

CSE 550: *Systems for all*

Au 2021

Ratul Mahajan

Distributed transactions

Atomic update of data across multiple nodes

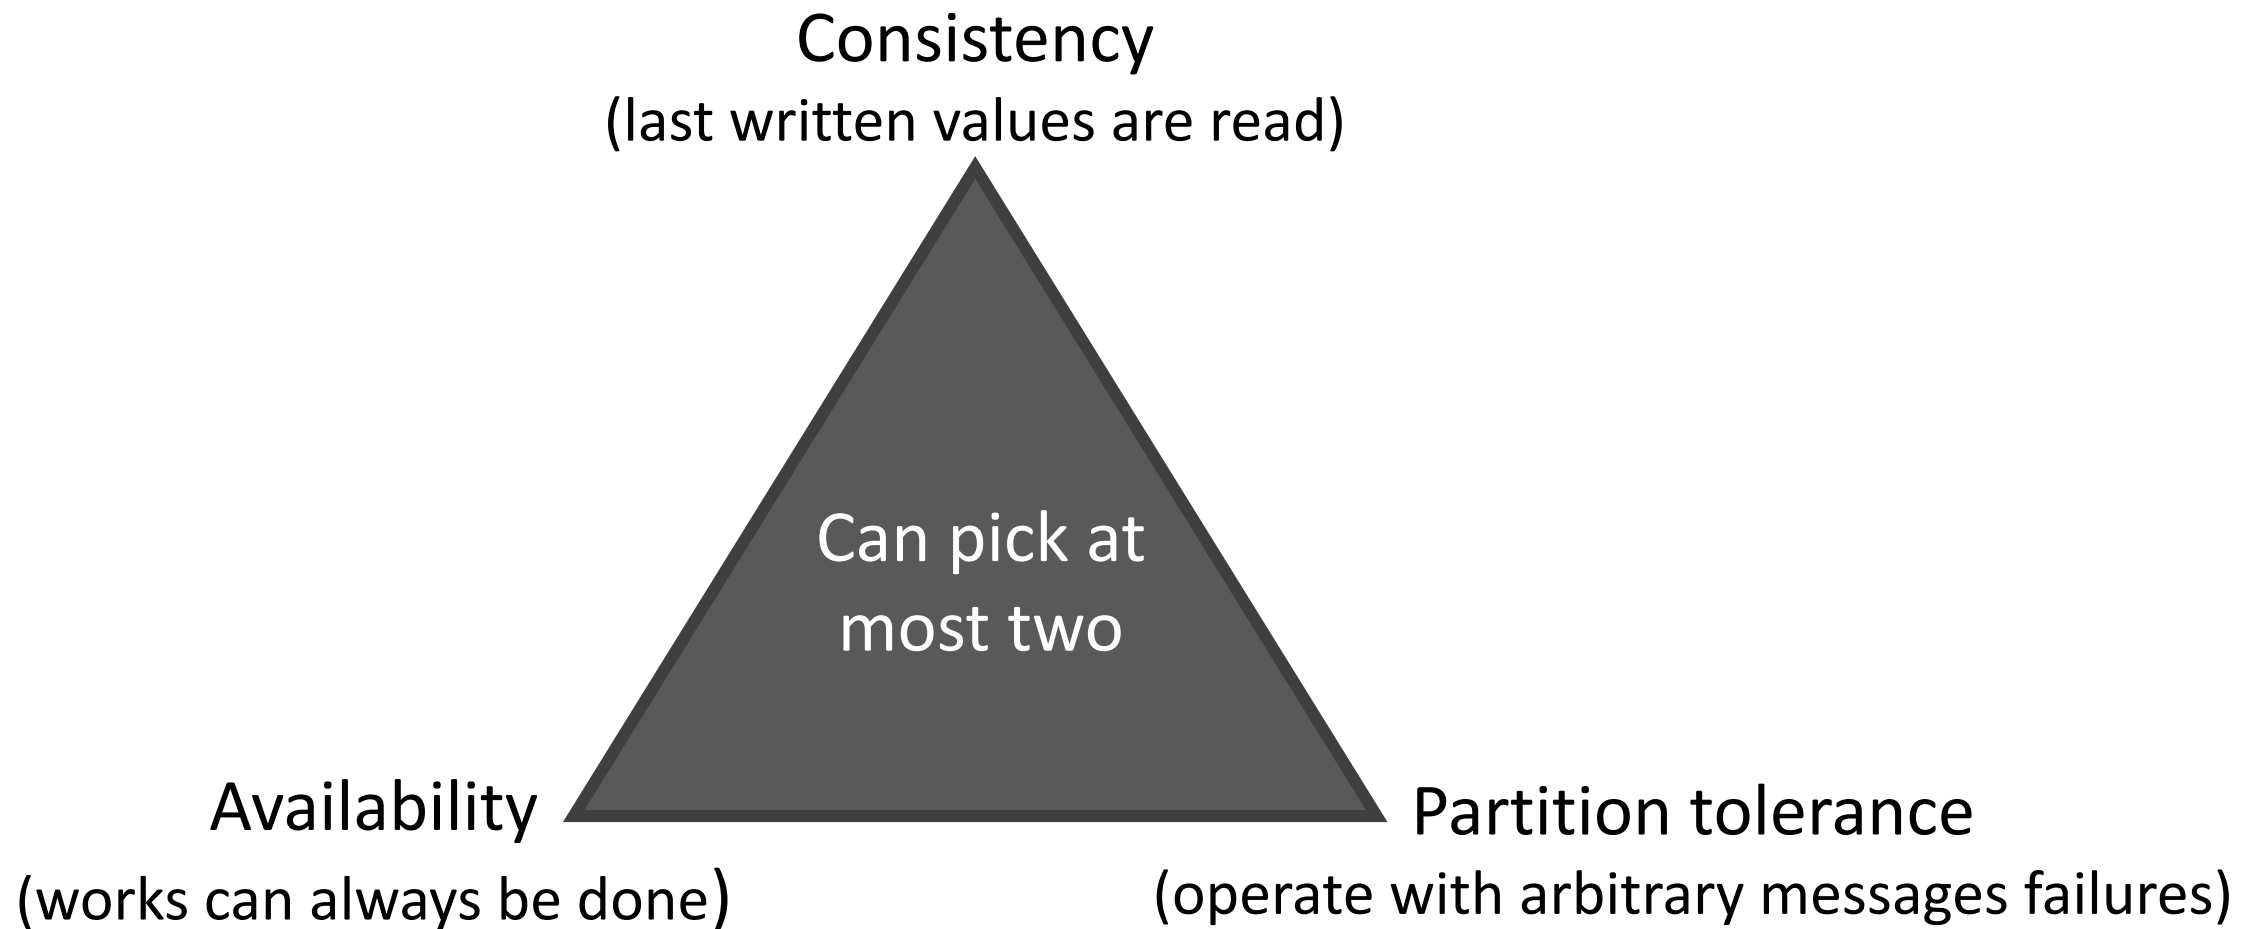
Why is the data distributed?

- Replication: Tolerate failures
- No replication: Need more compute / storage

Fault tolerance in distributed systems

	Replication	Non-replication
Fail stop	Paxos	2PC, 3PC
Byzantine	PBFT	N/A

CAP theorem: Fundamental trade-off for fault tolerance



Over to Diya and Peter

Next few weeks

Learned many of the fundamental techniques thus far

- Sharing resources, ordering events, handling failures, building large networks

Will focus on applications

- Different types data stores (files, KV stores, structured data)
- Different types of compute logic (streams, dataflow, ML)

Some fundamentals still coming

- Security, correctness reasoning