

Making Speculative BFT Resilient with Trusted Monotonic Counters

Lachlan J. Gunn, Jian Liu, Bruno Vavala, N. Asokan

- Current speculative BFT protocols have a performance-resilience trade-off
- Trusted hardware provides an efficient and secure ordering mechanism
- SACZyzyva breaks the trade-off: full performance and full fault-tolerance simultaneously

Byzantine Fault Tolerance (BFT)

- Replication of deterministic state machines
- Replicated system looks like one state machine, despite compromised replicas

Zyzyva

- Speculative BFT:
 - Very **fast**: no waiting for **coordination**
- Very **simple** protocol when no faults occur
 1. Leader sends requests to replicas
 2. Replicas respond immediately
- **Any fault triggers non-speculative fallback**
- Zyzyva5 sacrifices robustness for speed

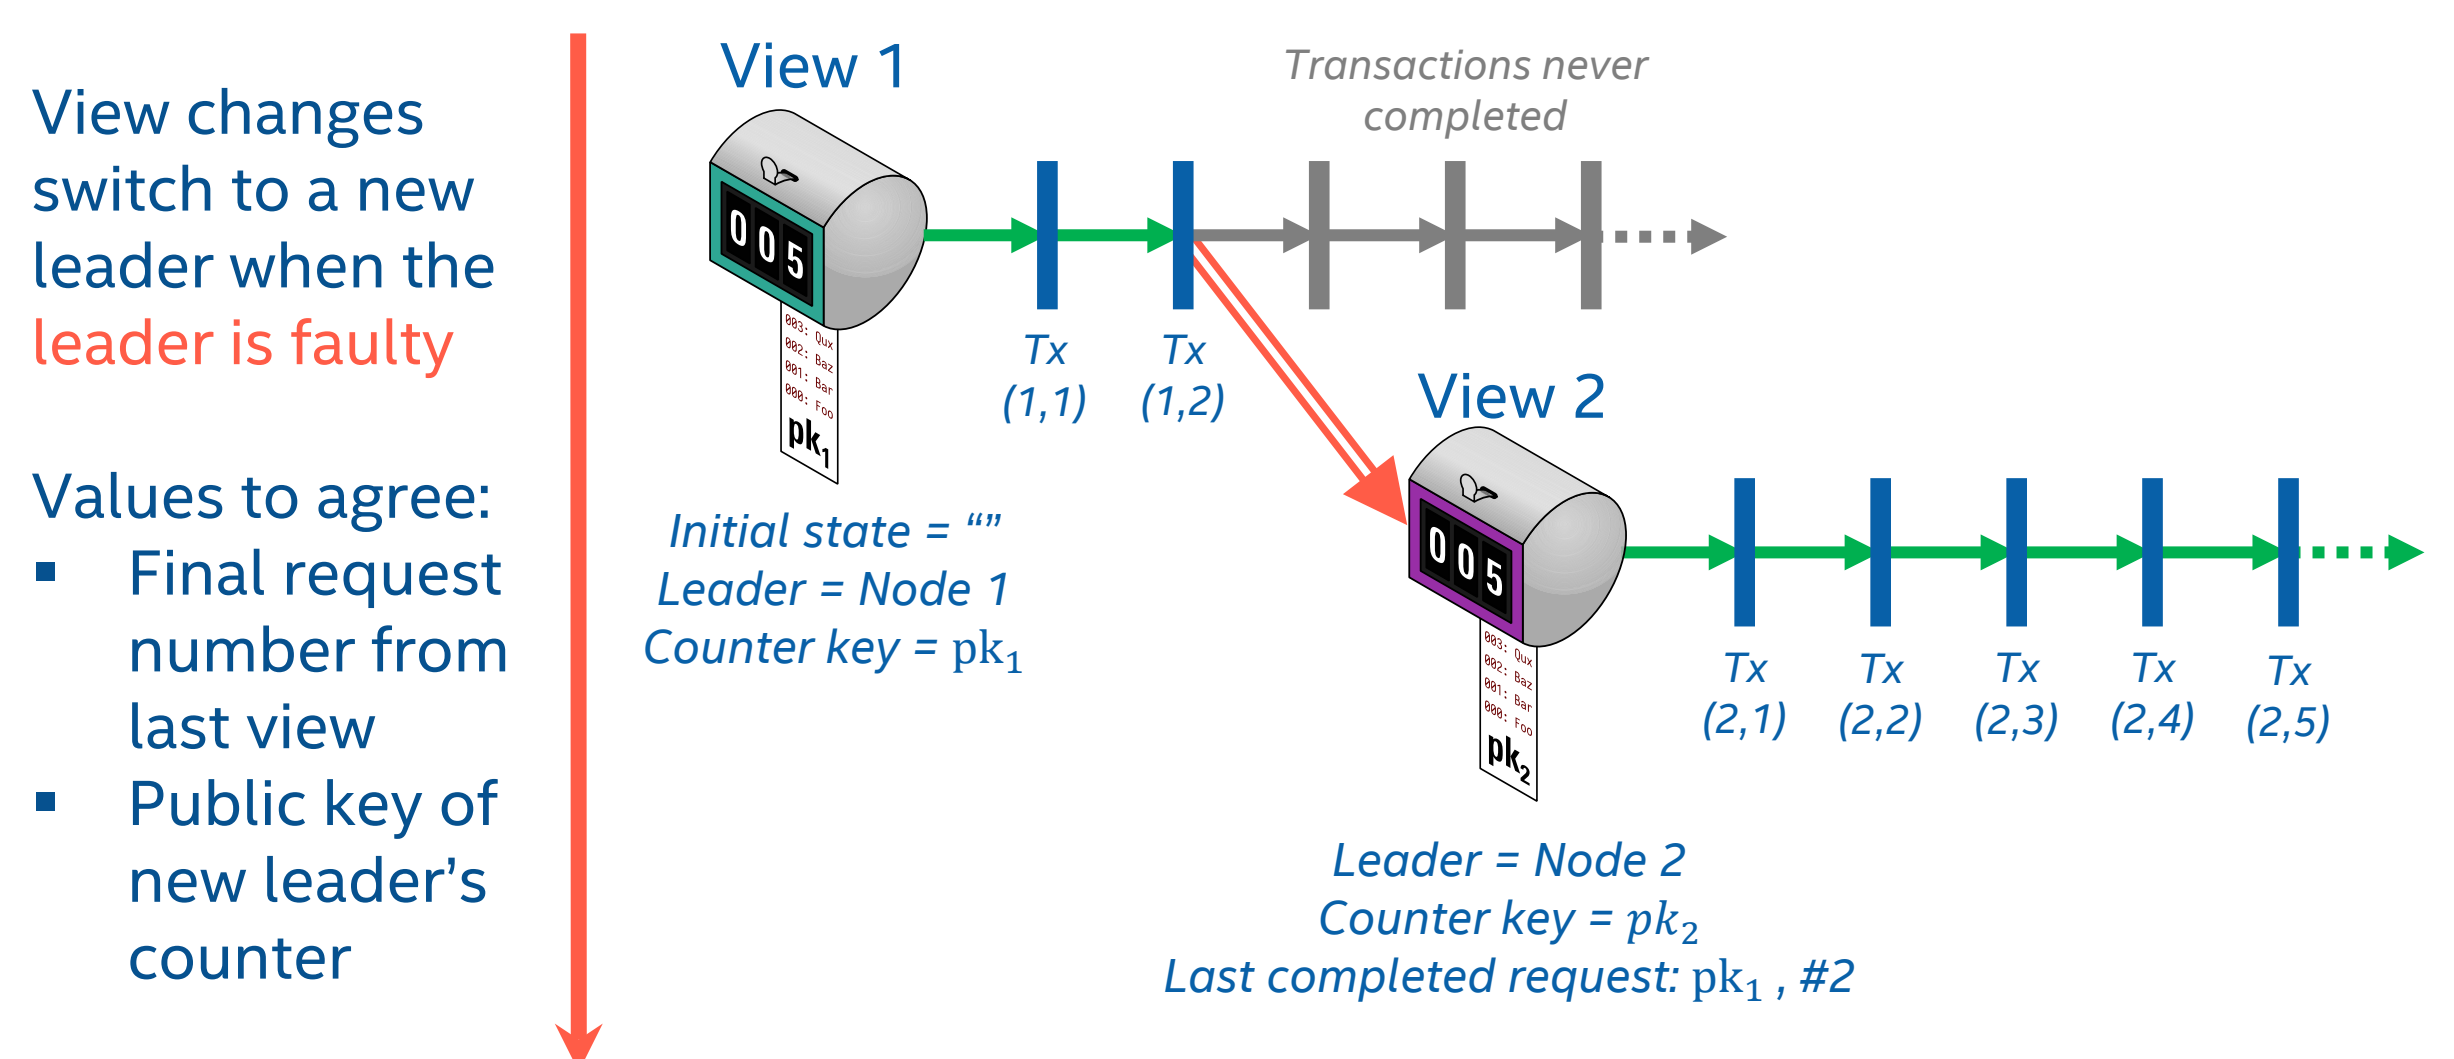
To tolerate f faults:

