

## Distributed Systems

2016 Exam 3 Review

Paul Krzyzanowski  
Rutgers University  
Fall 2016

### Question 1

Unlike the design of Chord, Amazon Dynamo stores the entire list of nodes on each system. Explain the advantage of doing this instead of using a finger table.

---

Each node knows exactly which node contains the desired key and can forward the query there directly. There is no need for multiple hops.

A finger table may require hops:  $O(\log N)$  vs.  $O(1)$

*Bad answers:*

- Better fault tolerance
- Easier to add new nodes (if every node stores the entire list, you still need to update all nodes)

December 15, 2016 © 2016 Paul Krzyzanowski 2

### Question 2

You have access to a file of class enrollment lists. Each line contains `{course_number, student_id}`.

Explain how you would use MapReduce to get information on how many classes students take.

For instance, you may discover that 1,495 students are enrolled in 6 courses; 13,077 students are enrolled in 5 courses; 14,946 students are enrolled in 4 courses; and 4,484 students are enrolled in 3 courses.

Explain each map and reduce operation. You may use pseudocode and assume that functions such as sum and count exist. Be sure to state the inputs & outputs of each step.

*Hint: you may need more than one iteration*

December 15, 2016 © 2016 Paul Krzyzanowski 3

### Question 2 (cont.)

You have access to a file of class enrollment lists. Each line contains `{course_number, student_id}`.

**First MapReduce: find # of courses each student takes**

Map\_1:  
input: { course\_number, student\_id } 1 for each course per student  
output: { key=student\_id, 1 }

Reduce\_1:  
input: { student\_id, courses[] }  
output: { student\_id, sum(courses) } We can also output { 1, sum(courses) }

**Second MapReduce: find # of students that take each course count**

Map\_2:  
input: { student\_id, course\_count } 1 for each student with course\_count courses  
output: { course\_count, 1 }

Reduce\_2:  
input: { course\_count, students[] }  
output: { course\_count, count(students[]) }

December 15, 2016 © 2016 Paul Krzyzanowski 4

### Question 2 (cont.)

<p><b>Input</b></p> <pre>213 130972375 416 162692062 416 534744968 416 021693896 417 162692062 417 130972375 519 130972375 ...</pre>	➔	<p><b>Output from map</b></p> <pre>key=130972375, 1 key=162692062 1 key=534744968 1 key=021693896 1 key=162692062 1 key=130972375 1 key=130972375 1 ...</pre>
<p><b>Input to reduce</b></p> <pre>130972375 1, 1, 1 162692062 1, 1 534744968 1, 1, 1, 1 021693896 1, 1, 1, 1 ...</pre>	➔	<p><b>Output from reduce</b></p> <pre>130972375 3 162692062 2 534744968 4 021693896 4 ...</pre>

Student 130972375 takes 3 classes

December 15, 2016 © 2016 Paul Krzyzanowski 5

### Question 2 (cont.)

<p><b>Input to map = output from reduce</b></p> <pre>130972375 3 162692062 2 534744968 4 021693896 4 ...</pre>	➔	<p><b>Output from map</b></p> <pre>key=3, 1 key=2, 1 key=4, 1 key=4, 1 ...</pre>
<p><b>Input to reduce</b></p> <pre>2, {1, 1, 1, 1, ... } 3, {1, 1, 1, 1, 1, ... } 4, {1, 1, 1, 1, 1, 1, ... } ...</pre>	➔	<p><b>Output from reduce</b></p> <pre>2, 1622 3, 4484 4, 14946 ...</pre>

"one student taking 2 courses"  
"one student taking 4 courses"  
"one student taking 4 courses"

December 15, 2016 © 2016 Paul Krzyzanowski 6

### Question 3

How does Spanner provide consistent lock-free reads of lots of data even if other transactions are modifying some of that data during the read?

Spanner stores multiple timestamped versions in each field.

Snapshot reads allow reading of data whose  
version  $\leq$  transaction start timestamp

December 15, 2016

© 2016 Paul Krzyzanowski

7

### Question 4

Explain why each of the following statements is incorrect.  
In Pregel, each compute process represents an edge of a graph along with the data that is sent on that edge.

