

Distributed Systems

07r. Pre-Exam 2 Review

Paul Krzyzanowski
Rutgers University
Fall 2016

November 2, 2016 2014-2016 Paul Krzyzanowski 1

Fall 2015 – Question 1a

(a) State one advantage of early binding over late binding.

- Ability to cache the name resolution and use it multiple times.
- Reduced latency when the binding is actually needed.

Bad answer:

- *More efficient*

November 2, 2016 2014-2016 Paul Krzyzanowski 2

Fall 2015 – Question 1b

(b) State one advantage of late binding over early binding.

- You get the most up-to-date results.
For example, better support for resources that move around.

Bad answer:

- *More efficient*

November 2, 2016 2014-2016 Paul Krzyzanowski 3

Fall 2015 – Question 2

How does virtual synchrony handle a fail-recover situation? That is, what happens when a failed node come back?

A node that fails is taken out of the group.
When it comes back, it has to re-join the group.
Doing this requires a **state transfer** from another group member.

November 2, 2016 2014-2016 Paul Krzyzanowski 4

Fall 2015 – Question 3

Eric Brewer discusses consistency, availability, and partition tolerance in the CAP Theorem. Explain how it applies to traditional ACID semantics and the two-phase commit protocol?

Note: the question does NOT ask to explain why ACID semantics are not always the best to use and why you might want BASE (eventual consistency) instead. It asks you to explain the trade-off in providing ACID semantics.

Availability is lost to preserve consistency & partition tolerance.
If a network is partitioned & some processes cannot be reached, the coordinator will wait, trying again & again.
Sub-transactions cannot commit and will hold on to their locks, sacrificing availability.
ACID: Locking & 2PC forces consistency across partitions.

November 2, 2016 2014-2016 Paul Krzyzanowski 5

Fall 2014: Question 1

NFS was initially designed as a stateless file system. The design principle of many REST-based web services today is to keep the servers stateless. Explain why server state is useful in a network file system.

Server state allows:

- Support for file locking (server keeps track of client locks)
- Client to keep a file open so it doesn't get deleted on the server (server keeps state that a given file is in use).
- Server keeping track of which clients need to be informed for cache invalidations.

NOT – server state allows caching

NFS did client caching without state, using validation. However, it wasn't always consistent. State allows for server-initiated cache invalidations

No broad answers, such as "shorter requests", "faster performance" or "improved fault tolerance" without explaining your answers.

November 2, 2016 2014-2016 Paul Krzyzanowski 6

Fall 2014: Question 2

What is a false deadlock and under what conditions could a false deadlock arise?

Improper message ordering when reporting a wait-for graph can result in a coordinator thinking that a circular dependency, and hence deadlock, exists when that is really not the case.

Not: a host with a lock dies and then recovers.

November 2, 2016

2014-2016 Paul Krzyzanowski

7

Fall 2014: Question 3

Explain what happens in the first phase of a three-phase commit protocol.

The coordinator gets everyone's decision on whether they can commit or have to abort.

More precisely:

- (a) Coordinator sends *canCommit?* queries to all participants
- (b) Waits for responses from everyone
- (c) If timeout on waiting for a response
 - Coordinator decides that the transaction must abort.

November 2, 2016

2014-2016 Paul Krzyzanowski

8

Fall 2015 – Question 5

Which of the following is an example of a pure name?:

- a) An email address (e.g., `pxk@cs.rutgers.edu`).
- b) An Ethernet address (e.g., `00:1d:72:ae:cd:c6`).
- c) A file pathname (e.g., `/usr/src/linux-headers/arch/arm/include/asm/syscall.h`).
- d) A URL (e.g., `http://www.cs.rutgers.edu/~pxk/417/exam/index.html`).

- A pure name has no context in it.

November 2, 2016

2014-2016 Paul Krzyzanowski

9

Fall 2015 – Question 6

The DNS (Domain Name System) server for the EDU domain knows:

- a) The list of names and owners of domains under `.edu` but not their IP addresses or name servers.
- b) The IP addresses for all domains under `.edu`, including subdomains (e.g., `cs.princeton.edu`).
- c) The IP addresses for all domains directly under `.edu` and name servers for subdomains.
- d) The domain name servers for each domain directly under `.edu`.

