

Project 3: Raft

Due: 11:59 PM, April 22, 2020

Contents

1	Getting Started	1
2	Introduction	2
3	Raft Overview	4
3.1	Leader Election	4
3.2	Log Replication	4
3.3	Log Commitment	4
3.4	Client Interaction	4
4	Implementation	4
4.1	Project Layout	4
4.2	Raft Protocol Implementation	5
4.3	CS 138 Specific Nuances	6
5	Testing	7
5.1	Building and Running	7
6	Style	8
7	Handing in	8
8	Extra Credit	8

1 Getting Started

Remember, if you write code on a department machine, you must use `go1.13` instead of just `go` (ie `go1.13 install` or `go1.13 test`). Add `alias go=go1.13` to your `~/.bashrc` for convenience.

```
go1.13 get -u -d github.com/brown-csci1380-s20/raft-<TeamName>
```

Use the command above to get your repo from GitHub Classroom. Fix the import path in all files in the `cmd` directory, `client` directory, and `raft/cluster.go`. Use `go1.13 get -u ./...` to fetch all dependency. Before you get started, please make sure you have read over, set up, and understand all the support code.

We highly encourage you to work in groups of two, but we understand that in some situations a group of three may be necessary. If you work in a group of three, you must implement an additional feature. These can be found in the Extra Credit section at the end of the handout. Stop by TA hours to learn more about what these are! If you work in a group of three, you must contact the TAs and let them know you intend to work in a group of three, and if you will be implementing additional features.

Working alone is not allowed for this project. If you do not have a partner for any reason, please attempt to find one thru the piazza partner search functionality, and if this is not successful please email the htas for assistance.

2 Introduction

An important part of creating total-ordered, fault-tolerant distributed systems is providing the ability for multiple nodes to come to a consensus about state.

The problem of distributed consensus has been around for a long time and has typically been solved using implementations of the popular **Paxos** algorithm, which was initially published in 1998. Paxos, however, has been shown to be difficult to fully understand, let alone implement. The difficulties related to Paxos have spawned much work over the years in trying to make a more practical protocol.

In this spirit, a group of researchers at Stanford (Diego Ongaro and John Ousterhout) developed the **Raft protocol** in 2014, which is what you will be implementing in this project. Raft is a consensus protocol that was designed with the primary goal of understandability without compromising on correctness or performance (when compared to protocols like Paxos).

For this project, the Raft paper will serve as the **central source of truth**. All implementation details are found in the paper and this handout will simply refer to relevant sections in the paper for your guidance.

This project is also meant for you to experience what it's like to implement an algorithm directly described from a paper. Raft is great for this in that it is very thorough in describing all the details necessary for its implementation. We hope that this skill will be useful for you in your future CS career.