No. Each process represents a vertex

December 15, 2016

© 2016 Paul Krzyzanowski

8

### Question 5

Explain why each of the following statements is incorrect.  
A shared nothing architecture does not allow two systems to access the same NFS server.

A shared nothing architecture does not allow two systems to share the same block-level storage (e.g., cluster file system).

*No credit for simply negating the statement!*

December 15, 2016

© 2016 Paul Krzyzanowski

9

### Question 6

Explain why each of the following statements is incorrect.  
Akamai uses dynamic DNS to look up a URL and return the address of the closest caching server that has that URL in its cache.

1. Akamai uses dynamic DNS to look up domain names, not URLs!  
DNS does not do URL lookups
2. Akamai does not necessarily return the closest server.  
It will return one that is available, not loaded, likely to contain content, ... and is closest (lowest latency)

December 15, 2016

© 2016 Paul Krzyzanowski

10

### Question 7

Explain why each of the following statements is incorrect.  
A digital certificate contains a hash that is encrypted with the certificate owner's private key.

Encrypted hash = signature.

A digital certificate is signed by the certification authority (CA), not the owner:

The hash is encrypted with the CA's private key

*Not: encrypted with the certificate owner's public key*

December 15, 2016

© 2016 Paul Krzyzanowski

11

### Question 7

Explain why each of the following statements is incorrect.  
A digital certificate contains a hash that is encrypted with the certificate owner's private key.

Encrypted hash = signature.

A digital certificate is signed by the certification authority (CA), not the owner:

The hash is encrypted with the CA's private key

*Not: encrypted with the certificate owner's public key*

December 15, 2016

© 2016 Paul Krzyzanowski

12

### Question 8

What is an advantage of Amazon Dynamo's *virtual nodes* over simply having each physical machine be a node?

- Performance is improved because you can create a massive number of virtual nodes.
- It's cheaper because the node runs on a virtual machine that can do other things as well.
- A virtual node can be easily migrated to another system.
- You can shed load from multiple systems when you add a new physical machine..**

Adding a new machine adds more virtual nodes that are scattered throughout the ring.

→ Each virtual node takes on some {key, value} sets from another node  
On average, these other nodes are spread across multiple physical systems

December 15, 2016

© 2016 Paul Krzyzanowski

13

### Question 9

Amazon Dynamo's *optimistic replication*:

- Uses Paxos to ensure that all replicas contain identical copies.
- Grabs locks on all replicas of objects while they are being updated.
- Forces each transaction to check at commit time whether the data was modified by another transaction.
- Allows some replication to fail, causing some of the replicas to have older versions.**

Optimistic replication = it's ok to fail on replication

Dynamo is designed to be always available and eventually consistent.

December 15, 2016

© 2016 Paul Krzyzanowski

14

### Question 10

MapReduce can process a lot of data relatively quickly because:

- Input data is broken up among many map workers.**
- Map workers run in parallel with reduce workers.
- Reduce workers shrink the amount of data that needs to be processed.
- All of the above.

MapReduce is a great example of divide-and-conquer.

(b) No. Reduce workers run after every single map worker is done

(c) Only sometimes with multiple iterations of MapReduce. In general, map workers get rid of unnecessary data and reduce workers compute results

December 15, 2016

© 2016 Paul Krzyzanowski

15

### Question 11

The *partitioning* function in MapReduce determines:

- How the original input data is divided into multiple shards.
- Which reduce worker will work on a specific key.**
- How the map worker parses input data to generate a key.
- The maximum number of values that will be assigned to each key.

The partitioning function is applied to the {key, value} results from map workers.

December 15, 2016

© 2016 Paul Krzyzanowski

16

### Question 12

A column family in Bigtable is:

- A group of columns within a single table.**
- A group of related columns among multiple tables.
- A set of columns that are shared across multiple tables.
- A preconfigured set of columns to make it easier to set up a Bigtable instance.

A column family is a storage unit for multiple columns in a table – usually related in function.

From a performance point, data in a row is retrieved by column families

December 15, 2016

© 2016 Paul Krzyzanowski

17

### Question 13

Rows in Bigtable:

- Are always stored in a sorted order by the row key.**
- Are ordered chronologically; the last row always gets added to the end of the table.
- Are not sorted but a row can be located by looking up its key in a hash table stored at the master.
- Are not sorted or searchable but an application can get a list of all rows that contain data for a specific column.

