15-440/15-640: Homework 3

Due: November 8, 2018 11:59pm

Name:
Andrew ID:

1 GFS FTW (25 points)

Part A (10 points)

The Google File System (GFS) is an extremely popular filesystem used by Google for a lot of its everyday tasks. AFS has also gained its fair share of fame especially among educational institutes. State whether AFS or GFS is better for each of the following cases. Give **one** line of explanation for your answer.

1. Tolerating high rate of failure of the file server. [**2 pts**]
2. Repeatedly accessing the same file. [**2 pts**].
3. Large file reads. [**2 pts**].
4. Multiple clients (several hundred) concurrently adding entries to the same todo file. [**2 pts**].
5. Low server storage requirements. [**2 pts**].

Part B (5 points)

Recall that the GFS master server is the central repository of a lot of information and serves an important purpose in GFS. The chunkservers store chunks of files.

1. The GFS master stores metadata such as the Namespace (directory hierarchy), the mapping from files to chunks etc. State **one** reason why the master server crashing is especially bad for **recovery**? [**1 pts**]
2. State **two** other tasks that the GFS master server does. [**2 pts**].
3. The Chunkservers store 64MB file chunks on local disk. State **one** disadvantage of choosing a smaller file chunksize. [**1 pts**].
4. Explain **one** reason why the GFS client does not cache data. [**1 pts**].

Part C (10 points)

The 440 staff uses GFS with 4 chunkservers (A, B, C, D) and one master server. We only store one file that consists of one chunk and its replicated across A, B and C. Assume that A is the chunk's primary. Zeleena and Amadou don't talk to each other and decide that it's a good idea to edit the same file at the same time. Zeleena wants to fill the chunks with "ZZZZZ...", while Amadou wants to fill the file with "AAAAA...". Zeleena and Amadou operate on the same chunk at the same time. The chunks are initialized to "MMMMM...". Initially, assume that there are no failures.

1. If Zeleena and Amadou use WRITE, is it possible that Amadou writes to a chunk on B "AAAAA..." but Amadou's data gets overwritten by Zeleena with "ZZZZ..."? [**2 pts**]

2. If Zeleena and Amadou use APPEND, is it possible that Amadou writes to a chunk on B "AAAAA..." but Amadou's data gets overwritten by Zeleena with "ZZZZ..."? **[2 pts]**.

3. Now B suddenly fails. **Name** the mechanism the master uses to detect this failure. State the next steps once the master has detected the failure. **[3 pts]**.

4. B is back up again and everything is back to normal. Now Zeleena wants to delete the file. In **three short** bullet points explain the steps once the client (Zeleena) has issued the delete request. **[3 pts]**.

2 The Return Of BurgerNet! (20 points)

Part A (15 points)

By popular demand we have decided to change the name of BergerNet to BurgerNet. BurgerNet is a social network that has suddenly gained a lot of popularity. BurgerNet hosts its services in its own proprietary servers located in Pittsburgh. We now need your help to scale BurgerNet.

1. Sam has travelled to London and is trying to access his pics on BurgerNet but its significantly slower to load than when he was in Pittsburgh. Explain the possible reason. **[2 pts]**.

2. Sam is trying to load his profile picture but it takes X milliseconds when he is in London, while it only takes 5 ms in NYC. The speed of information travel is 10^5 mile/s. The distance between Pittsburgh and London is 3000 miles and the distance between Pittsburgh and NYC is 300 miles. The bandwidth is 1 Gb/s. Give an exact value for X. Assume there is no ACKing and to 'load' a picture you only have to send the picture across the network. **[6 pts]**

3. We talked in class about something called a CDN. Define the term CDN. In **two short sentences** explain how a large CDN like **Akamai** could potentially help alleviate the problem described above. **[3 pts]**.

4. BurgerNet users normally access the same picture multiple number of times. What other technique could help make BurgerNet even faster? Explain in **one sentence** how. **[2 pts]**.

5. BurgerNet decides to add a collaborative document and picture editing. This feature is now very popular. Answer and explain if Akamai CDNs would still help? **[2 pts]**.

Part B (5 points)

BurgerNet has decided to adopt the model of Sequential Consistency over its backend infrastructure. One day you want to post comments on a controversial event. You first remove some people from your list of friends (transaction T1), and then post your comments (transaction T2).

