

Computer Science Department

New York University

## G22.3033-001 Distributed Systems: Fall 2010

# Quiz I

All problems are open-ended questions. In order to receive credit you must answer the question *as precisely as possible*. You have 80 minutes to answer this quiz.

Some questions may be much harder than others. Read them all through first and attack them in the order that allows you to make the most progress. If you find a question ambiguous, be sure to write down any assumptions you make. Be neat. If we can't understand your answer, we can't give you credit!

**THIS IS AN OPEN BOOK, OPEN NOTES QUIZ.**

I (xx/20)	II (xx/10)	III (xx/20)	IV (xx/15)	V (xx/10)	VI (xx/10)	Total (xx/85)

**Name:**

## I Multiple choice questions:

Answer the following multiple-choice questions. Circle *all* answers that apply. Each problem is worth 4 points. Each missing or wrong answer costs -2 point.

**A.** Which of the following statements are true about the design of YFS and NFS?

1. A buggy NFS client program might corrupt the shared file system data structure.
2. A buggy YFS client program might corrupt the shared file system data structure.
3. There are many more different types of RPC messages between a NFS client and the NFS server than there are between a YFS client and the extent/lock server.
4. If only one YFS client is allowed to perform file system writes (e.g. create file, append data etc.) while all other YFS clients can only perform read-only file system operations, there is no need for using a lock service.

**B.** Which of the following statements are true for sequential consistency (SC)?

1. SC is the strongest and yet practical consistency model.
2. SC allows for high availability during periods of network disconnectivity.
3. SC allows a client thread to read a stale value of some variable.
4. SC is only applicable to distributed shared memory systems and not applicable to distributed storage systems.
5. Suppose there are two concurrent write operations (executed by different threads): one changes variable A to 10, the other one changes variable B to 20. With SC, it's possible that one thread reads A=10, B=0 and another concurrent thread reads A=0, B=20 (assume both variables have zero as their initial value).

**C.** Which of the following things are true for Bayou and Tra?

1. All storage replicas in Bayou and Tra eventually have identical state.
2. Bayou is able to detect update conflicts (i.e. when updates happened in a non-sequential manner.)
3. When using Tra, one must keep a replica for all data in the entire file system tree on every node.
4. One can and should build a banking application on top of Tra or Bayou.

**D.** Which of the following statements are true?

1. Serializability requires the database manager to sequentially execute transactions, one after another.
2. Serializability demands that, if the database manager receives transaction A before transaction B, it will commit (or abort) transaction A before committing (or aborting) transaction B.
3. Snapshot isolation is a multiversion concurrency control scheme that realizes serializability.
4. No multi-version concurrency control scheme can realize serializability; one must use 2-phase locking for serializability.

## II Remote Procedure Call

After attending the lecture on failure recovery, Ben Bitdiddle becomes paranoid about failures. He is now concerned that the RPC library he completed in Lab 1 might not preserve at-most-once guarantee after the server reboots from a crash. After all, when the server recovers from a crash, it loses all its state about previous received RPCs and is at risk of executing a retransmitted RPC that it has already executed before the crash.

**1. [10 points]:** Does your completed RPC server in Lab 1 preserve at-most-once guarantee across reboots? If not, give a counterexample. If yes, please explain.

Hint: Recall that in our RPC library, each client creates a `rpcc` object and `bind()`s it to the corresponding server before using it. Further, the RPC packet header contains the following fields:

```
// add RPC fields before the RPC request data
req_header h(ca.xid, proc, clt_nonce_, srv_nonce_, xid_rep_window_.front());
req.pack_req_header(h);
```

### III Threads and locks

Ben Bitdiddle wants to implement a reader-writer mutex (R/W lock) using ordinary pthread library. Ben has come up with the following simple implementation:

```
struct rwlock_t {
    int nreaders; //number of readers who have grabbed the R/W lock, initialized to 0
    int nwriters; //number of writers who have grabbed the R/W lock, initialized to 0
    pthread_mutex_t m;
    pthread_cond_t cv;
    rwlock_t() { //some initialization code here... }
};

void
read_lock(rwlock_t *rw) {
    pthread_mutex_lock(&rw->m);
    if (rw->nwriters > 0) {
        pthread_cond_wait(&rw->cv, &rw->m);
    }
    rw->nreaders++;
    pthread_mutex_unlock(&rw->m);
}

void
read_unlock(rwlock_t *rw) {
    pthread_mutex_lock(&rw->m);
    rw->nreaders--;
    if (rw->nreaders == 0) {
        pthread_cond_signal(&rw->cv);
    }
    pthread_mutex_unlock(&rw->m);
}

void
write_lock(rwlock_t *rw) {
    pthread_mutex_lock(&rw->m);
    if (rw->nreaders > 0 || rw->nwriters > 0) {
        pthread_cond_wait(&rw->cv, &rw->m);
    }
    rw->nwriters++;
    pthread_mutex_unlock(&rw->m);
}

void
write_unlock(rwlock_t *rw) {
    pthread_mutex_lock(&rw->m);
    rw->nwriters--;
    pthread_cond_signal(&rw->cv);
    pthread_mutex_unlock(&rw->m);
}
```

**2. [10 points]:** Ben has noticed multiple problems when testing the correctness of his R/W lock: 1) sometimes, there are multiple readers and writers simultaneously holding the lock, violating correctness. 2) sometimes, while multiple readers should be able to simultaneously hold the lock in read mode, he found the code is causing them to complete the `read_lock` operation serially, one after another.

Please correct Ben's errors for him. (You may directly mark your corrections on the previous page.)