December 15, 2016

© 2016 Paul Krzyzanowski

18

### Question 14

Spanner transactions are kept isolated by:

- Using the TrueTime API to ensure no two transactions start at the same time.
- Using strict two-phase locking on resources.
- Using a two-phase commit protocol.
- Waiting until all previous transactions have committed.

(a) Transactions may start at the same time.

(c) This deals with commits, not isolation during the transaction.

(d) No – but this would also deal with commits rather than isolation

December 15, 2016

© 2016 Paul Krzyzanowski

19

### Question 15

To provide external consistency in Spanner, commits are delayed:

- Until all earlier transactions have released all their locks.
- All the locks used by the transaction are released.
- Up to the estimated clock skew between clocks in different locations.
- Until all earlier transactions have committed.

External consistency = order of transactions reflects their true time order

Commit wait: before a transaction commits, it acquires a commit timestamp:

$$t = TT.now().latest$$

$t$  = the latest possible value of the true time across all servers in the system.

Do not release any locks until that time is definitely in the past.

This means waiting until the earliest possible current time on any system is greater than the transaction timestamp:

$$TT.now().earliest > t$$

December 15, 2016

© 2016 Paul Krzyzanowski

20

### Question 16

In the Pregel framework, the entire job is complete when:

- Every compute function vote to halt.
- Every compute function entered an inactive state.
- There are no incoming messages for any compute function.
- Every compute function entered an inactive state and has no incoming messages.

Voting to halt == entering an inactive state

BUT:

- A compute function can send a message before voting to halt
- If a compute function receives a message, it is reawakened – placed back into an active state

December 15, 2016

© 2016 Paul Krzyzanowski

21

### Question 17

A combiner in Pregel:

- Processes several messages to create a single input to a compute function.
- Allows the graph to keep global state.
- Combines multiple edges into a single edge.
- Combines multiple vertices into a single vertex.

Combiner:

User-defined function applied to a set of messages targeted to the same vertex – creates a single message to that vertex

Aggregator: global variable available to all vertices

December 15, 2016

© 2016 Paul Krzyzanowski

22

### Question 18

A Spark resilient distributed dataset (RDD) is fault tolerant because:

- It is stored as replicas on multiple servers.
- It is partitioned across many servers in the cluster.
- It can be regenerated using the transformation that created it.
- It is cached in memory with a backup on the disk of the servers that created it.

(a) No. An RDD is sharded but generally not replicated.

(b) Yes, it could be partitioned but that does not make it fault tolerant.

(d) Also possible but disk is generally used as a RAM overflow.

RDDs are created by transformations and can be re-created by those same transformations.

December 15, 2016

© 2016 Paul Krzyzanowski

23

### Question 19

Which statement best describes Spark's processing?

- A sequence of transformations followed by an action.
- A sequence of actions followed by a transformation.
- An alternating sequence of transformations and actions.
- An arbitrary mix of transformations or actions.

Transformations read an RDD and create a new RDD.

Example transformations are *map*, *filter*, *groupByKey*, and *reduceByKey*.

Actions are operations that evaluate and return a value to the driver program.

Example actions are *reduce*, *grab samples*, and *write to file*.

December 15, 2016

© 2016 Paul Krzyzanowski

24

### Question 20

A System Area Network (SAN) is typically:

- An overlay network that identifies the nodes that make up the cluster.
- A dedicated network that is used only for cluster management messages, such as heartbeats.
- A high bandwidth, high reliability, low latency network used to connect nodes of a cluster.**
- A high-speed network used to connect processors within a multiprocessor system.

High performance, low-latency, high-reliability LAN interconnect for cluster members.

Examples: Infiniband, Myrinet, ethernet with Data Center Bridging

December 15, 2016

© 2016 Paul Krzyzanowski

25

### Question 21

A Storage Area Network (SAN) is

- A network that provides block level access to storage devices via a network.**
- A network dedicated to distributed file system servers, such as NFS servers.
- A set of connected servers that provide file system services.
- An overlay network to allow systems in a cluster to identify storage servers.