1. Is it possible that the people you remove in T1 could potentially see the post in T2? Explain your answer in one sentence. **[2 pts]**.

2. Define External Consistency. **[1 pts]**

3. Spanner is a globally distributed database that uses "TrueTime" to guarantee external consistency. With Spanner is it possible for the people you remove in T1 to see the post in T2? Explain your answer in **one sentence**. **[2 pts]**.

3 A Side of Consistent Hashing (10 points)

You are the primary software engineer who manages and maintains a new CDN. You're responsible for choosing where files reside in the servers and therefore, how to direct user requests. You decide to distributed the content based on the hash value of the file ID (int), with the goal of having a good load balancer that doesn't have to move too many files on server failures. Initially you decide to use separate chaining as the underlying technique, as shown below:

Part A (4 points)

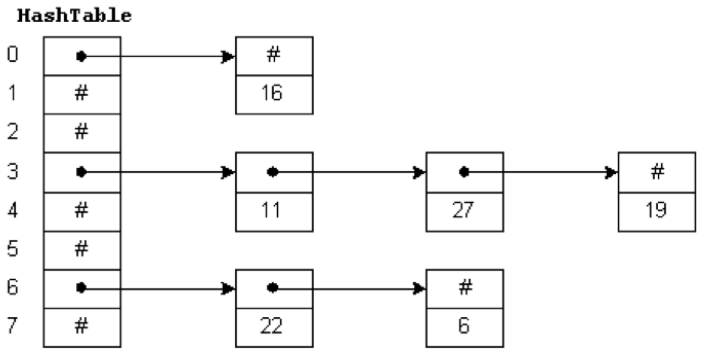

The hash function for the table is simply f(x) = x. Hence a number k will get placed at index f(k) % sizeOfArray. Example: f(22) = 22. 22 % 8 = 6. Hence 22 is placed at index 6.

1. Due to increased load, we decide to increase the size of the hash-table. A re-size causes the array of the hashtable to be increased by size 1. Assume that there is a re-size operation that takes place right after the state shown above. How many keys would need to be remapped? **[2 pts]**.
2. State the array index where numbers 11 and 6 would go after resizing. **[2 pts]**.

Part B (6 points)

Your colleague who took 15-440 mentions that you could load balance in a better way, and recommends using consistent hashing, as seen below:

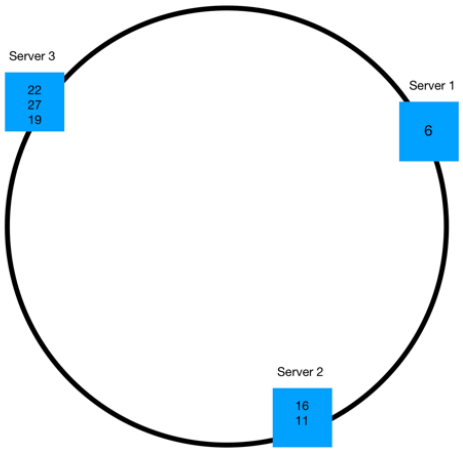

The hashing function is f(x) = (x*10)° mod 360°. The ring above represents a 360° circle. Each number is mapped to the closest server whose corresponding degree is great than or equal to its own angle. Server 1 is at 65°, Server 2 is at 170° and Server 3 is at 270°. Example: f(6) = 60° and hence it would be mapped to Server 1.

In each question below assume the state starts off as shown in the figure above.

1. Which servers would the numbers 17, 0, 28 and 18 be mapped to? **[2 pts]**.
2. Now assume that Server 3 is overloaded. Hence we decide to add a new Server 4 at 240°. State the numbers that need to be remapped to the new server? **[2 pts]**.
3. Suddenly Server 2 fails. State the numbers that need to remapped and which server would they be mapped to. **[2 pts]**.

