

Q1 Networking

Part A


Imagine you are in charge of a company-wide data transfer project where your company needs to transfer 100TB of data stored in the NYC datacenter to the LA datacenter. Come up with two proposals, where one approach has bigger bandwidth but higher latency, and the other approach has lower latency but smaller bandwidth. Describe each proposal in one or two sentences. (Note: You cannot simply use one proposal with different latency/bandwidth assumptions)

Part B

For each of the scenarios, state whether you should use UDP or TCP. Briefly explain why.

1. Streaming live 15-440/640 lectures
2. A remote SSH session on AFS to test your P0 code
3. A time server broadcasting the correct time to all machines on a LAN

Part C



Alice (A) is trying to communicate with Bob (B) via a network with one router (R) in between. Data from A to B needs to be broken into 10Kbit packets. Alice knows that the link 1 between A and R has 200Mbps capacity, but she doesn't know that of link 2. Alice wanted to test and find out the capacity of link 2.

Test 1: Alice sent a very small piece of data (negligible size) to Bob and had Bob send it back immediately, and she recorded that the time interval is 0.5s.

Test 2: Alice sent over a 1GB file and had Bob send back a checksum of the file immediately, and she recorded that this time interval is 50.0s.

Assume that computing checksum on Bob's computer has a speed of 10Gbps, and Bob can only start computing the checksum after the entire file is received. Ignore all other processing times on the router and Bob.

For all math in the following questions you must show your work to receive full credit.
(Reminder: 1B = 8 bits)

1. Assume that the router operates in cut-through mode. What is the capacity of link 2 (in Mbps, keep two decimal points)?
2. If the router is reconfigured to be in store-and-forward mode, what will the time interval in Test 2 be?
3. The router is configured back to cut-through mode. Assume that connection is lossy and each packet has a 0.001% chance of being corrupted in transit. Assume that any corruption will result in a different checksum. After Test 2, Alice compares the checksum received from Bob with the checksum she has. If they are different, Alice resends the entire file. (For simplicity, ignore the possibility of the checksum being corrupted.) What is the expected number of times that Alice needs to resend the file (round to nearest int)?
4. Propose a different protocol where Alice can shorten the expected time that Bob receives the file without corruption. You don't need to show any calculations.

Q2 Classic Synchronization

“Office Hours during the pandemic”

Due to the COVID-19 pandemic, the university has published new policies for students attending Office Hours. While a TA can help with exactly one student at a time, other students must wait patiently in a separate waiting room. The waiting room has a maximum capacity of 10 people, and there will only be 10 seats available in the room.

When the TA finishes answering a student's question, he/she dismisses the student and goes to the waiting room. If the TA sees there are other students waiting, he/she will ask one of the students to leave the seat and go to the office, then the TA starts answering the student's question. If there are no students in the waiting room, the TA will return to the office, relaxing for a while until a student comes to the office and asks questions.

When a student arrives, he/she should first check if there are available seats in the waiting room. If the room is full and there are no seats left, the student should leave and attend another OH

later. If the student finds a free seat in the waiting room, he/she will grab that seat and then go to the office to check whether the TA is there. If the TA is not there or the TA is busy with answering another student's question, the student will go back to the waiting room, sitting on his/her seat and waiting for the TA to come here. If the TA is taking a rest, the student will release his/her seat in the waiting room and directly start asking questions.

Note: The conditions inside each sub-question below are independent. They only hold in the scope of a sub-question, unless specified otherwise.

1. Tony, a student taking the *Distributed Systems* course this semester, wants to implement the scene above using pseudocode. But there is one thing he disagrees with in the university policy. He believes that TAs and students shouldn't just check each other's availability only on arrival and then wait. When they are waiting, they should both check each other's availability on a frequent basis, for example, every 5 minutes. He starts from drafting some pseudocode of students' and TA's actions:

```
// Student actions
...
lookForTA:
    goToOffice;
    if taAvailable          // If the TA is waiting in the office.
        askQuestion;
    else
        goToWaitingRoom;
        waitFor5Minutes;
        goto lookForTA;    // Check again after 5 minutes.
                           // Only apply to this sub-question.
...

// TA actions
...
lookForStudent:
    goToWaitingRoom;
    if studentWaiting      // If the TA sees a student waiting in the waiting room.
        answerQuestion;
    else
        goToOffice;
        waitFor5Minutes;
        goto lookForStudent; // Check again after 5 minutes.
                             // Only apply to this sub-question.
...
```

Clearly, without using a synchronization algorithm, the above pseudocode will have many concurrency issues. From what you have learned in class, please explain the concept of a "livelock", and show an example of a potential livelock that might occur in the above scenario.