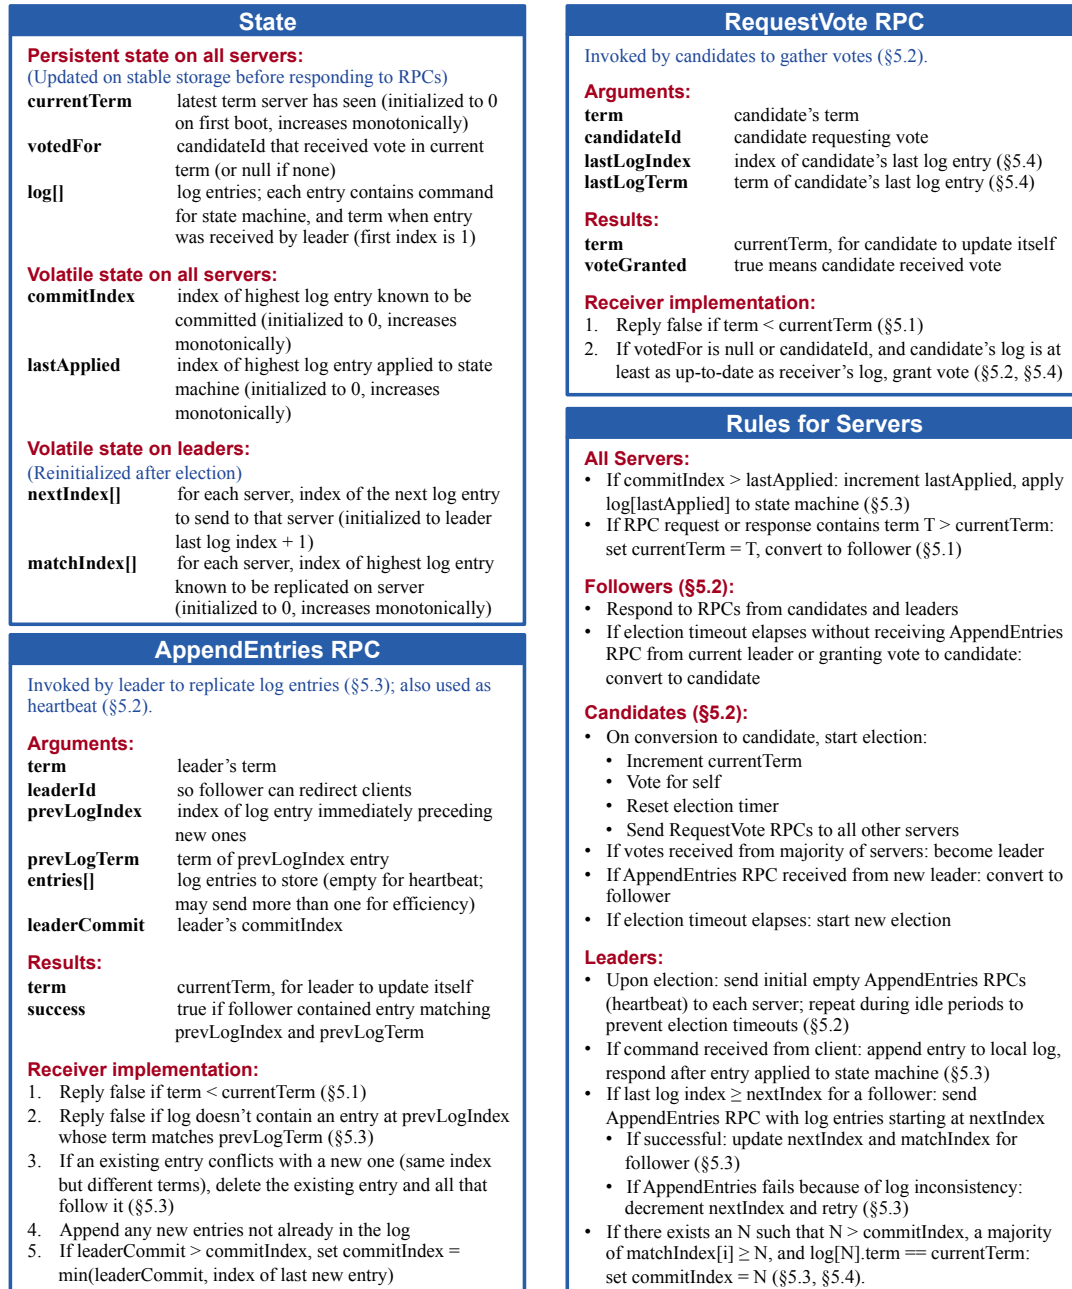


Figure 2: A condensed summary of the Raft consensus algorithm (excluding membership changes and log compaction). The server behavior in the upper-left box is described as a set of rules that trigger independently and repeatedly. Section numbers such as §5.2 indicate where particular features are discussed. A formal specification [31] describes the algorithm more precisely.

Figure 1: This figure is from the Raft paper, and helps summarize a lot of the important details in the protocol. If you find a note particularly confusing, refer to the section specified (e.g., §3.5).

3 Raft Overview

We hope the last homework provided you a nice introduction to the Raft protocol. If you would like to refresh your memory on an overview of Raft, refer to section 5.1.

