

Raft User Study


Diego Ongaro & John Ousterhout
Stanford University
SEDCL Retreat

June 6, 2013

Intro


- ▶ Last year's retreat: talk on Paxos
- ▶ John started a competing algorithm
- ▶ Designed Raft to be easier to understand
- ▶ Our reviewers didn't believe us
- ▶ Conducted an experiment to demonstrate that Raft is easier to understand than Paxos

Context: replicated state machines



- ▶ State machine defines data structure
 - ▶ Interface is application-specific
- ▶ Replicated log feeds commands to state machine
- ▶ Same log \Rightarrow same sequence of states, outputs
- ▶ Raft and Multi-Paxos are two consensus algorithms to manage the replicated log

1. Serial operation



- ▶ Basic Paxos defines consensus on just one log entry
- ▶ Multi-Paxos forms a log and optimizes across entries
- ▶ Each log entry can proceed concurrently
- ▶ **What's the advantage of concurrent operation?**
 - ▶ Ultimately the state machine must consume entries serially
- ▶ Raft appends entries to the log in order

2. Strong leader

- ▶ Raft first elects a cluster leader
 - ▶ Only the leader appends to the replicated log
 - ▶ Inconsistencies arise only on leader changes
- ▶ Basic Paxos is symmetric (p2p)
- ▶ Multi-Paxos introduces a leader as an optimization

A few tips from Scott

- ▶ Don't just ask people their opinion: measure it
- ▶ Record the lectures
- ▶ Pilot everything twice
- ▶ *Doing Psychology Experiments* by David W. Martin

From the participants' view

- ▶ Participants: undergrad and grad students
 - ▶ Stanford's Advanced OS class: 32 (5% participation grade)
 - ▶ Berkeley's Distributed Computing class: 16 (obvious bluff)
- ▶ Get randomly assigned to group
- ▶ Log onto web site, watch 1 hr Paxos (Raft) video, take 1 hr Paxos (Raft) quiz
- ▶ sleep($60 \times 60 \times 24 \times 5$)
- ▶ Log onto web site, watch 1 hr Raft (Paxos) video, take 1 hr Raft (Paxos) quiz
- ▶ Take short survey


Lecture challenges

- ▶ Same lecturer or expert on each algorithm?
- ▶ Which Paxos do we teach? How much do we improve it?
- ▶ What material do we include?

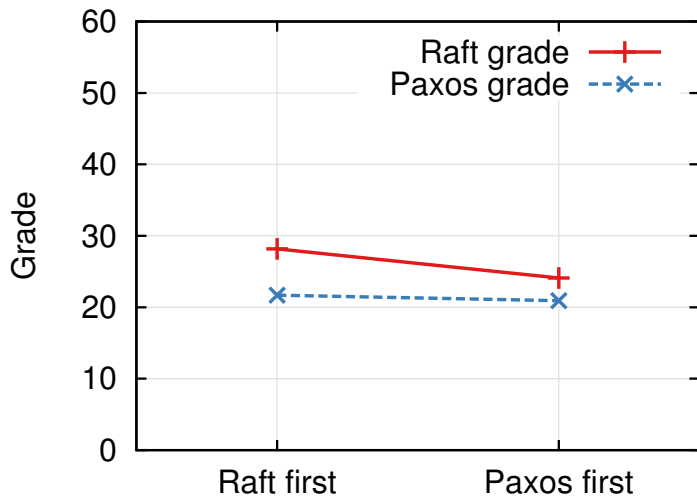
Quiz challenge: maintain equal difficulty

- ▶ Easy questions (4 points): warm-up
- ▶ Medium questions (26 points): apply algorithm
- ▶ Hard questions (30 points): not clear which algorithm to apply
- ▶ Paired question difficulty across exams


Quiz results



Ordering effects



Survey results



Recent popularity

A bunch of open-source Raft implementations:

Bloom	3
C++	1
Erlang	4
F#	1
Go	2
Haskell	1
Java	2

Upcoming talks:

- ▶ StrangeLoop (Ben Johnson of go-raft, September)
- ▶ RICON West (Diego, October)

Conclusions

- ▶ Really hard to measure understandability
 - ▶ 99% of effort before getting any results
- ▶ Students averaged 23% better on Raft quiz
- ▶ Survey showed overwhelming support for Raft
- ▶ Recent academic and industrial interest is encouraging
- ▶ Under submission...

<http://ramcloud.stanford.edu/raft.pdf>

Acknowledgements

- ▶ Scott Klemmer for help with experimental design
- ▶ David Mazières and Ali Ghodsi for their guinea pigs
- ▶ “Volunteers” for pilots: Aleks, Amit, Ankita, Arjun, Daniel, David, and Vimal
- ▶ The 48 participants
- ▶ Nelson Ray for his stats help