 - Zyzyva: $3f+1$ replicas, **slow after 1 fault**
 - Zyzyva5: $5f+1$ replicas, **never slow in fast path**
- Goal: $3f+1$ replicas, **never slow in fast path**

Single Active Counter Zyzyva (SACZyzyva)

- Leader orders requests and sends to replicas
- Ordering by trusted monotonic counter **in the leader only**
- **Never slow in fast path: only a faulty leader can slow progress**
- Same methodology applies to other protocols

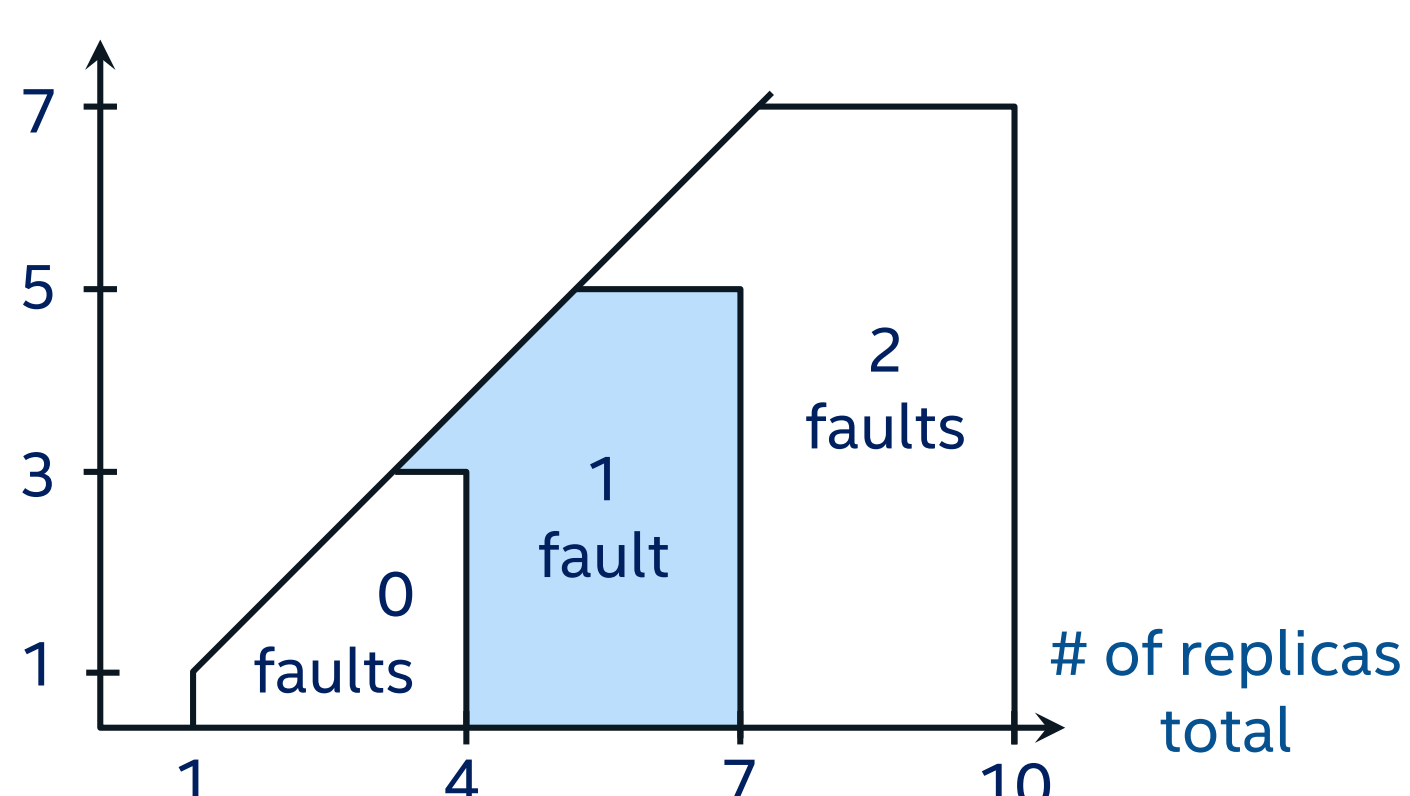
Ordering within a view guaranteed by trusted monotonic counter
Very fast, **always speculative**