4 MoCkInG sPoNgEbOb (20 points)

Chelsea wants to de-stress before her exam and she checks out the newest SpongeMock memes online by directing the browser to **spongemock.memes.com**. The local DNS server then performs an iterative lookup. The diagram below shows some of the DNS records contained in each DNS server. Note that DNS responses are cached in the local DNS server.

localdns.cmu.edu (S1)

Record Number	Name	Value	Type	TTL
R1	.	a.root.net	NS	1 day
R2	a.root.net	198.40.4.8	A	1 day

a.root.net (S2)

Record Number	Name	Value	Type	TTL
R3	com.	b.gtld.net	NS	12 hours
R4	b.gtld.net	198.31.2.90	A	12 hours

b.gtld.net (S3)

Record Number	Name	Value	Type	TTL
R5	memes.com.	ns-9.memes.com	NS	2 hours
R6	ns-9.memes.com	83.102.188.3	A	2 hours

ns-9.memes.com (S4)

Record Number	Name	Value	Type	TTL
R7	spongebob.memes.com	83.102.188.4	A	30 minutes
R8	spongemock.memes.com	83.102.188.5	A	30 minutes
R9	spoderman.memes.com	83.102.188.6	A	30 minutes

1. Fill in following table to indicate the sequence of queries and responses exchanged among the servers. **[8 pts]**

	Sender	Receiver	Type (Query/Response)	Data
1	Chelsea's PC	S1	Query	spongemock.memes.com
2				
3				
4				
5				
6				
7				
8				
9				
10				

2. Fill in any new DNS records in the local DNS server right after the sequence of queries and responses in (a). Label the new records with record numbers starting from R10. **[4 pts]**.

Record Number	Name	Value	Type	TTL
R1	.	a.root.net	NS	1 day
R2	a.root.net	198.40.4.8	A	1 day
R10				
R11				

3. After six hours, Chelsea is done with her exam and checks spongemock.memes.com again to see if there are any updates. Fill in the DNS records in the local DNS server right before any queries and responses are performed for her second request. **[4 pts]**.

Record Number	Name	Value	Type	TTL

4. Once again, fill in following table to indicate the sequence of queries and responses exchanged among the servers for Chelsea's second request. **[4 pts]**.

	Sender	Receiver	Type (Query/Response)	Data
1	Chelsea's PC	S1	Query	spongemock.memes.com
2				
3				
4				
5				
6				
7				
8				
9				
10				

5 MapReduce (10 points)

Freddie just bought a MapReduce compute cluster and wanted to run the following tasks on his new cluster. For each of the following tasks, state whether it is a good fit for MapReduce and explain your answer. If MapReduce is a good fit, describe how you will define the mappers and reducers for the task; otherwise, explain why MapReduce does not work well.

1. Support a Bitcoin mining service that is similar to the one you implemented in Project 1. There are over 10^{10} requests, and each request is consists of message M and unsigned integers N1, N2. The Bitcoin service needs to find the unsigned integer pair (n1, n2) that generates the largest hash value for all $0 <= n1 <= N1$, $0 <= n2 <= N2$ when concatenated with M. **[2 pts]**.
2. Solve an object recognition task using stochastic gradient descent method. This involves over 10^5 iterations. **[2 pts]**.
3. Calculate the average amount of time spent on coursework in the past week among all CMU students. **[2 pts]**.
4. Analyze the similarities of 10^9 websites. Specifically, for each website A, find the website B that is most similar to A. **[2 pts]**.
5. Maintain a trading simulation app that updates stock prices every minute based on 10^8 real-time transactions. **[2 pts]**.

6 MapReduce's Limitations & Spark (7 points)

1. A key performance challenge in MapReduce is called "stragglers". What's a practical technique to alleviate the straggler problem? (One sentence) **[1 pts]**
2. List and explain in a single sentence two additional limitations of MapReduce that are addressed by Spark (one sentence each). **[2 pts]**
3. Spark enables fault-tolerance in-memory cluster computation using Resilient Distributed Datasets (RDD). **[2 pts]**
   (a) Explain the idea of lineage in enabling fault tolerance in two sentences.
   (b) Why are RDDs immutable? Please give a one-sentence example.

4. Give a specific scenario where Spark is preferred over Hadoop MapReduce. Briefly explain your answer. **[1 pts]**.
5. Give a specific scenario where Hadoop MapReduce is preferred over Spark. Briefly explain your answer. **[1 pts]**.

7 We are Industrial Fans (8 points)

The purpose of this question is to expose you to some of the famous distributed systems used in industry. These questions are intended to be **short** answers. Discuss in 2-3 sentences what each one of these technologies is and **one famous use-case**. The links to the papers have been provided to you.

1. BigTable **[4 pts]**.
2. DynamoDB **[4 pts]**.