

Conflict Free Replicated Data-types in Eventually Consistent Systems

Joel Jacobson, Basho Technologies

@joeljacobson



Riak

Distributed, Masterless, Key/Value Database
+ Extras

Key	Value
Key	Value

Buckets, Keys and Values

Simple operations;
GET, PUT, DELETE

Distributed & Scalable

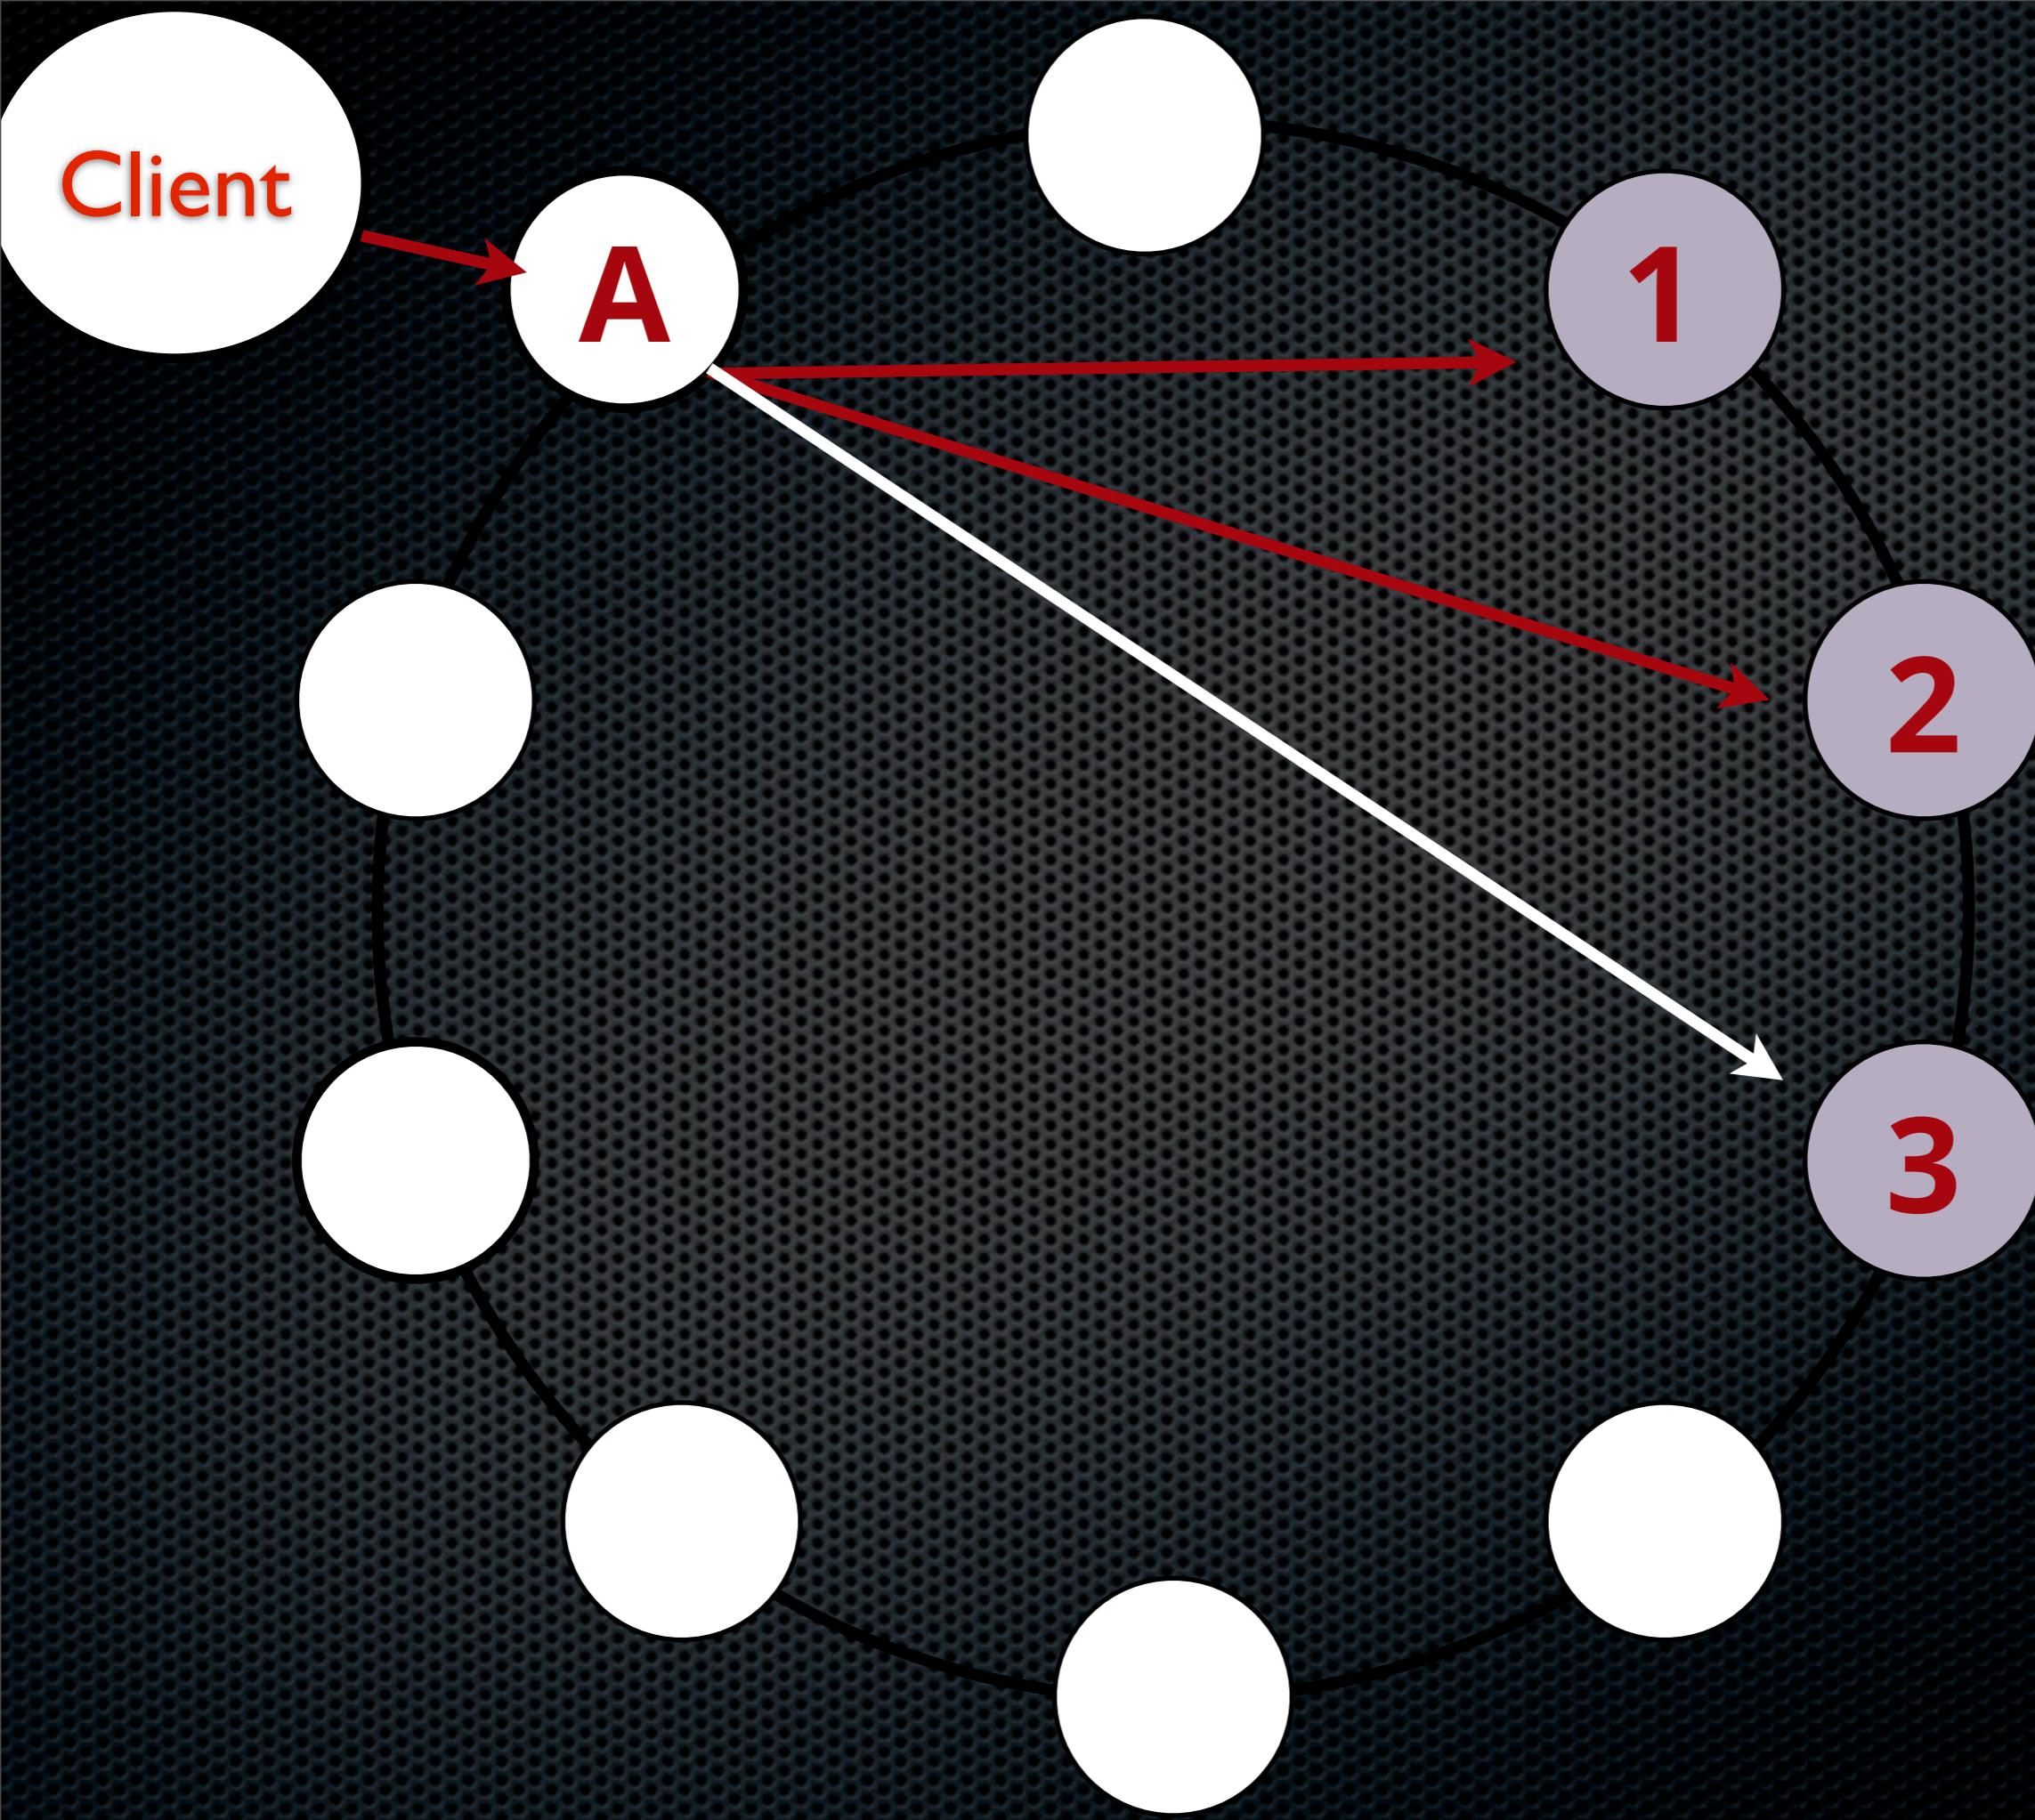


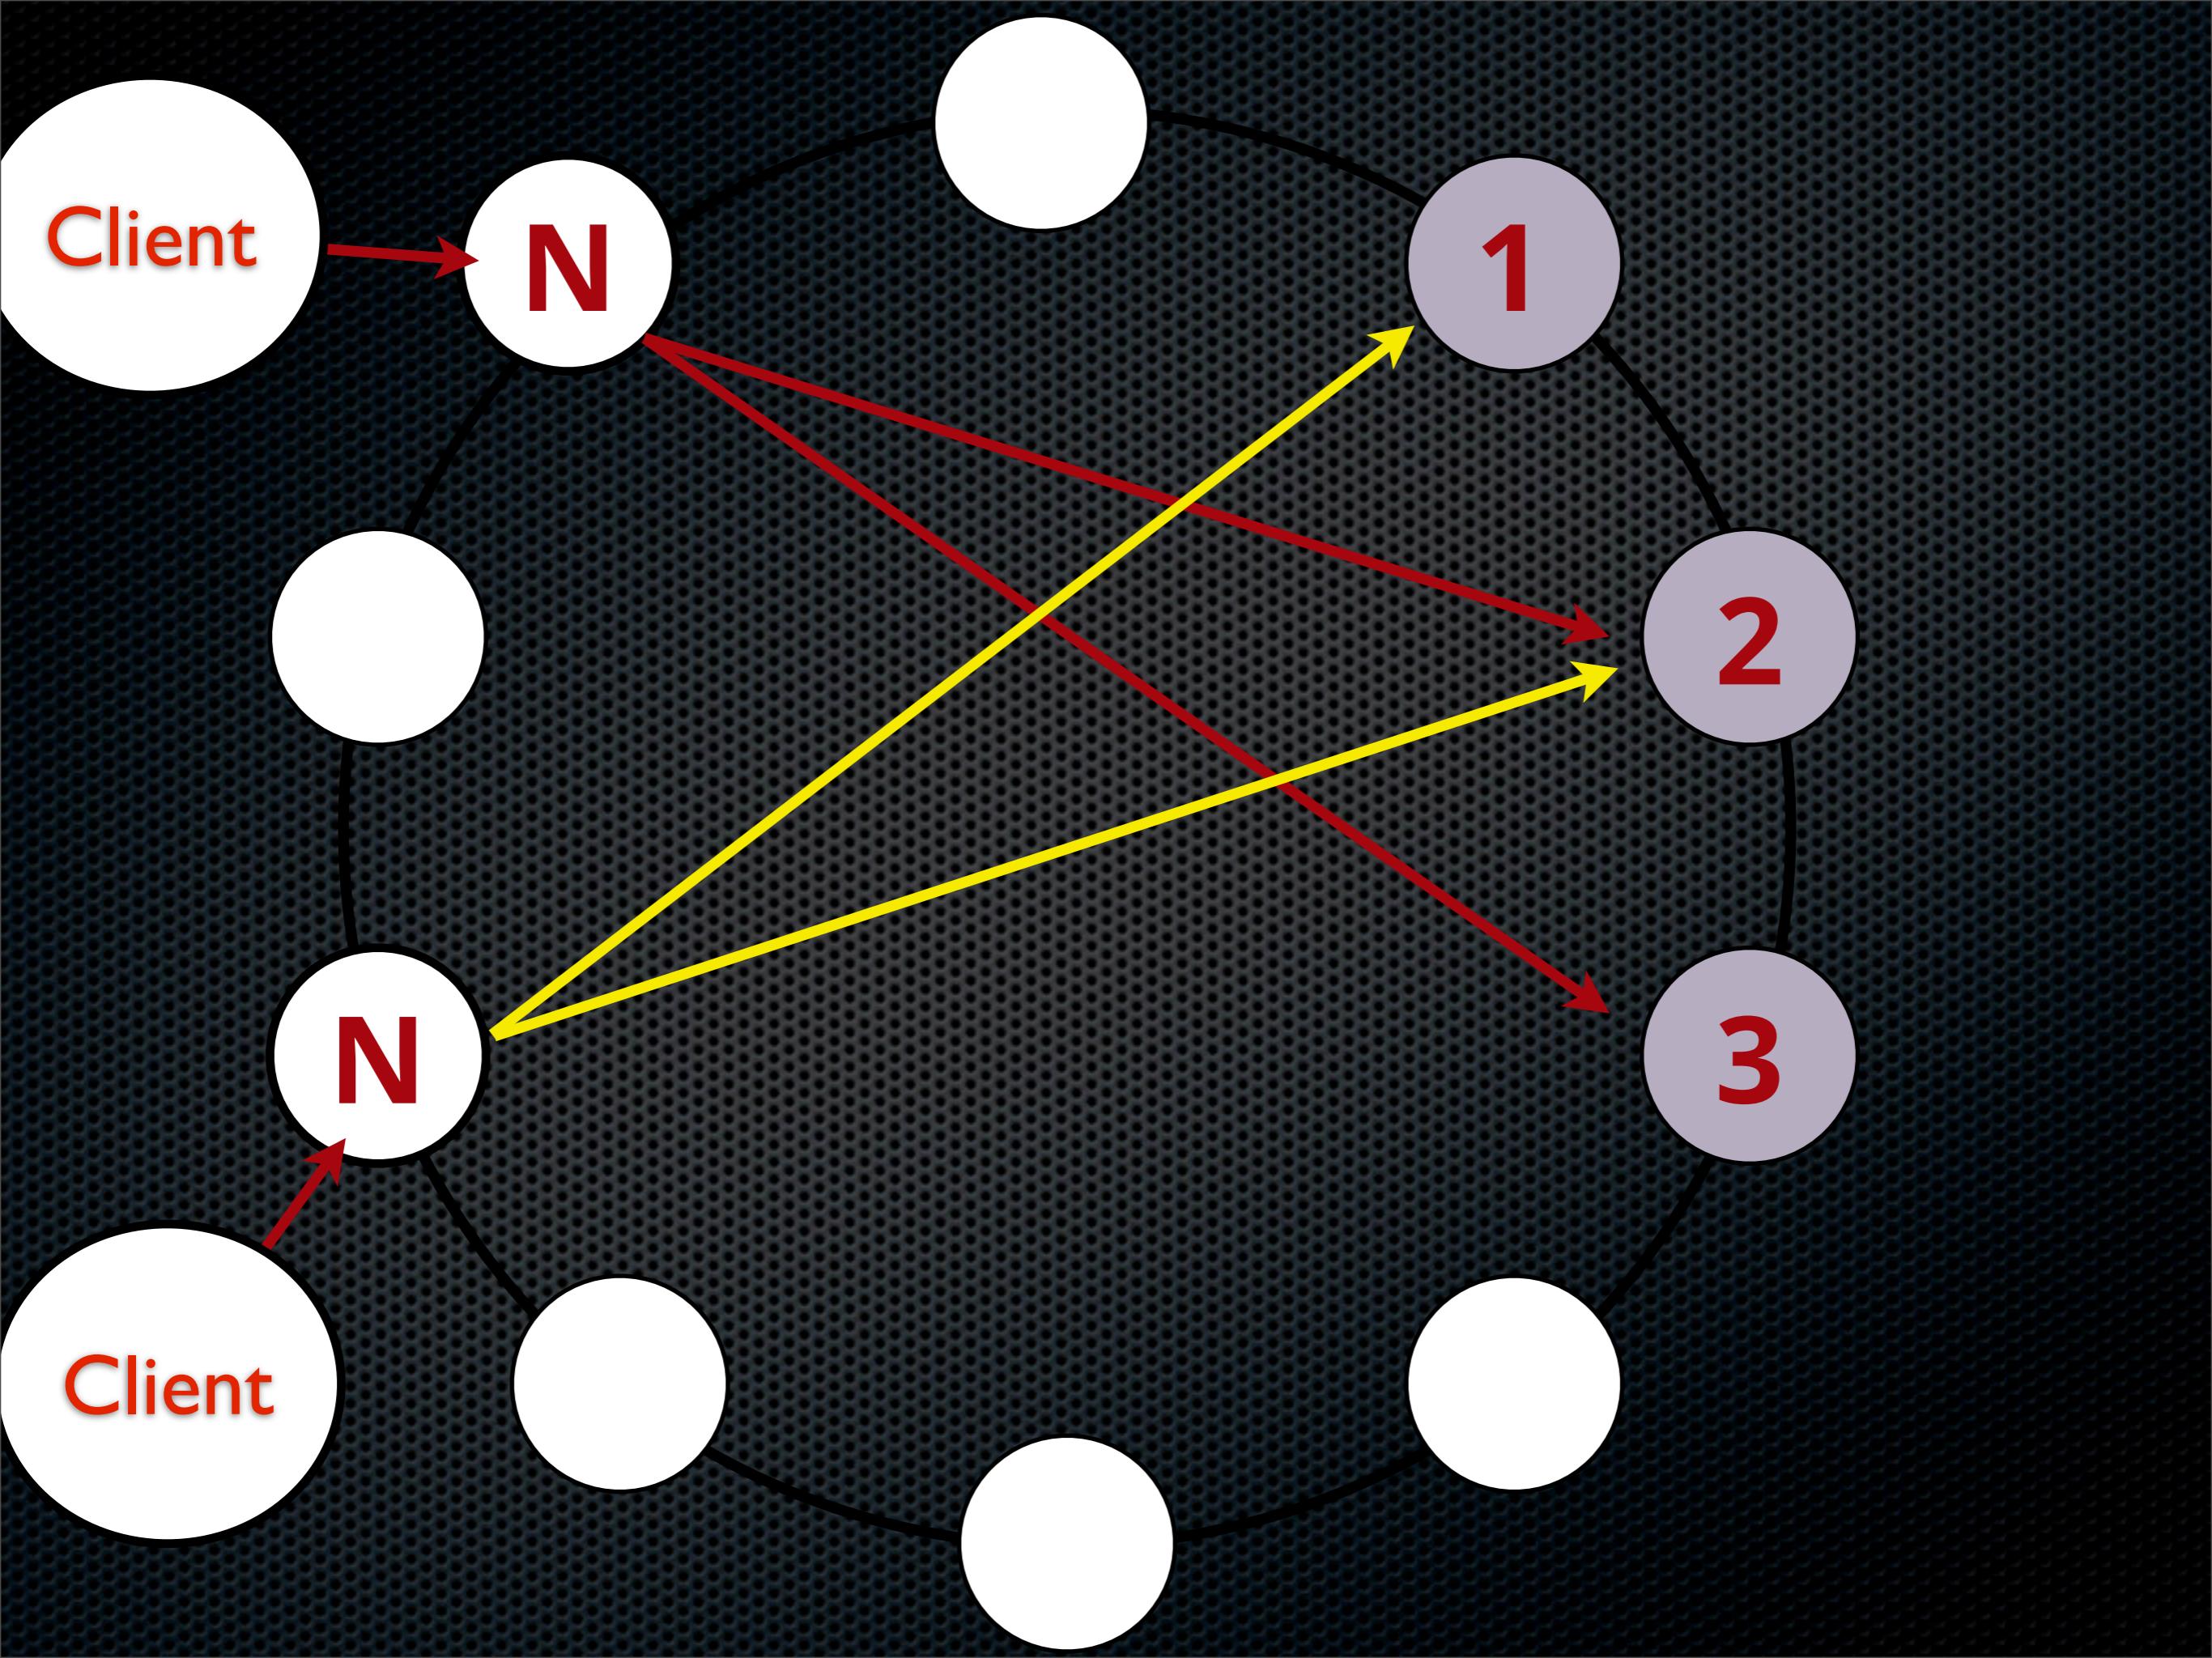
Fault Tolerance

- Any node can serve R/W requests
- Data is replicated
- Fallback Node
- Hinted Handoff

Dynamo Systems

- Distributed
- High Availability
- Masterless





Conflicts

Conflict Resolution

- Last-Write Wins (LWW)
- Vector Clocks (Siblings)

Sibling Resolution

Conflict Free Replicated Data Types



INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET EN AUTOMATIQUE

CRDTs: Consistency without concurrency control

Mihai Letia — Nuno Preguiça — Marc Shapiro

[http://hal.upmc.fr/docs/00/55/55/88/PDF/
techreport.pdf](http://hal.upmc.fr/docs/00/55/55/88/PDF/techreport.pdf)

Primitives

- Counters
- Sets
- Maps

Counters

Sets

Maps

The Problem of Absence

Use-Cases

- Web page visits (G-Counter)
- Shopping Cart (Modified OR-Set)
- “Like” button (U-Set)
- Maps can be made up of multiple CRDTs

Common Questions

- What do they cost me?
- How big are they?
- What can't they do?

Thank you

joel@basho.com