Byzantine Faults in a Rational World

Amitanand Aiyer, Allen Clement, Jean-Philippe Martin, Carl Porth, Mike Dahlin and Lorenzo Alvisi

LASR
UT Austin

Two motivating observations

- Dependability more pressing need than performance
- Distributed systems increasingly span multiple administrative domains

How should nodes be modeled?

Traditionally, a node is modeled as either::

- *Correct:* the node follows its specification
- *Faulty:* the node deviates from its specification
  - benign
  - Byzantine

A new classification

A node is either:

- *Altruistic*: the node follows the assigned protocol
- *Rational*: the node is not malicious, but will deviate from the assigned protocol to maximize its benefits and minimize its costs
- *Byzantine*: the node deviates from assigned protocol even when not "in its interest" because of malfunction, misconfiguration, or malice

Nodes may be subject to benign faults

Our goal

Develop the theory and practice of building distributed systems that tolerate both rational and Byzantine behavior

Our approach

- Adapt low-level BFT primitives (state machine replication, quorum replication, reliable broadcast) to tolerate rational behavior
  - create suite of building blocks
  - avoid ad-hoc reasoning for each application
- Develop end-to-end BRFT applications on top of these primitives
  - challenge: integrate low-level BRFT mechanism with end-to-end incentive structure of the application

Our assumptions

1. Byzantine nodes are few, but no bounds on the number of rational nodes
2. Cost: bandwidth, storage, computation, power, etc.
3. Long term repeated interactions
   - only way to achieve equilibrium in Prisoner's dilemma
4. Strong identities and restricted membership
   - prevent Sybil attack
   - enable *internal* and *external* disincentives to deter misbehavior
   - reasonable for our target applications

Our target application

Peer-to-peer backup system

- stresses BRFT in multiple dimensions
  - multiple resources integration
  - requires achieving BRFT at different timescales
  - range of provisioning may require to break simple symmetry between pairs of nodes
  - applicable to deployment scenarios with different trust models
- useful!
  - lab, dorm, Box Populi

Incentive compatible backups

- System links storage available to a node with storage contributed by the node
- To enforce quotas
  - peers publish signed lists of teh data they store and of the data that is stored on their behalf
  - *receipts* used to detect and prove lies
  - *witnesses* provide incentives against "passive-aggressive" nodes
  - witnesses implemented as BRFT replicated state machines

Status: protocols

- Studied two protocols:
  1. Lamport's Byzantine agreement with using unforgeable signatures
  2. Srikanth and Toueg's Byzantine agreement without signature s (a.k.a. the *echo* protocol)
- Proved both protocols are vulnerable to the "tragedy of the commons"
- Derived and proved incentive compatible versions of these protocols
- Working on BRFT state machine replication

Status: application

- Authors are trusting their iTunes library (or whatever else is vital to them) to initial prototype
- On schedule for lab-wide deployment in 2 weeks (about 20 users)
- Working on dorm deployment in 6 weeks