- When you register a domain name, you supply the **domain name registrar** with your ID info and IP addresses for name servers for your domain
- The **domain name registry operator** keeps track of name servers for the domains registered under it; it sends **referrals** to those servers.

November 2, 2016

2014-2016 Paul Krzyzanowski

10

Fall 2015 – Question 7

The two-army problem demonstrates this about communication over faulty lines:

- a) **Reliable communication can never be achieved.**
- b) Reliable communication can be achieved but requires having the receiver send an acknowledgement.
- c) Reliable communication can be achieved but may require several back-and-forth acknowledgements.
- d) Reliable communication can be achieved only in one direction.

- To be 100% certain, you'll need an infinite # of back and forth ACKs

November 2, 2016

2014-2016 Paul Krzyzanowski

11

Fall 2015 – Question 8

In virtual synchrony, a view change is defined as the event that takes place when:

- a) The system hosting the group membership service changes.
- b) A different process starts sending messages.
- c) A message has been received by all group members.
- d) **The group membership changes.**

November 2, 2016

2014-2016 Paul Krzyzanowski

12

Fall 2015 – Question 9

In virtual synchrony, a message is considered stable:

- When the sender promises not send any updates for that message.
- Once it has been delivered to the application.
- Once an incoming message has been fully received and validated to be error-free.
- If it has been received by all group members..**

- A message that has been confirmed to have been received by all group members is stable.

November 2, 2016

2014-2016 Paul Krzyzanowski

13

Fall 2015 – Question 10

In a Paxos system with P proposers and A acceptors, how many proposers need to be alive for the algorithm to work?

- 1.
- $\lceil (P+1)/2 \rceil$ (a majority of proposers).
- $\lceil (A+1)/2 \rceil$ (one proposer for a majority of acceptors)
- A (one proposer for each acceptor)

- Paxos can use any number of proposers – but only one needs to be functioning
 - Usually, one is elected as a *leader* and the others are not used
 - One proposer avoids message rejection due to proposing earlier sequence numbers

November 2, 2016

2014-2016 Paul Krzyzanowski

14

Fall 2015 – Question 11

In a Paxos system with P proposers and A acceptors, how many acceptors need to be alive for the algorithm to work?

- 1.
- $\lceil (P+1)/2 \rceil$ (one acceptor for a majority of proposers).
- $\lceil (A+1)/2 \rceil$ (a majority of acceptors)**
- P (one acceptor for each acceptor)

- Paxos requires a majority of acceptors to be functioning
 - Ensures that at least one acceptor in the current group of acceptors has knowledge from previous proposals

November 2, 2016

2014-2016 Paul Krzyzanowski

15

Fall 2015 – Question 12

Which is NOT a property of a transaction?

- All-or-nothing: a transaction cannot complete partially.
- Available: a transaction cannot keep other transactions from running.**
- Serializable: the result of concurrent transactions must be the same as if they ran in some serial order.
- Permanent: once completed, the results of a transaction are made permanent.

Availability is not part of transactional (ACID) semantics

- "all or nothing" = atomic
- "serializable" = isolated
- "permanent" = durable

November 2, 2016

2014-2016 Paul Krzyzanowski

16

Fall 2015 – Question 13

The first thing that happens in the second phase of the two-phase commit protocol is (assume it will commit):

- The participant responds to the commit query message received in phase 1.
- The coordinator sends a commit message to all participants.**
- The coordinator tells all participants to free resources held by the transaction.
- The coordinator asks the participants if they have completed the transaction, which started in phase 1..

