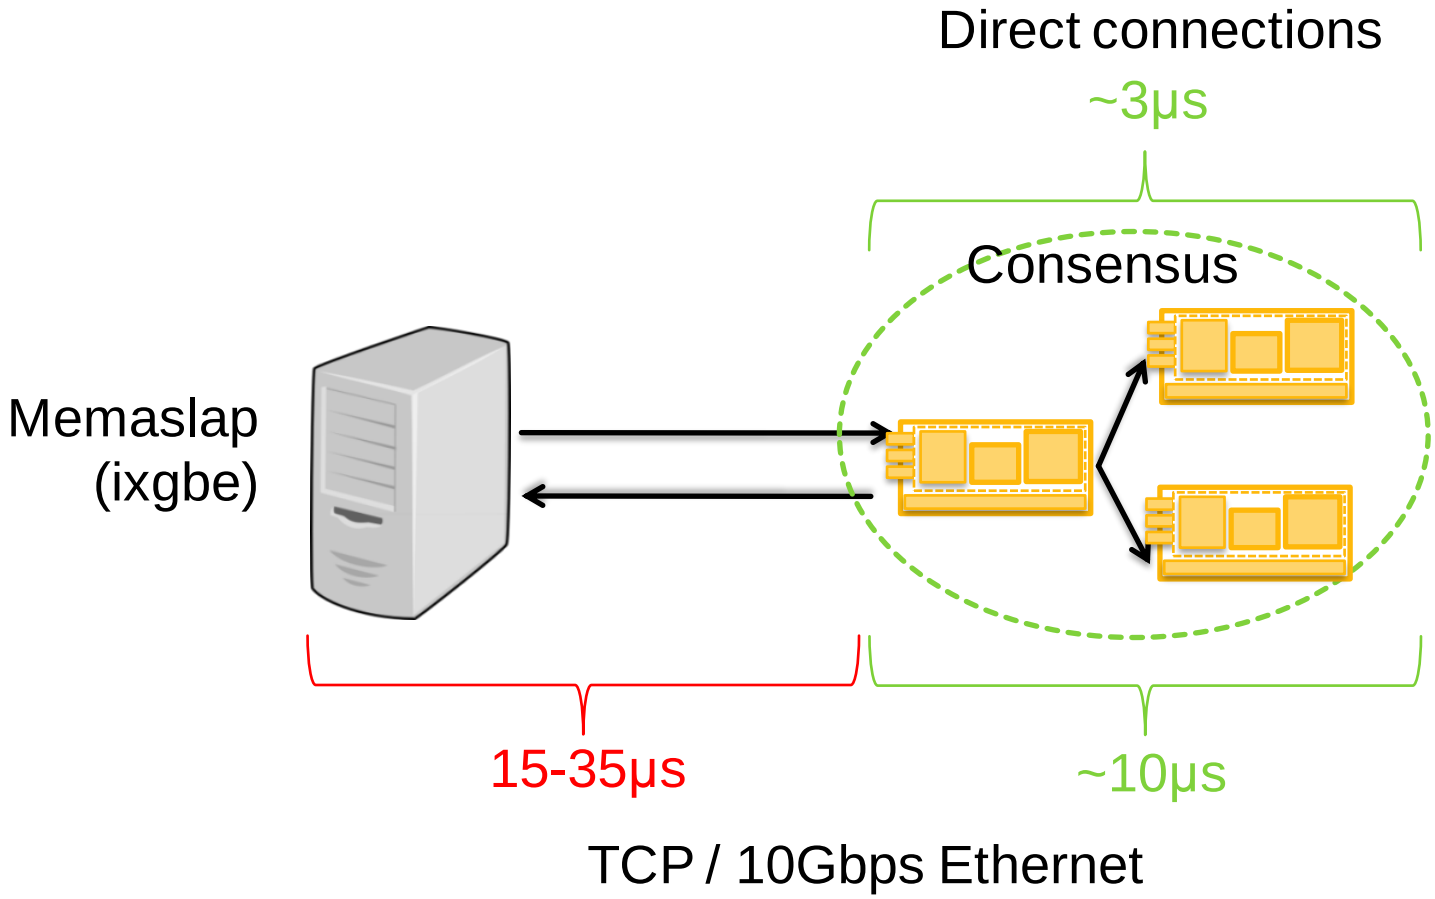

The benefit of specialization…

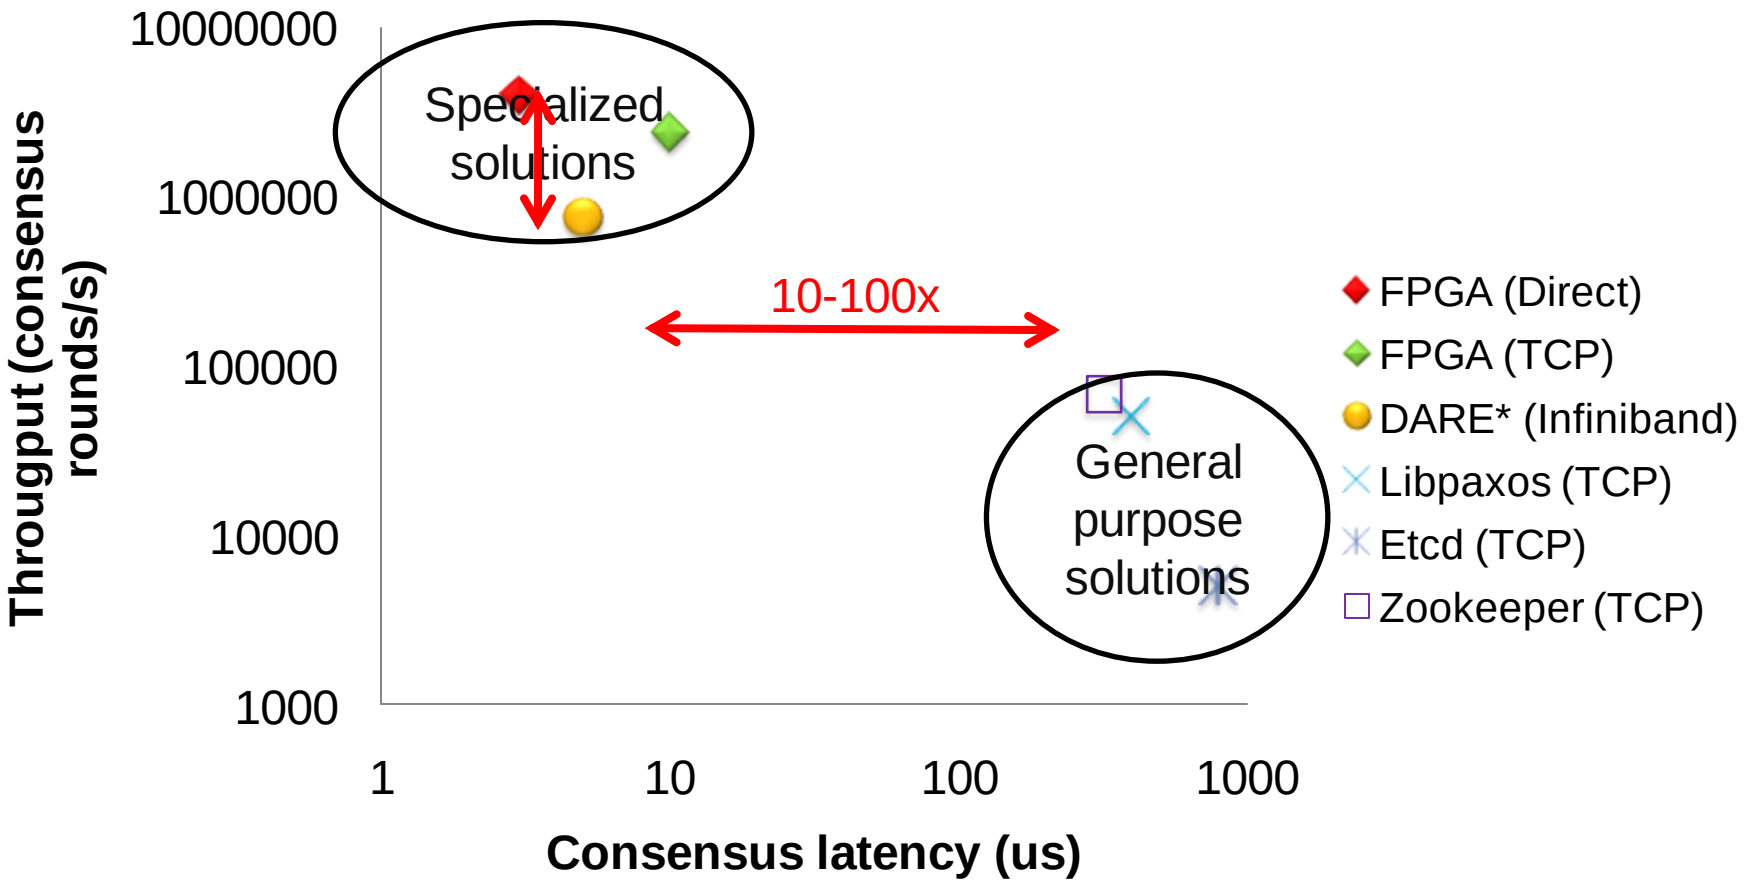

[1] Dragojevic et al. FaRM: Fast Remote Memory. In NSDI'14.
[2] Poke et al. DARE: High-Performance State Machine Replication on RDMA Networks. In HPDC'15.
*=We extrapolated from the 5 node setup for a 3 node setup and removed estimated client overhead.