**3. [10 points]:** Alice P. Hacker is concerned that Ben's implementation is not fair to writers. In other words, readers might starve a waiting writer: there could be arbitrarily many readers coming and holding the lock while a writer is waiting for the lock. Please give a R/W lock implementation that favors the writer. In other words, any waiting writer can grab the R/W lock after a bounded amount of time.

continued from previous page....

## IV Crash Recovery

In the class, we have simply assumed that writing to a single disk sector is atomic with respect to failures. In other words, the write operation either succeeds in its entirety (leaving the new data) or not at all (leaving the old data intact). Ben Bitdiddle is using a disk for which this assumption does not hold. The disk may fail in the middle of an sector write, leaving partially written new data and corrupting existing data on that sector.

Ben would like to implement `atomic_write()` and `atomic_read()` on top of this disk. The correct behavior of these functions should be: if `atomic_write(addr, data)` succeeds, all subsequent `atomic_read(addr, buffer)` returns the new data written. If `atomic_write(addr, data)` fails because of a disk crash, all subsequent `atomic_read(addr, buffer)` returns the old data stored at *addr* before the crashed `atomic_write` started.

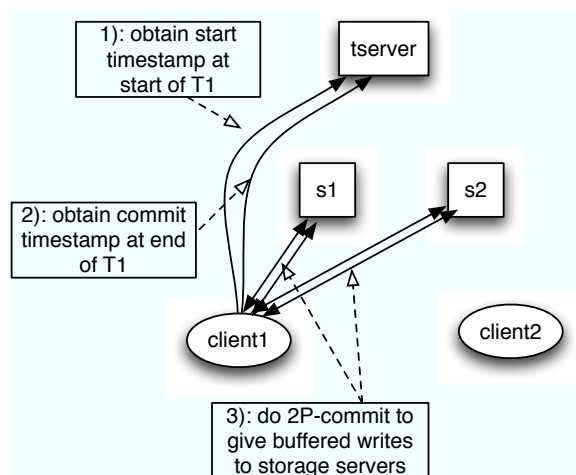
- 4. [5 points]:** Ben's first idea is to only write data in the first 500 byte of a 512-byte sector and reserve the last 12 byte as the checksum. Ben implements the `chsum_write(int addr, char[500] data)` function which writes both the 500-byte *data* and its checksum, `chsum(data)`, into the underlying 512-byte sector with address *addr*. The `bool chsum_read(int addr, char *buffer)` function reads the underlying 512-byte sector and verifies that the last 12-byte checksum is correct for the preceding 500-byte data block. If so, the data is returned, otherwise a failure is returned. Can Ben directly use `chsum_write` and `chsum_read` for atomically reading and writing data blocks of 500-byte in size? Please explain.

**5. [10 points]:** Please help Ben design and implement the `atomic_write` and `atomic_read` functions. Hint: you might want to consider storing two copies of checksummed data at one address.



## V 2P commit and Snapshot isolation

Ben Bitdiddle wants to implement snapshot isolation in a distributed storage service. Ben's system consists of a single timestamp server (*tserver*) and a number of storage servers (e.g. *s1* and *s2*) (see Figure below).



Here's an example of how Ben's system works. Suppose *client1* is executing a transaction *T1* that writes two items *X* and *Y*. *client1* first obtains a start timestamp for *T1* from *tserver*, say,  $T1.sts = 1$ . During execution of the transaction, *client1* buffers writes to *X* and *Y* locally. After *T1* commits, *client1* obtains a commit timestamp from *tserver*, say,  $T1.cts = 10$ . Since *s1* is the server responsible for storing *X* and *s2* is the server responsible for storing *Y*, *client1* performs a 2P-commit (as the coordinator) with *s1* and *s2* to 1) check whether *T1* suffers from write-write conflicts with other concurrent transactions and 2) write both items to *s1* and *s2* atomically if *T1* can commit. In particular, *client1* sends  $T1.sts$  and  $T1.cts$  as well as the buffered write of *X* to *s1* in a 2P-prepare message. Likewise, it also sends  $T1.sts$  and  $T1.cts$  as well as the buffered write of *Y* to *s2* as a 2P-prepare message. Both *s1* and *s2* check for write-write conflicts locally and vote accordingly.

Let's examine how an ongoing 2P-commit affects concurrent reads. For example, suppose *T1* is in the middle of 2P-commit and *s1* has just voted "yes" to commit *T1*'s write of *X* (but it has not known the final outcome of the 2P-commit yet). Suppose another concurrent transaction *T2* wants to read *X* from *s1* with start timestamp  $T2 = 11$ .

- 6. [10 points]:** Which of the following action(s) is correct for *s1* to serve *T2*'s reads?
1. *s1* returns the write of *T1* (i.e. the version of *X* with timestamp 10) to *T2* immediately.
  2. *s1* returns the latest version of *X* that's no bigger than *T1*'s commit timestamp, e.g. suppose there's a version of *X* with timestamp 9, then *s1* can return this version as the result of read to *T2* immediately.
  3. *s1* blocks the read of *T2* until it has received the outcome of the 2P commit.

Please explain your answer. Specifically, if you think a choice is incorrect, please explain why it is wrong.

continued from previous page....

## VI G22.3033-001

7. [5 points]: Describe the most memorable error you have made so far in one of the labs. (Provide enough detail so that we can understand your answer.)

We would like to hear your opinions about the class so far, so please answer the following two questions.

8. [3 points]: What is the best aspect of this class?

9. [2 points]: What is the worst aspect of this class?

# End of Quiz I