2. Some students believe that they could wait outside the waiting room if the room is full. They have proposed a piece of code describing a student's actions:

```
// Global Variables
int numFreeSeats = 10;

// Student actions
void student() {
    onArrive:
        while (numFreeSeats <= 0) {
            // Wait outside the waiting room.
            ;
        }
    // Double-check if there are free seats in the room.
    if (numFreeSeats > 0) {
        numFreeSeats--;
        waitForMyTurn(); // Block until the student's turn. Assume that it is
                        // implemented correctly with no concurrency issues.

        askQuestion(); // On TA's side, there will be a corresponding:
                       //      numFreeSeats++;
                       // to make sure the number of free seats increases by
                       // one when the student's question is answered.
    } else {
        // Wait outside again.
        goto onArrive;
    }
}
```

Assume the TA's actions are implemented correctly. Are there any potential concurrency issues on the students' side in the above code snippet? Please explain your answer.

3. Jane, one of the top students in the class, comes up with a more complex implementation of the behavior of the TA and the students, by using semaphores:

```
// Global Variables
int numFreeSeats = 10; // The capacity of the waiting room.
semaphore taReady = 0; // Initialized to 0.
semaphore checkSeats = 1; // Initialized to 1.
semaphore studentReady = 0; // Initialized to 0.

// TA actions
void ta() {
    while (true) {
        P(studentReady);
```

```

        P(checkSeats);
        numFreeSeats++;
        V(taReady);
        V(checkSeats);
        answerQuestion(); // Block until finish answering a student's
                           // question.
    }
}

// Student actions
void student() {
    P(checkSeats);
    if (numFreeSeats > 0) {
        numFreeSeats--;
        V(studentReady);
        V(checkSeats);
        P(taReady)
        askQuestion(); // Block until the question is answered.
    }
    // Leave otherwise.
}

```

However, when she tests her code, as more and more students come to the OH, the program blocks at some point. Please identify this concurrency issue, explain why it would happen, and propose a fix to it by making changes to the code. (Hint: You shouldn't add/delete/modify more than 5 lines of code.)

4. Consider the code in sub-question 3 with your fix to the bug. Assume there are two students coming to the OH, one after another. The second student arrives while the TA is answering the first student's question. The second student waits for 10 minutes in the waiting room before the TA dismisses the first student and picks up the second student. Please fill in the following table, where each line represents an event of changing in the value of a semaphore. The events are happening in sequential order.

Note: There will be only one correct answer to receive full points. During the procedure, there can be cases where two events can happen in any order. For those cases, we will specify one possible case and fill in those rows for you.

| Event Index | Action Taker (TA / Student1 / Student2) | checkSeats | taReady | studentReady |
|-------------|---|------------|---------|--------------|
| Initial | \ | 1 | 0 | 0 |

| | | | | |
|----|----------|---|---|---|
| 1 | | | | |
| 2 | | | | |
| 3 | Student1 | 1 | 0 | 1 |
| 4 | TA | 1 | 0 | 0 |
| 5 | | | | |
| 6 | | | | |
| 7 | TA | 1 | 1 | 0 |
| 8 | Student1 | 1 | 0 | 0 |
| 9 | | | | |
| 10 | | | | |
| 11 | | | | |
| 12 | | | | |
| 13 | | | | |
| 14 | | | | |
| 15 | TA | 1 | 1 | 0 |
| 16 | Student2 | 1 | 0 | 0 |

5. Consider the code in sub-question 3 with your fix to the bug. Do you see another remaining problem with this implementation? (Hint: Is this implementation fair to every student?) If you identify the problem, briefly explain it and describe a possible solution. (Plain text only, no pseudocode needed.)