Predictable hardware performance

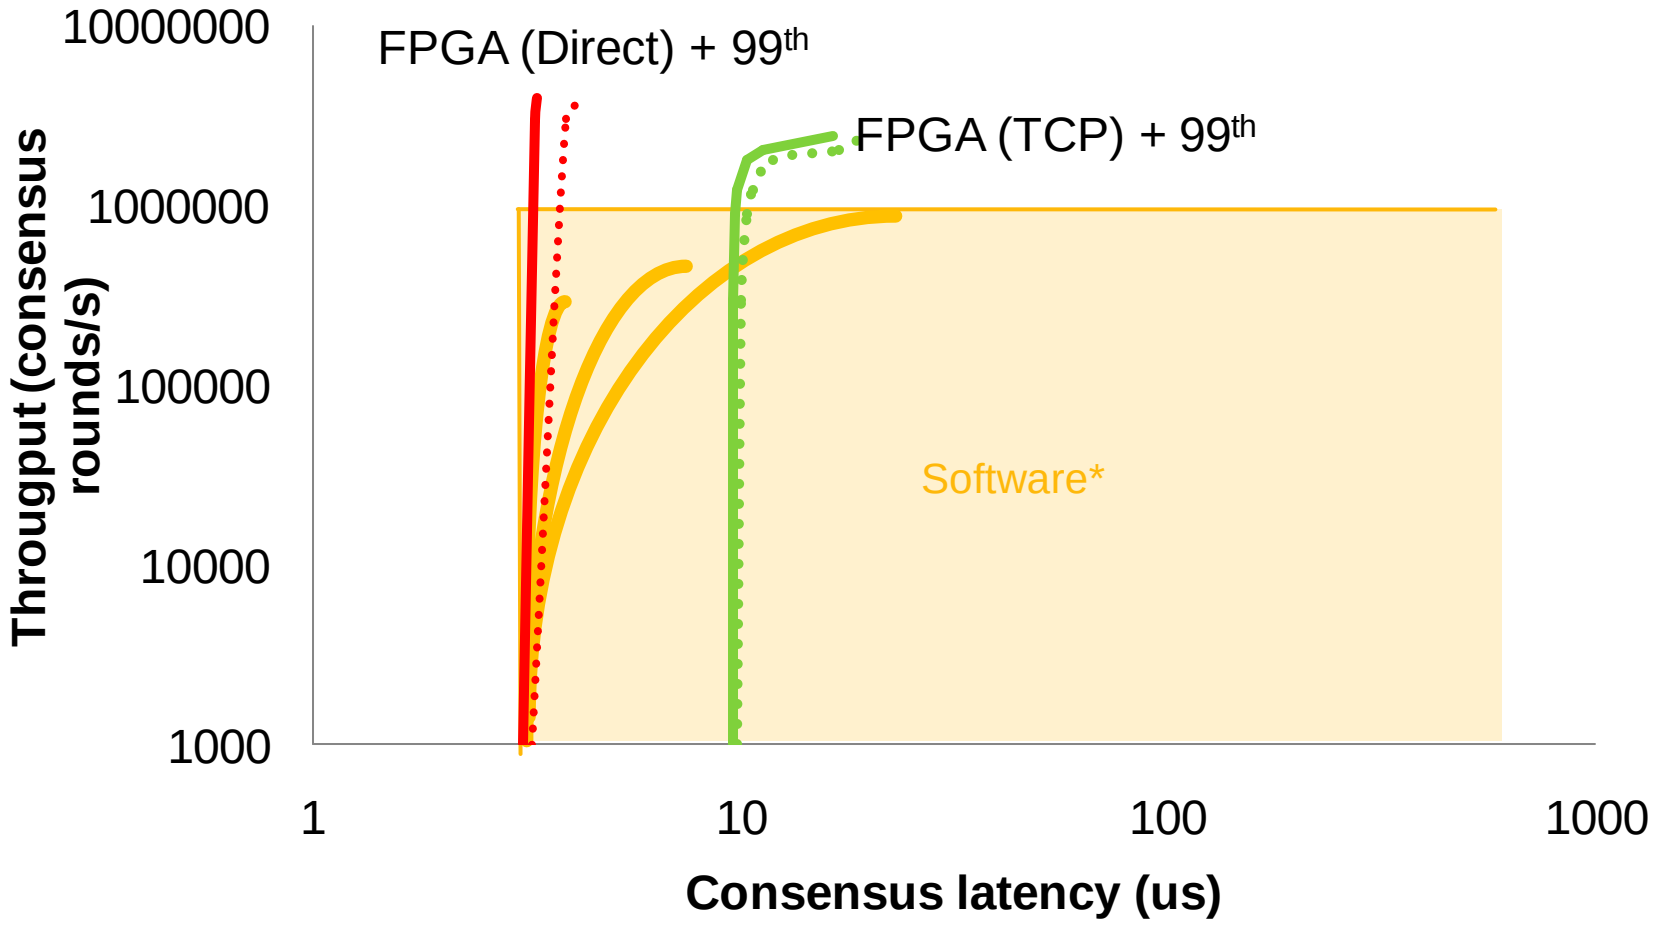

Throughput overview

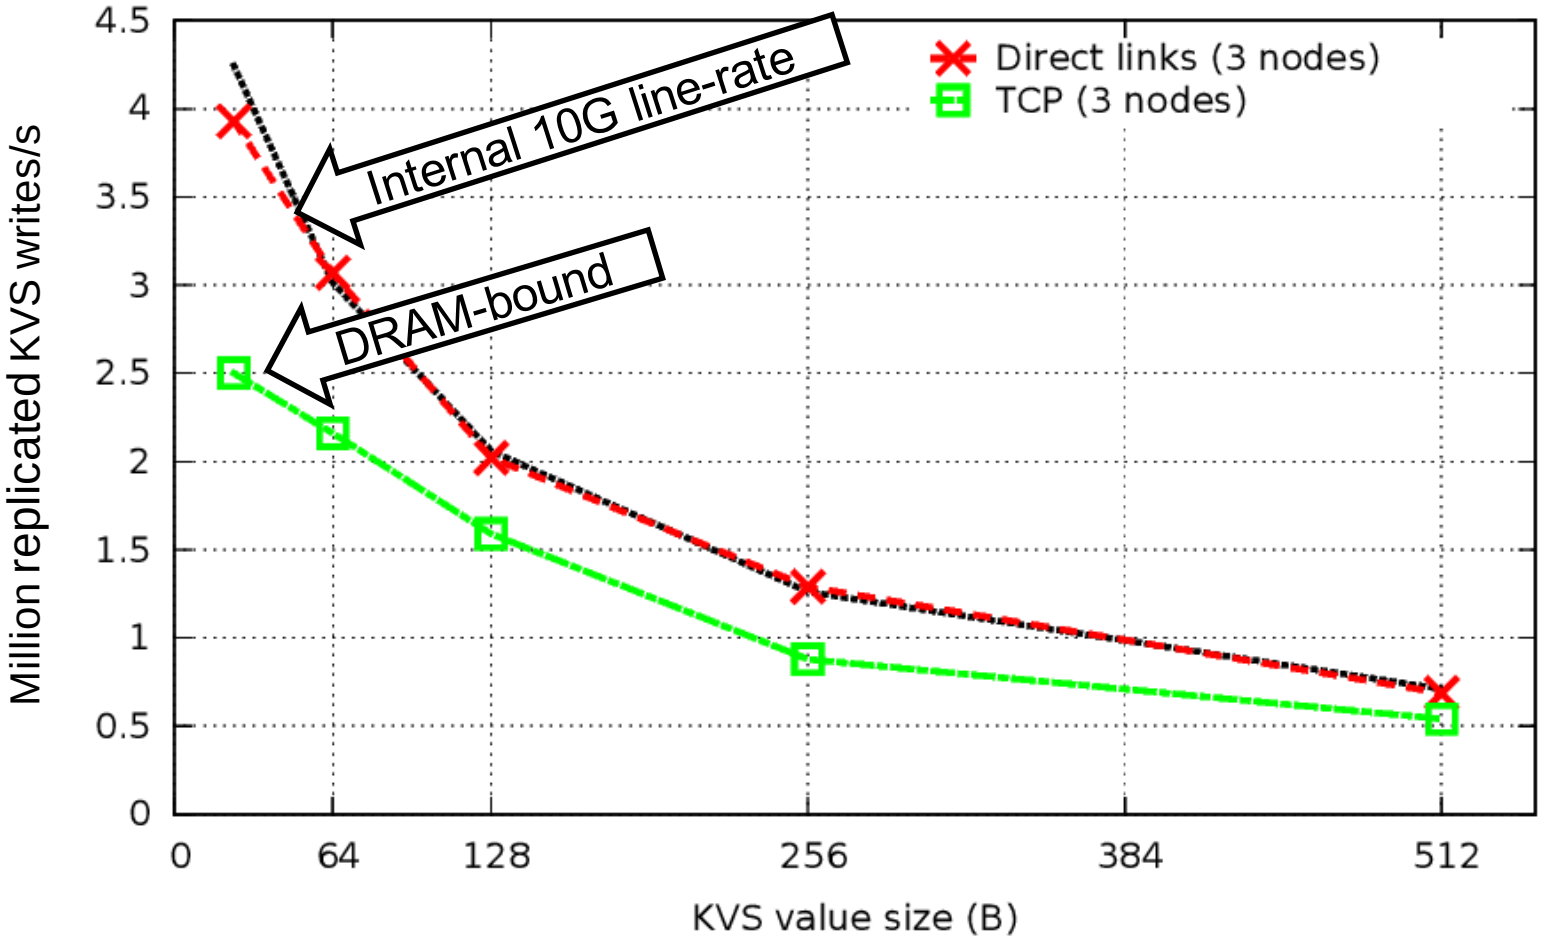

Conclusion