SAN: network for remote disk access

December 15, 2016

© 2016 Paul Krzyzanowski

26

### Question 22

The difference between a clustered file system and a distributed file system is

- A clustered file system is only made available to a limited set of systems.
- Distributed file systems may span multiple servers.
- Clustered file systems provide replication for fault tolerance.
- All nodes access the clustered file system at the block level.**

Each system maintains its own file system and uses a shared block device for storage.

Distributed file system: server(s) maintain the file system and provide file system APIs for remote access (e.g., *read bytes from a file, write bytes to a file, create a file, delete a file*)

December 15, 2016

© 2016 Paul Krzyzanowski

27

### Question 23

Checkpointing enables

- Cold failover.
- Warm failover.**
- Hot failover.
- Cascading failover.

Cold failover = application restart

Warm failover = restart but restore state from last checkpoint

Hot failover = backup node is always synchronized and ready to take over

Cascading failover = failover node failing over

December 15, 2016

© 2016 Paul Krzyzanowski

28

### Question 24

Akamai's overlay network is primarily designed to:

- Enable load balancing.
- Enable load shedding.
- Provide secure transport of requests within Akamai's systems.
- Find faster routes than might be available via Internet routing.**

Akamai collects statistics about its servers and their availability & connectivity

Core function of Akamai: content caching

Akamai's dynamic DNS enables load balancing & load shedding

Overlay network enables Akamai to find faster routes to origin servers

December 15, 2016

© 2016 Paul Krzyzanowski

29

### Question 25

The Diffie-Hellman algorithm was designed to:

- Allow two parties to come up with a shared secret key while using a public network.**
- Allow one party to use public key cryptography to send a secret key to the other party.
- Ensure that messages are signed as well as encrypted.
- Enable an arbitrary number of parties to agree on a common secret key so they all can communicate.

The Diffie-Hellman algorithm allows two parties to create a common key while conversing over an insecure network.

- Does not use public key cryptography
- Does not use digital signatures
- No. Only two parties can come up with a common key.

December 15, 2016

© 2016 Paul Krzyzanowski

30

### Question 26

A session key is:

- A key derived from one user's private key and another user's public key.
- A randomly generated key designed to be used for one communication session.
- A key that is used only for signing messages.
- A key that automatically expires after a certain number of messages have been sent.

Session key = random, throw-away key used for one session.

December 15, 2016

© 2016 Paul Krzyzanowski

31

### Question 27

Salt in a password is:

- A permutation of a password to ensure that it does not match any dictionary words or well-known passwords.
- A secret transformation of the password performed prior to storing it.
- Non-secret data added to a password prior to hashing it.**
- An encryption of the password that will require an administrator's secret key to decrypt it.

Salt is data added to a password to thwart attacks based on precomputed hashes

- Salt does not permute a password; it adds to it.
- It is not secret.
- It does not encrypt the password.

December 15, 2016

© 2016 Paul Krzyzanowski

32

### Question 28

For Alice to authenticate Bob, she sends him a bunch of bytes that:

- Bob encrypts with his private key.**
- Bob encrypts with his public key.
- Bob encrypts with Alice's private key.
- Bob encrypts with Alice's public key.

- Alice can decrypt the result using Bob's public key to validate that only Bob could have performed the encryption
- Anyone can encrypt the data with Bob's public key.
- Bob does not have access to Alice's private key.
- Anyone can encrypt the data with Alice's public key.

December 15, 2016

© 2016 Paul Krzyzanowski

33

### Question 29

A hybrid cryptosystem:

- Uses two layers of encryption for extra security.
- Combines the convenience of public key cryptography with the speed of symmetric cryptography.**
- Both encrypts and signs each message to prevent tampering.
- Alternates among multiple encryption algorithms throughout communication.

December 15, 2016

© 2016 Paul Krzyzanowski

34

### Question 30

Tunneling is the technique of

- Encrypting the data part of a packet before transmitting it.
- Encrypting the entire packet before transmitting it.
- Using the data part of a packet to transport another packet.**
- Signing a packet to prevent unauthorized modifications to it.

Tunneling is simply the encapsulation of one packet within the payload of another.

VPNs use tunneling and add encryption and/or signatures.

December 15, 2016

© 2016 Paul Krzyzanowski

35

The End

December 15, 2016

© 2016 Paul Krzyzanowski

36