

Byzantine Replication Under Attack

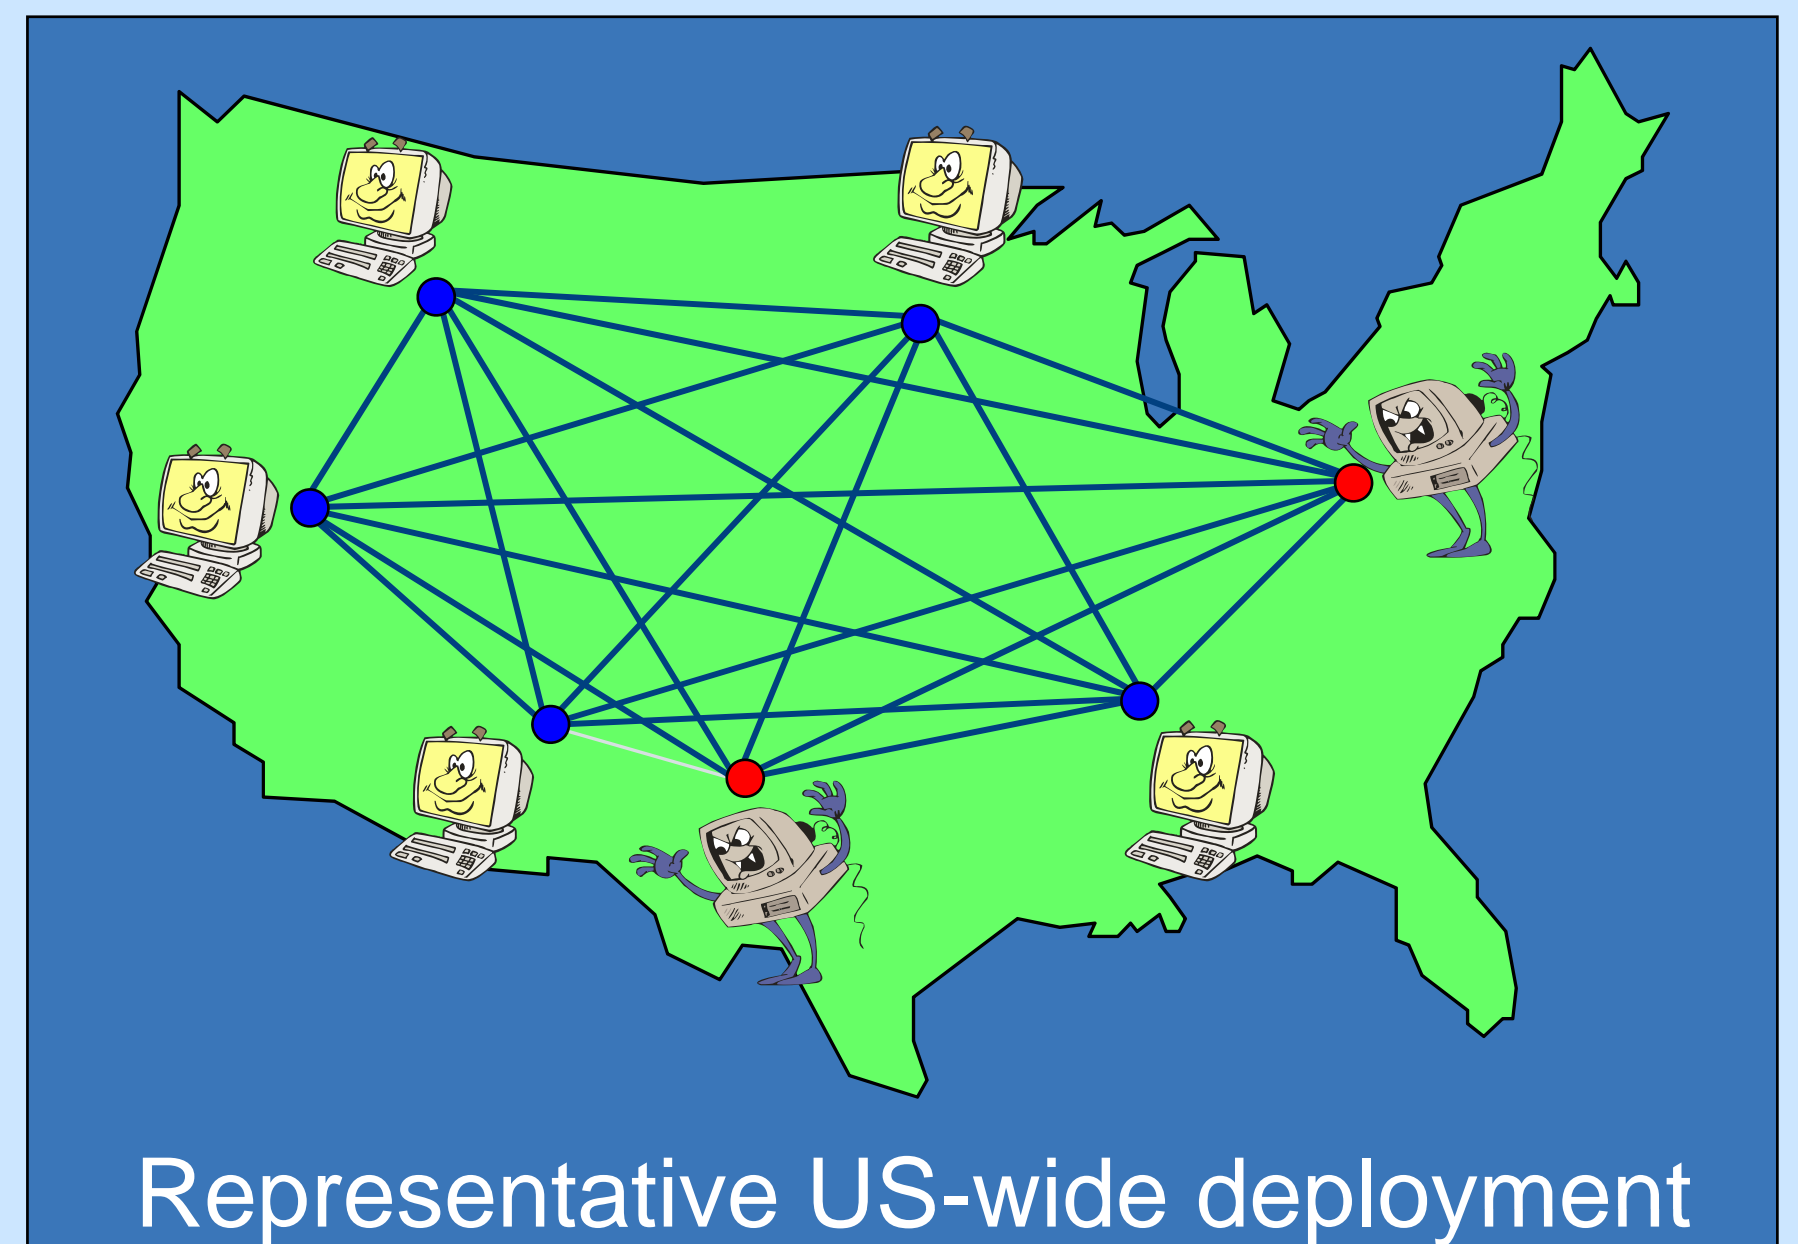
Yair Amir, Jonathan Kirsch, John Lane
 Johns Hopkins University (<http://www.dsn.jhu.edu>)

