

Distributed consensus and fault tolerance - Lecture 1

Tuesday 7 March 2017 10:00 (1 hour)

In a world where clusters with thousands of nodes are becoming commonplace, we are often faced with the task of having them coordinate and share state. As the number of machines goes up, so does the probability that something goes wrong: a node could temporarily lose connectivity, crash because of some race condition, or have its hard drive fail.

What are the challenges when designing fault-tolerant distributed systems, where a cluster is able to survive the loss of individual nodes? In this lecture, we will discuss some basics on this topic (consistency models, CAP theorem, failure modes, byzantine faults), detail the raft consensus algorithm, and showcase an interesting example of a highly resilient distributed system, bitcoin.

Presenter: BITZES, Georgios (CERN)