TACT: Tunable Availability and Consistency Tradeoffs for Replicated Internet Services

Haifeng Yu and Amin Vahdat

{yhf, vahdat}@cs.duke.edu

Computer Science Department

Duke University

Durham, NC 27708

Dynamic Replication in the Internet

- Replication for increased availability and performance
  - Internet makes the need real
- Static content replication: We know the answers
  - Static web pages
  - Images
  - Audio/Video
- Dynamic content replication? Still hard after decades of research
  - Airline reservation system
  - Online bookstore
  - Stock-trading system
- **Key question: How to get consistency without sacrificing availability?**

The World Today: A Switch

- Two choices

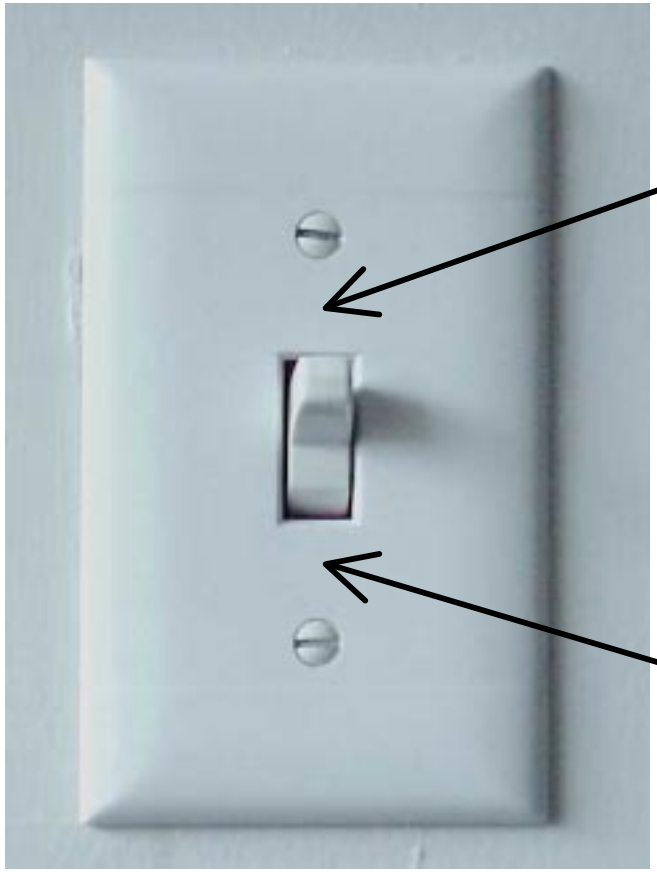

one-copy availability,
weak consistency
(Bayou, Coda, Ficus, etc)

strong consistency,
reduced availability
(replicated databases)

A Better Tomorrow: A Knob

- TACT: **continuously tunable tradeoffs**

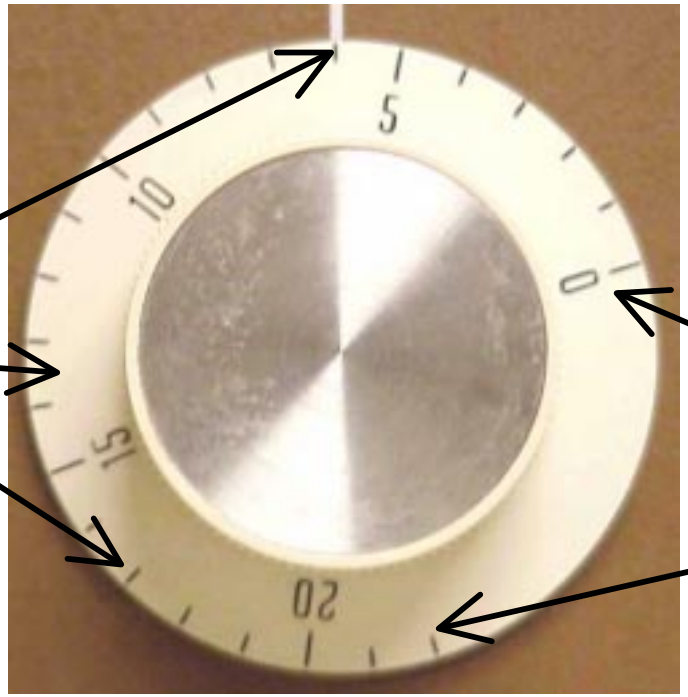

medium-level availability, medium-level consistency

one-copy availability, weak consistency

strong consistency, reduced availability

- Consistency is **continuous** rather than binary for many Internet services
  - Airline reservation system: percentage of aborted reservations
  - Application -specific quantification of consistency

Quantifying Consistency in TACT

- In TACT

  Consistency of a replica = (*Unseen Writes, Uncommitted Writes, Staleness*)

  - *Unseen Writes*: Number of updates not seen by a replica
  - *Uncommitted Writes*: Number of uncommitted updates on a replica
  - *Staleness*: The age of the "oldest" update not seen by this replica
  - These metrics are directly related to the consistency observed by apps

- Arbitrary consistency level achieved by bounding three metrics
- At two extremes, TACT demonstrates the behavior of
  - Voting algorithm (read quorum 1; write quorum n)
  - Optimistic replication system
  - TACT's goal is to provides everything in between

Trading Consistency for Availability

- Access to a service replica may be denied if specified consistency cannot be obtained due to partitions

- The weaker the consistency level =>

  The smaller set of servers that must be contacted to get consistency =>

  The higher probability these servers can be contacted =>

  The higher the availability

- Weaker the consistency level =>

  Improved performance

Status and Future Work

- We are using simulation to:
  - Validate our consistency metrics
  - Validate TACT toolkit design
- Next step:
  - Prototype implementation
  - Sample Internet service implementation on top of TACT
- Further investigation of availability/performance/consistency tradeoffs

TACT: Turning the Switch into a Knob

Email:{yhf, vahdat}@cs.duke.edu

QUESTIONS?