Q3 Time Synchronization

Consider three processes. The system has totally ordered clocks by breaking ties by process ID. It uses the Ricart & Agrawala algorithm. The timestamp for each process of ID i is $T(p) = 10 * L(p) + i$, where $L(p)$ is a regular Lamport clock.

Each message takes 2 ‘real time’ steps to get delivered. The critical section takes 4 real-time steps. This means that if a process receives all replies at real time $t-1$, it starts executing the critical section at time t , and finishes executing and sends messages at time $t+4$.

Assume that if a process receives messages from the other two processes at the same time, the message that comes from the lower process ID will be received first.

Fill in the table with the messages that are being broadcast, sent, or received between the processes until all nodes have executed their critical sections. The first few rows have been filled for you. Then use your filled table to answer the questions below.

Action types: Broadcast (B), Receive (R), Send (S), Execute Critical Section (ExecCS), Exit Critical Section (ExitCS).

Initial timestamps:

$L(p1) = 19, T(p1) = 191,$

$L(p2) = 27, T(p2) = 272,$

$L(p3) = 6, T(p3) = 63$

| Real Time | Process ID | Lamport Time | Action | Message | Queue at P1 | Queue at P2 | Queue at P3 |
|-----------|-----------------------|---------------------------------|--|---|-------------|------------------|------------------|
| 1 | 1 3 | 201 73 | B B | request 201 request 73 | 201 | | 73 |
| 2 | 2 | 282 | B | request 282 | 201 | 282 | 73 |
| 3 | 1 2 2 3 | 211 292 302 213 | R from 3 R from 1 R from 3 R from 1 | request 73 request 201 request 73 request 201 | 73 201 | 73 201 282 | 73 201 |
| 4 | 1 1 2 2 3 | 291 301 312 322 293 | R from 2 S to 3 S to 3 S to 1 R from 2 | request 282 reply 73 reply 73 reply 201 request 282 | 201 282 | 282 | 73 201 282 |
| 5 | | | | | 201 282 | 282 | 73 201 282 |
| 6 | | | | | | | |

| | | | | | | | |
|----|--|--|--|--|--|--|--|
| 7 | | | | | | | |
| 8 | | | | | | | |
| 9 | | | | | | | |
| 10 | | | | | | | |
| 11 | | | | | | | |
| 12 | | | | | | | |
| 13 | | | | | | | |
| 14 | | | | | | | |
| 15 | | | | | | | |
| 16 | | | | | | | |
| 17 | | | | | | | |
| 18 | | | | | | | |
| 19 | | | | | | | |
| 20 | | | | | | | |
| 21 | | | | | | | |
| 22 | | | | | | | |
| 23 | | | | | | | |
| 24 | | | | | | | |
| 25 | | | | | | | |

Questions:

1. At what real time does each process start executing the critical section?
2. At what timestamp does each process start executing the critical section?
3. What does the queue for each process look like at real time 12?

Q4 Distributed Mutual Exclusion

1. Describe the fairness of each of the following distributed mutual exclusion algorithms: Ricart & Agrawala, token ring, decentralized mutual exclusion, and centralized mutual exclusion. Please use 2 or fewer sentences for each algorithm.
2. Describe how each of the following distributed mutual exclusion algorithms is impacted by node failure: Ricart & Agrawala, token ring, decentralized mutual exclusion, and centralized mutual exclusion. Please use 2 or fewer sentences for each algorithm.
3. Explain why the Ricart & Agrawala algorithm is starvation and deadlock free.

Q5 Remote Procedure Calls

For each of the following scenarios, state whether you would use an RPC guaranteeing at-most-once, at-least-once, or exactly-once semantics (exactly-once should work in all cases, but if it is overkill you should pick the weakest guarantee that is enough for your application). In some cases multiple answers can be accepted, but you need to make sure your explanation is clear and coherent with your answer, and your (reasonable) assumptions are made explicit. Explain in no more than 3 sentences.

1. Depositing money into a bank account
2. Uploading your proof of vaccination to SIO
3. Making a login attempt online where your account will be frozen after 3 incorrect attempts
4. Cancelling a fraudulent order that's been placed on your credit card

5. Booking a cross-country airplane ticket to go to your internship