Problem

As network environments become increasingly hostile, even well-protected distributed information systems, constructed with security in mind, are likely to be compromised. An attacker can turn the infected machines into malicious entities that interfere with the system's functionality. **Byzantine replication** has emerged as a promising direction for addressing server compromises. While current state of the art protocols maintain consistency when under attack and perform well when there are no compromises, they are vulnerable to significant performance degradation when under attack, particularly by malicious servers that act slowly but without triggering defense mechanisms.

The Need for New Metrics

- Existing solutions satisfy **Safety** (consistency) and **Liveness** (eventual progress).
- Liveness is a necessary but insufficient correctness condition in the face of attacks, because a sophisticated adversary can make progress but at a very slow rate.
- We propose **average throughput** and **bounded delay** as alternative performance-oriented metrics.




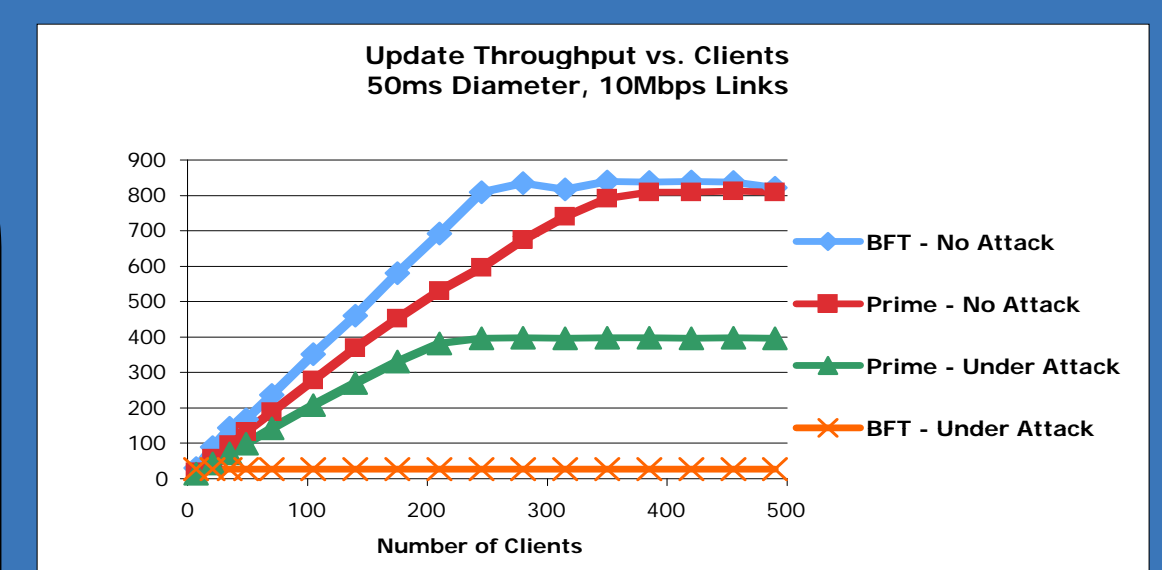
Comparison to State of the Art

Standard leader-based protocol (BFT)

- Achieves high performance in fault-free and benign-fault executions.
- Malicious leader can slow throughput down to one update per timeout in the worst case.
- Satisfies standard liveness criterion, where each update is eventually ordered but can have arbitrarily high latency, even if the network is stable.

Our approach- Prime

- Achieves comparable fault-free performance to BFT, and outperforms BFT by an order of magnitude under attack.
- Leader requires outgoing bandwidth independent of system throughput for ordering, enabling aggressive performance monitoring.
- Adapts threshold level of acceptable performance based on current network conditions.
- Achieves Bounded-Average-Delay criterion when network is stable.