Trusted hardware and BFT

- Trusted hardware can increase robustness
- Common primitive: monotonic counter
- **New result**: tolerating f faults **always** requires at least one of the following:
 - $2f+1$ replicas with trusted hardware
 - $3f+1$ replicas total

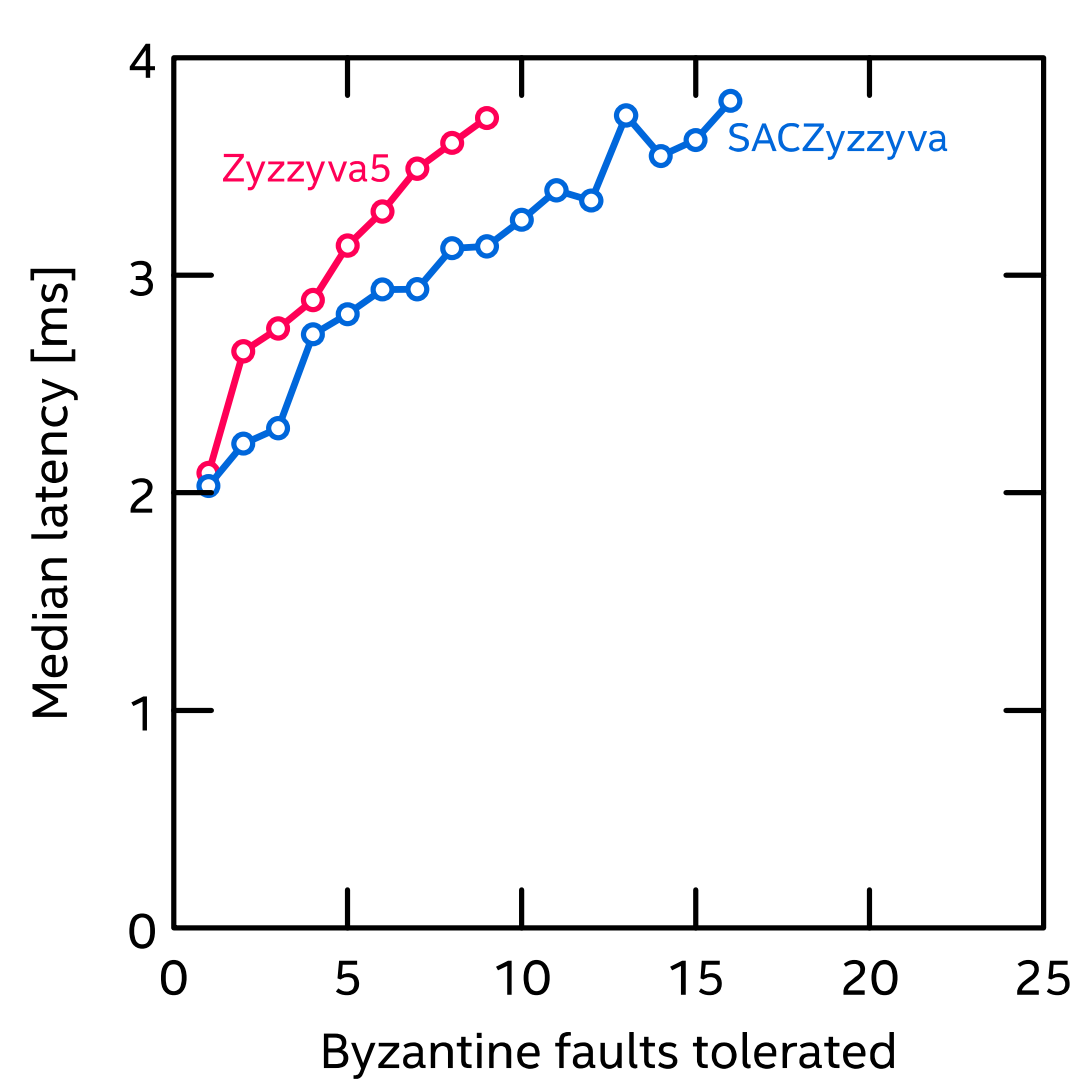
of replicas with trusted hardware



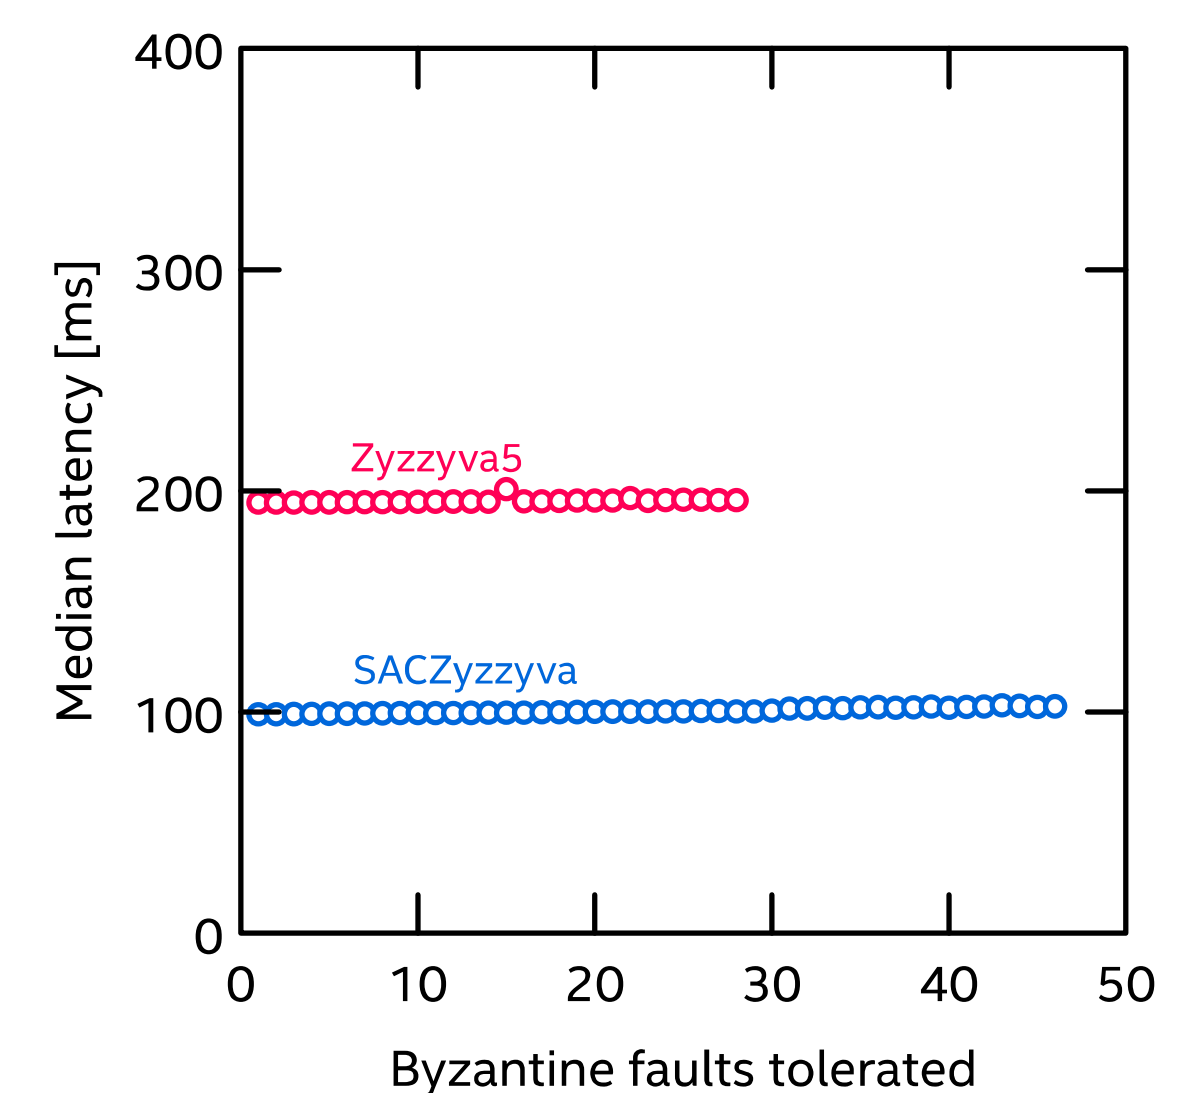
Theoretical limit of tolerable faults with partial availability of trusted hardware

Performance

- We outperform Zyzyva5 at the **same level of robustness**
- C++ Implementation of fault-free path for Zyzyva5 & SACZyzyva
- **Low- and high-latency** experiments using Amazon EC2
- Marginal latency increase for additional replicas: $<100\mu\text{s}/\text{replica}$



Latency with all replicas in one EC2 region



Latency with replicas spread across Frankfurt, Ohio, and Sydney regions