CS425 Fall 2023 – Homework 2

(a.k.a. "Once upon a time in Distributed Hollywood")

Out: Sep 13, 2023. Due: 2 pm US Central on Oct 2 (Monday!), 2023

Topics: P2P Systems, Key-value Stores, Time and Ordering (Lectures 7-12)

Instructions:

1. **Attempt any 8 out of the 10 problems** in this homework (regardless of how many credits you're taking the course for). If you attempt more, we will grade only the first 8 solutions that appear in your homework (and ignore the rest). Choose wisely!
2. Please hand in **solutions that are typed** (you may use your favorite word processor. We will not accept handwritten solutions. Figures (e.g., timeline questions) and equations (if any) may be drawn by hand (and scanned).
3. **All students (On-campus and Online/Coursera)** – Please submit PDF only! Please submit on Gradescope. [https://www.gradescope.com/]
4. Please **start each problem on a fresh page**, and **type your name at the top of each page**. **And on Gradescope please tag each page with the problem number!**
5. Homeworks will be **due at time and date noted above. No extensions. For DRES students only:** once the solutions are posted (typically a few hours after the HW is due), subsequent submissions will get a zero**. All non-DRES students must submit by the deadline time+date.**
6. Each problem has the same grade value as the others (10 points each).
7. Unless otherwise specified in the question, the only resources you can avail of in your HWs are the provided course materials (slides, textbooks, etc.), and communication with instructor/TA via discussion forum and e-mail.
8. You can discuss lecture concepts and the questions on Piazza and with your friends, but you cannot discuss solutions or ideas on Piazza.

Prologue: You have just been made the technical head in a production company that is producing a new Hollywood movie. The movie is sure to be a blockbuster, with a lot of well-known actors and actresses hired to star in it. Amazingly many of them know distributed systems! You run into them every day on the set. Here is what ensues.

All characters and their actions used in this homework are meant to make the homework fun! Any resemblance of their actions or opinions to real events, or places, is purely coincidental. Any stories involving real actors or actresses are fictional.

Problems:

1. Walt Disney and Pokemon wants to build a virtual theme park where young customers will all wear VR goggles and be connected via a Gnutella P2P system. At one point of time, the Gnutella topology looks like a virtual ring with 15 processes, with each process connected to its immediate three clockwise neighbors and immediate three anticlockwise neighbors (thus each process has six neighbors). All links are bidirectional. Note that any node can be the original query-sending node (for each of the three parts below). Answer three parts:
    a. A process sends a Query message with TTL=4. How many of the processes, apart from the sender, receive this Query message?
    b. What is the minimum TTL in a Query message to ensure all nodes in the system receive the Query (no matter the timing of the Query forwarding)?
    c. If we add a 16th process that is directly connected to all the old 15 processes, what is the minimum TTL in a Query message to ensure all nodes in the system receive the Query?
2. Three wealthy industrialists, Lo Skum, Frank Bozo, and Sugar Mountain, all want to make a movie aboard a very unsafe spaceship called "Solargate" that travels to all 8 planets (go figure!). Anyway, you are responsible for managing the spaceship as long as it is active. Lo and behold, the different planets use BitTorrent! While working through an example, you find a case where a file has 7 shards: A, B, C, D, E, F, G. The querying node (peer, leecher) is on a spacecraft, connected to all the planets directly (the SBM group figured out a way to have high bandwidth, low latency communication that beats the speed of light!). The planet servers at the 8 planets each have the following shards: Mercury (ACEG), Venus (BDF), Earth (ABCDEFG), Mars(ABCDEF), Jupiter (ABCDE), Saturn (ABCD), Uranus (ABC), Neptune (AB).
    a. Which shard will be fetched *first*?
    b. What is the *order* in which shards will be fetched by the querying node? (assume that no shards are being fetched by any other nodes). If there are ties, you can use the alphabetically lower one.
    c. If one introduced a new ~~10th~~ renegade planet called DeGrasseTyson (and the spacecraft was connected to it), assign a set consisting of *just one shard* to DeGrasseTyson that changes the *first* fetched shard by the spacecraft (ties broken by lower letter)?
3. One of the Producers, Leo Bloom, and his start Orlando Bloom, both like Bloom filters. But being a producer, he wants to create a new Bloom filter-based data structure. A (regular) Bloom filter's false positive rate is given as $\left(1 - e^{-\frac{kn}{m}}\right)^k$

where *k* is the number of hash functions, *n* is the size of the input set and *m* is the size of the Bloom filter in bits. Leo says that instead of using a single Bloom filter B with 4096 bits and 2 hash functions, his new datastructure, called Leo Bloom filter, uses 4 Bloom filters B1, B2, B3, and B4, each with 1024 bits, and each using 2 hash functions (each hash function different from each other, and different from the above 2 hash functions). There are two variants: When checking for an item, it returns true only if the item is present in {Variant A: *all*, Variant B: *any*} of B1, B2, B3, and B4. When inserting an item, for both variants A and B, it is inserted into *all* of B1, B2, B3, and B4. Which of the above three approaches—original using Bloom filter B vs. Leo–any vs. Leo–all —gives the best (lowest) false positive rates? Answer this for two cases: (1) when there are typically 10 elements inserted into the datastructure, (2) when there are typically 1000 elements inserted into the datastructure. (We recommend though, that you solve the problem with the variables *k, n, m, etc.,* and then apply these values. But solving with only these two values of 10, 1000, would be ok as well.)