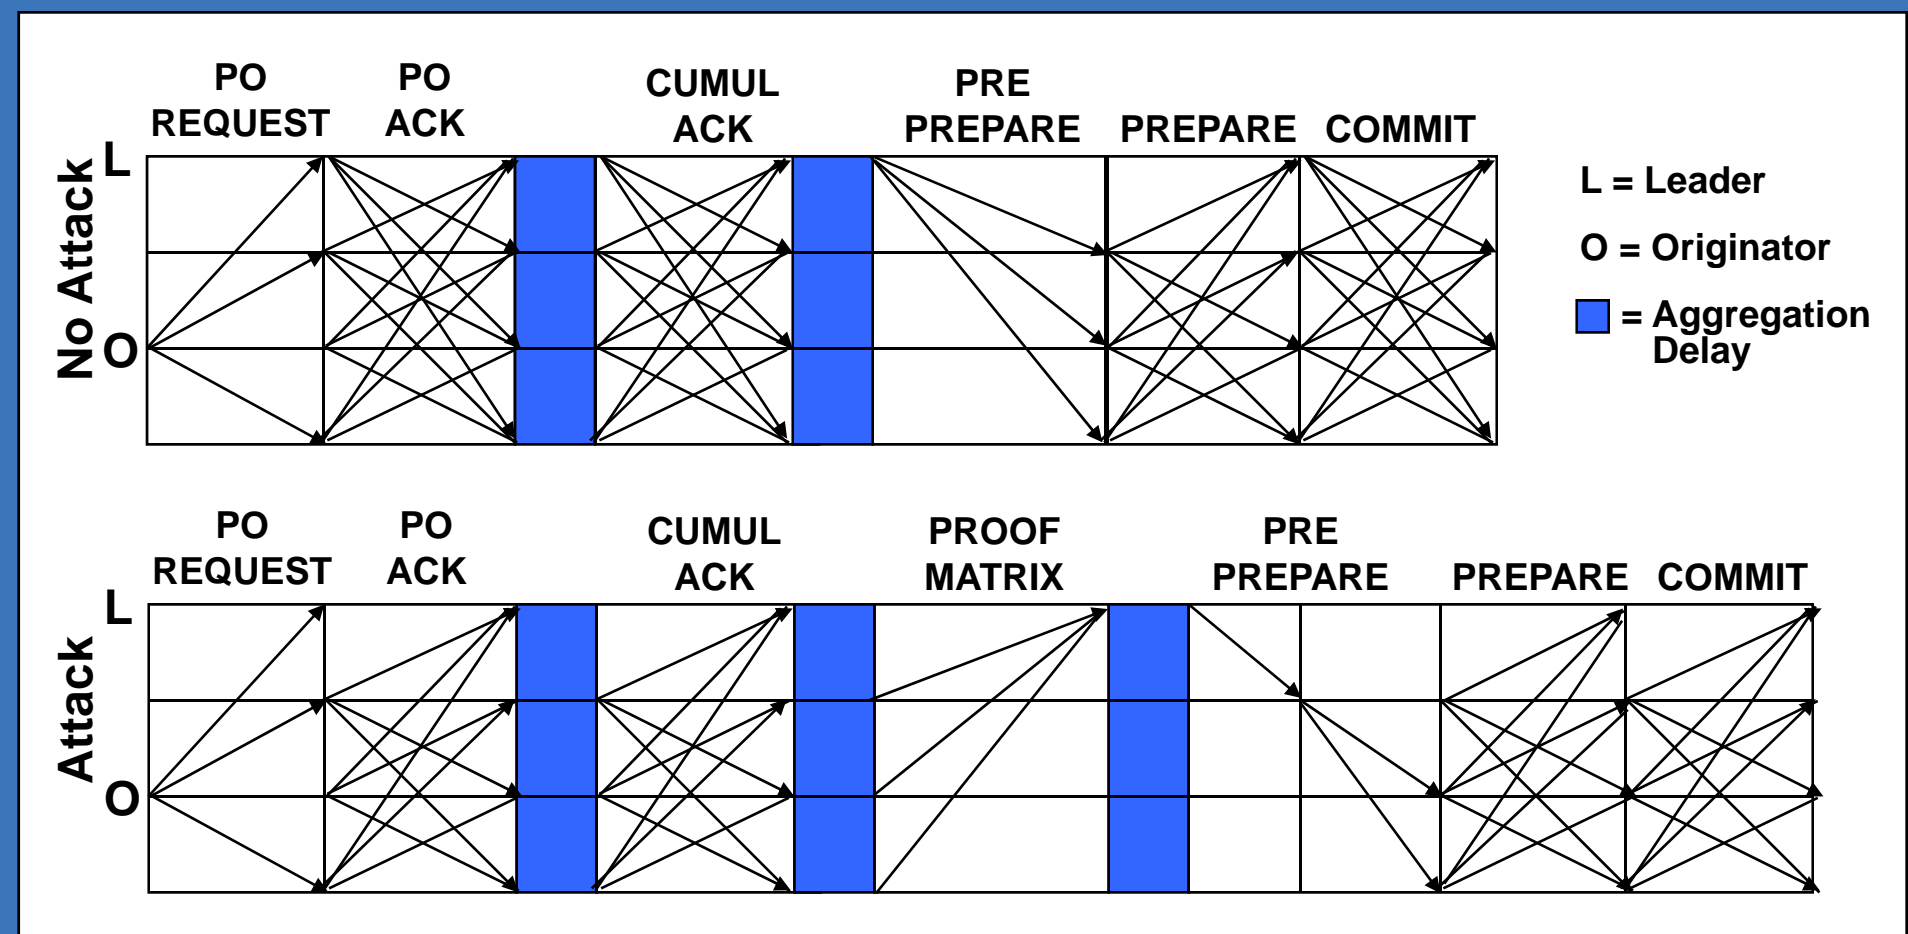
Prime: Performance-Oriented Replication in Malicious Environments

Pre-Order (PO) Phase

- Each server originates and disseminates updates from its local clients.
- Originator coordinates a protocol to locally order, or preorder, its updates.
- Cumulative acknowledgements enable aggregation of preorder information into small messages.

Ordering Phase

- Similar to BFT.
- Leader strings preordered updates into a global total order.
- Aggregation makes size of global ordering messages independent of system throughput.
- Companion protocol forces leader to send Pre-Prepare to at least one correct server in a timely manner to avoid being replaced.



In Prime, a malicious leader can add at most two rounds and one aggregation delay into the ordering path without being suspected.

Joint work with Dr. Brian Coan, Telcordia Technologies