

# Problem Set 4

CSE 452 / CSE M552

May 28, 2017

Submit short, **typeset** answers to the following questions. Please work on this individually. You may not use skip days on this assignment.

## Problem 1: Chubby

Consider a system with many clients simultaneously accessing Chubby. Each client frequently reads, and occasionally writes, the same file. Reviewing the logs, the sysadmin notices that the writes are usually resolved quickly, but occasionally display latency spikes of up to 12 seconds over normal latency. What is the origin of these latency spikes? Be specific.

## Problem 2: BigTable

Consider a BigTable client whose cache is empty—it has never talked to this BigTable deployment before. In order to read a single row with key  $k$  from table  $T$ , what requests does it need to make? Be specific.

## Problem 3: GFS

Why does GFS's "record append" operation sometimes insert duplicate records into the log? Give one scenario (a sequence of message sends and failures) that can cause this behavior.

## Problem 4: Spanner

Suppose you execute the following two transactions against Spanner:

```
start read/write transaction
  write("a", "1")
end transaction
start snapshot read at T
  print read("a")
end snapshot read
```

The initial value associated with the key "a" is "0". Can the code print "0" if:

- (a) `T=TT.now().earliest`
- (b) `T=TT.now().latest`
- (c) T is the timestamp at which the previous transaction was committed