4. The director, C. Nolan, likes to deal with time travel, which means he asks a lot of "What if?" questions. He asks you several questions about Cassandra – for each of these, give one advantage and one disadvantage:
    a. What if there were no Bloom filter (but the rest of the system works the same way)?
    b. What if Cassandra did not use Memtable at all, and instead wrote directly into the "latest" SSTable? (but the rest of the system works the same way)
    c. What if SSTables were mutable, i.e., when a Memtable were flushed, instead of creating a new SSTable, it would go and modify the corresponding values in already-present SSTables? (but the rest of the system works the same way)
5. Help, the stars are stuck in a Black Mirror episode. In a system of N processes (N large enough, etc.), k quorum sets Q1, Q2, Q3, Qk are selected in an arbitrary manner, each Qi of the same size M. k is a constant, much smaller than N. What should this <u>minimum</u> value of M be so that at least one (any) process belongs to ALL three of Q1, Q2, Q3,…Qk (that is, the intersection of all k sets is non-null)? If this problem is too abstract (and daunting), we recommend you start with k=2 (as discussed in class), and quantitatively reason/think about why it has to be N/2+1. Then solve for k=3, then solve k=4, and so on. And then notice the patterns.
6. (For this question you can search resources on the Web.) One of the actresses, named Meryl, is consistently a good actress and consistently wins awards. It's no surprise that she is very interested when you tell her about consistency models. She asks you about the differences between *linearizability, sequential consistency,*

causal consistency, and eventual consistency (for key-value stores with get/put operations on keys).

  a. Can you say briefly, and clearly what the differences are between the above models? (50 words or less)
  b. Give an example (using 2-3 clients writing and reading objects), where, for a particular read, using each of these models above gives a completely different return value (so 4 different answers). While you can search the Web to clarify differences between these models, you *cannot* borrow an example from the Web. Be concise.

7. An overly-religious person suddenly wants to fund the show, and they only worship Cristian's algorithm. They find that the round-trip time for one round of synchronization messages is exactly 0.78 ms. They are trying to calculate the error for this run, and so they calculate some minimum delays. On the server side, they find that there is a delay of at least 56.0 microseconds for a packet to get from an application to the network interface, and the delay on the opposite path (network interface to application buffer) is at least 0.12 ms. On the client side, they measure that the minimum time to get from the network interface to the application buffer is at least 0.34 ms, and the minimum time on the reverse path is X microseconds. But they forget to measure X. They figure that for the error in this Cristian's algorithm run to be zero, X should lie in an interval [A microseconds, B microseconds]. That is, A and B are respectively the minimum and maximum values possible for X so that the error for this run is zero. What should the values of A and B be, given these measurements?
8. Katherine Johnson was a Black computer mathematician and scientist who worked at NASA in building/writing the first NASA computer programs to do mathematical calculations for the geometrical trajectories of the Apollo moon spacecraft, being one of the first to write computer programs involving space and time. The astrophysicist consultant for your current film, a Dr. N. Tyson, is also a big fan of space and time. He is intrigued by the NTP algorithm. He poses the following question – suppose you are dealing with a network where latencies are *symmetrical* (e.g., an enhanced and controlled ATM network), i.e., latencies are identical in reverse directions, i.e., A → B message latency and B → A message latency (for same message) are known to be identical (except of course you don't know their absolute values), for every pair of processes A,B. From the equations, for NTP, can you derive the new error bound? Don't just show the last portion, please work through the entire derivation just like in the slides, from the beginning to the end.
9. Spiderman and his doppelgangers in the metaverse are trying to figure out the communication among the different metaverses. They have the following

timeline of messages exchanged, where each dot represents an instruction. Can you mark *Lamport timestamps* on each event? It is ok to print out this and hand-draw/write the timestamps (then scan or photograph your solution, and insert it into your solution doc).

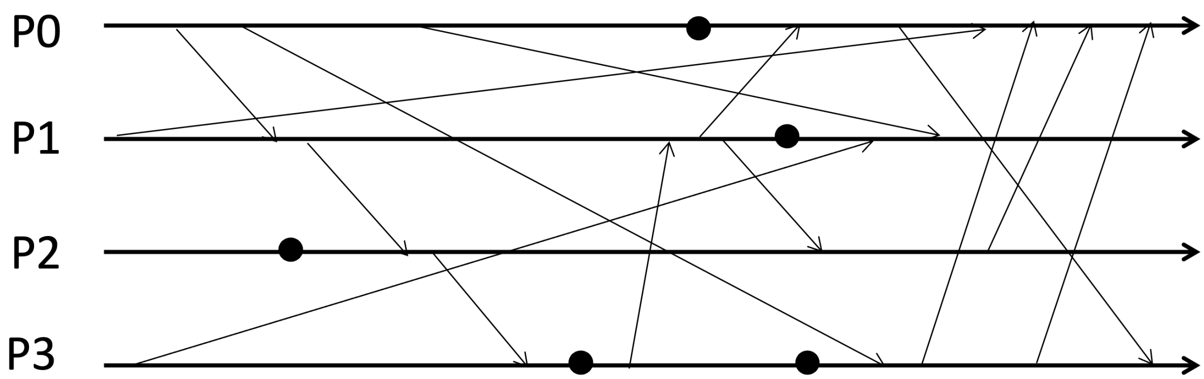

10. Whoops, the simultaneous and sudden arrival of all of the concurrent villains (Doc Ock, Green Goblin, Electro, Sandman, and the others) means Lamport timestamps are unusable for distinguishing concurrent events from causal events. Can you mark *vector timestamps* for the same timeline as in the previous question? It is ok to print out this and hand-draw/write the timestamps (then scan or photograph your solution, and insert it into your solution doc).

====== END OF HOMEWORK 2 =====