The Raft protocol can be broken down into four major components that you will have to implement: Leader Election, Log Replication, Log Commitment, and Client Interaction.

For your reference we have included a cheatsheet summary of the consensus algorithm in Figure 1.

3.1 Leader Election

Leader election consists of a Raft cluster deciding which of the nodes in the cluster should be the leader of a given term. Refer to the cheatsheet, section §5.2 and §5.4.1 for implementation details.

3.2 Log Replication

Log replication consists of making sure that the Raft state machine is up to date across a majority of nodes in the cluster. Refer to the cheatsheet, section §5.3, §5.4.2, and §5.5 for implementation details.

3.3 Log Commitment

Log commitment is responsible for ensuring consistency of data in the Raft cluster since leaders are allowed to overwrite follower logs. Refer to the cheatsheet, section §5.3, §5.4.2, and figure 8 from the paper for implementation details.

3.4 Client Interaction

Client interaction consists of a node outside the Raft cluster making requests to modify the state machine of the cluster. Refer to the cheatsheet and section §8 for implementation details.

4 Implementation

4.1 Project Layout

In the stencil code, you'll see three high level packages:

- **raft**: the core Raft protocol
- **hashmachine**: a state machine based on a hash chain
- **client**: a client API for the **raft** package

You'll see a `cmd` directory with three packages inside: `raft-cli`, which is a CLI for a Raft client (it uses the `client` package above) and `raft-node`, a CLI to create and control a Raft node, and `raft-cluster`, a CLI to simultaneously start a group of Raft nodes to test leader election. Note that the Raft client is specific to the hash chain implementation we provide.

4.2 Raft Protocol Implementation

You should implement the following functions:

- `utils.go`
 - `func randomTimeout(minTimeout time.Duration) <-chan time.Time`
- `node.follower_state.go`
 - `func (r *Node) doFollower() stateFunction`
 - `func (r *Node) handleAppendEntries(msg AppendEntriesMsg) (resetTimeout, fallback bool)`
- `node.candidate_state.go`
 - `func (r *Node) doCandidate() stateFunction`
 - `func (r *Node) requestVotes(electionResults chan bool, fallback chan bool, currTerm uint64)`
 - `func (r *Node) handleCompetingRequestVote(msg RequestVoteMsg) (fallback bool)`
- `node.leader_state.go`
 - `func (r *Node) doLeader() stateFunction`
 - `func (r *Node) sendHeartbeats() (fallback, sentToMajority bool)`

Each function that returns a `stateFunction` should contain the logic for the Raft node being in one of the three Raft states: FOLLOWER, CANDIDATE, LEADER. You can transition to another state by returning that state function. For example, `doLeader()` could be written to always transition to the FOLLOWER state:

```
func (r *RaftNode) doLeader() state {  
    return r.doFollower  
}
```

When an RPC is received, the request is forwarded over a channel so that the function of the appropriate state can determine how it should be interpreted. For example, the following code would always reply successful to an `AppendEntriesRPC`:

```
for {
  select {
    case msg := <-r.appendEntries:
      msg.reply <- AppendEntriesReply{
        r.GetCurrentTerm(),
        true,
      }
  }
}
```

The full list of channels you should handle are:

- `RaftNode.appendEntries`
- `RaftNode.requestVote`
- `RaftNode.registerClient`
- `RaftNode.clientRequest`
- `RaftNode.gracefulExit`

Feel free to add whatever support code you need. **DO NOT CHANGE THE PUBLIC APIs.**

4.3 CS 138 Specific Nuances

- `DefaultConfig()` specifies in-memory log storage by default. This means that if you shutdown your raft node, all of its logs and state will be lost. Raft requires persistence of logs and some state such as a node's last vote, so in-memory storage would violate those conditions. However, using in-memory storage will actually be incredibly convenient when you start testing your Raft implementation. When you would like to test a node failing and restarting, set the config's `inMemory` field to false.
- We use **BoltDB**, an embedded DB written in Go, as the underlying stable storage for Raft. BoltDB is widely used across the software industry and is featured in projects such as **Consul**, Hashicorp's service-discovery platform, and **InfluxDB**, a popular DB for metrics and analytics.
- A Raft node is uniquely identified by its listener port in this scheme. So if you start up a new node that has the same port as a node that was running before which had its state saved to disk, this new node will appear with the same state. To avoid this, you can either use in-memory storage, ensure that new nodes you start have distinct ports, or you can call `func (r *Node) RemoveLogs()` to remove the persisted data.