- Consensus is expensive but often necessary
  - Solution: specialization and tight integration with networking
- We built high-throughput, low-latency consensus in hardware

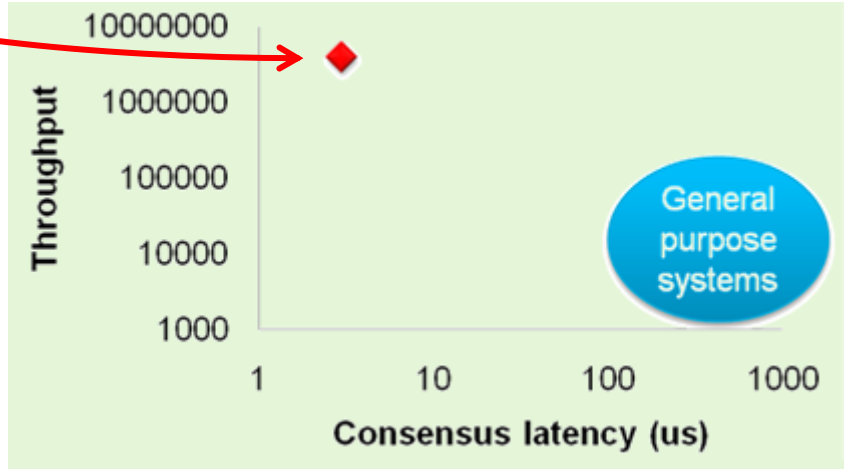

- Specialized hardware opens up new opportunities for smarter networks

{zsolt.istvan},{david.sidler}@inf.ethz.ch