15-440 Recitation 7: The CAP Theorem

Vijay Vasudevan

Outline

- Questions on Part 3?
- The CAP Theorem

Acids and Bases

- Classical Distributed systems: ACID semantics
  - **A**tomic
  - **C**onsistent
  - **I**solated
  - **D**urable

- Modern Internet systems: BASE semantics
  - **B**asically **A**vailable
  - **S**oft-state
  - **E**ventually consistent

Consistency matters

- (example)

Availability matters

- Amazon being down for 1 minute = loss of $10,000 in revenue

Partitions happen

- In the Internet
  - (prefix hijacking, fiber line cuts, etc.)
- Even in datacenters
  - Google, Amazon report network partitions

Can we have it all?

- Strong Consistency
- High Availability
- Partition Tolerance

The CAP Conjecture

- Eric Brewer at Berkeley made the CAP conjecture in 2000

- Conjecture: you can have only two of the three properties

Consistency and Availability

- Examples include:
  - Tightly coupled clusters, databases
- 2PC, cache consistency protocols abound

Consistency and Partition Tolerance

- Distributed databases, locking, quorums
- Techniques:
  - Strict locking
  - Partitions with minority don't work

Availability and Partition Tolerance

- Design of many internet services
- Examples: DNS, caches, large-scale key-value stores
- Use TTL/state for eventual consistency, out-of-band conflict resolution

CAP Theorem

- Proven in 2002 by Seth Gilbert, Nancy Lynch

- We'll go over a part of the proof

In reality

- Systems are a neutral mix of ACID and BASE
- ACID *required* for monetary transactions
- BASE for things like shopping cart, DNS