5 Testing

We expect to see several good test cases. This is going to be worth a portion of your grade. You can check your test coverage by using [Go's coverage tool](#).

To aid you with testing, we have provided sample tests inside `example_client_test.go`, `example_election_test.go`, and `example_partition_test.go`. These are not exhaustive, but test some core components of your implementation. Consider them a starting point for devising more complete tests of your own.

We have also provided the helper functions used by our tests in `test_utils.go`. Feel free to use them within your own tests.

To help you test out the behavior of your implementation under partitions, we have provided you with a framework which allows you to simulate different network splits. You can find the relevant code in `network_policy.go`, and see how it can be used in `cmd/raft-node/main.go` and `example_partition_test.go`.

5.1 Building and Running

Once you're in your repo's higher level `raft` directory, to get Raft to build, you must first update your dependencies, like so:

```
$ go get -u ./...
```

Then, you can build and install our three binaries, `raft-node`, `raft-client`, and `raft-cluster` by running:

```
$ cd cmd
$ go install ./...
```

This generates three CLIs and places them in your `$GOPATH/bin`:

You are welcome to improve any of the CLIs as you see fit. For a list of available commands in each CLI, type `help` after entering.

- **raft-node**

This is a CLI that serves as a console for interacting with Raft, creating nodes, and querying state.

You can pass the following arguments to `raft-node`:

- `-p <port>`: The port to start the server on. By default selects a random port.
- `-c <addr>`: Address of an existing Raft node to connect to.
- `-d=true`: Enable or disable debug. Default is false.
- `-m=true`: Enable or disable in-memory store. Default is true.

- **raft-client**

This is a sample client which allows you to connect to a Raft node and issue commands to the hash state machine.

You can pass the following arguments to **raft-client**:

- **-c <address>**: A raft node address for raft-client to connect to

- **raft-cluster**

This is a CLI that will simultaneously start up a cluster of raft nodes and is meant primarily to test for leader-election

You can pass the following arguments to **raft-cluster**:

- **-n <number>**: The number of nodes in the raft cluster. Default is 3.
- **-m=true**: Use in-memory storage for raft cluster. Default is true
- **-d=true**: Enable or disable debug. Default is false

6 Style

You should use Go's formatting tool **gofmt** to format your code before handing in. You can format your code by running:

```
gofmt -w=true *.go
```

This will overwrite your code with a formatted version of it for all go files in the current directory.

7 Handing in

You need to write a README documenting any bugs in your code, any extra features you added, and anything else you think the TAs should know about your project. Document the test cases you created and briefly describe the scenarios you covered.

When you are done, push to your GitHub repo. We will pull your latest commit in the master branch for grading.

8 Extra Credit

If you're in a 3-person group, your group must implement one additional feature and demonstrate it works (via tests, a CLI, or benchmarks). 2-person groups can start thinking about these too, as you may want to implement them for an A-level design of your final project, Puddlestore.

Some ideas for extra credit include:

- Membership Changes (§6 of the Raft paper and §4 of the dissertation)
- Log Compaction (§7 of the Raft paper and §5 of the dissertation)
- Web application to control and interface with a Raft node (using gRPC).
(Refer to this [repo](#) for an example on node.js and grpc-web.)
- Snapshots, using the Chandy-Lamport or another algorithm

If you choose to implement one of these, please drop by TA hours or email the TA list to discuss your plan first and get any questions answered.

Feedback

Please let us know if you find any mistakes, inconsistencies, or confusing language in this or any other CS138 document by filling out [the anonymous feedback form](#).