- Phase 1
 - Send a commit query to all participants
 - Get a response so we compute the outcome of the vote
- Phase 2
 - Send a commit/abort request to all participants
 - Get an acknowledgement that the commit/abort completed

November 2, 2016

2014-2016 Paul Krzyzanowski

17

Fall 2015 – Question 14

The purpose of an additional phase in the three-phase commit protocol is to:

- Get unanimous agreement from all participants on whether to commit or abort.
- Tell all participants the outcome of the commit vote before telling them to commit or abort.**
- Have the coordinator receive acknowledgement of whether a commit or abort was successful.
- Log the status of the commit or abort decision to a write-ahead log.

- 3PC enables the use of a recovery coordinator even if some participants are down as well
 - Propagate knowledge of the commit vote BEFORE any participant takes action on it

November 2, 2016

2014-2016 Paul Krzyzanowski

18

Fall 2015 – Question 15

Differing from two-phase locking, in strict two-phase locking:

- A transaction must promise to release every lock that it is granted.
- A transaction must grab a lock for every resource that it needs to access.
- A transaction must release all of its locks only at the end.**
- A transaction cannot acquire locks after it has released a lock.

- 2PL

- Grab locks as you need them
- Release locks once you don't need the resource anymore
- BUT: promise NOT to acquire any lock if you already released a lock

- Strict 2PL

- Don't release any locks until you're done with the transaction
- Avoids **cascading aborts**

November 2, 2016

2014-2016 Paul Krzyzanowski

19

Fall 2015 – Question 16

The reason that strict two-phase locking was introduced as an enhancement of two-phase locking was to:

- Increase the potential amount of concurrency possible.
- Ensure that transaction executions are serialized.
- Avoid cascading aborts when transactions read uncommitted data.**
- Achieve consistency by forcing transactions to get locks for any resources they need.

- With 2PL, a concurrent transaction may read data modified by another transaction if that transaction released its lock.
- That transaction may not have yet committed.

November 2, 2016

2014-2016 Paul Krzyzanowski

20

Fall 2015 – Question 17

The edge chasing used in the Chandy-Misra-Haas algorithm is used to:

- Prevent deadlock by ensuring that a transaction can only wait on resources held by an older transaction.
- Construct a global wait-for graph to determine whether deadlock exists.
- Detect whether there will be a circular dependency on waiting for resources.**
- Ensure deadlock cannot occur by aborting the transaction that is currently using the needed resource.

- Send a probe message to the holder of the resource you want.
- That holder does the same ...
- If the message comes back to you, then you know there will be a circular dependency if you get the lock (hence, deadlock)

November 2, 2016

2014-2016 Paul Krzyzanowski

21

Fall 2015 – Question 18

Deadlock will never occur if:

- A resource can be held by at most one process.
- Processes that hold resources are not allowed to wait for another resource.**
- A resource, once granted, cannot be taken away.
- Two or more processes are waiting for resources held by one of the other processes.

- Four conditions for deadlock

- Mutual exclusion
- Hold & wait
- Non-preemption
- Circular wait

- (a) mutual exclusion – we might get deadlock
- (c) non-preemption – we might get deadlock
- (d) circular wait – we might get deadlock
- (b) hold & wait is NOT permitted – the four conditions will not be met

November 2, 2016

2014-2016 Paul Krzyzanowski

22

Fall 2015 – Question 19

A callback in AFS:

- Enables a server to control bandwidth by telling a client when it can issue a request.
- Informs a client when its requested operation has completed.
- Informs a client that a requested file lock is available.
- Informs a client that a cached file is now invalid.**

- When a client downloads a file, the server makes a **callback promise**
- When another client uploads that file, the server notifies each client that downloaded the file that its cached copy is invalid

November 2, 2016

2014-2016 Paul Krzyzanowski

23

Fall 2015 – Question 20

Which file system is not implemented at the operating system level?

- Network File System (NFS).
- Andrew File System (AFS).
- Microsoft SMB.
- Google File System (GFS).**

- NFS, AFS, and SMB are available to normal apps and look like a native file system.
- GFS was designed to be accessible via user-level APIs

November 2, 2016

2014-2016 Paul Krzyzanowski

24

Fall 2015 – Question 21

Coda's client modification log enables:

- Reintegration.
- Read-write replication.
- Consistent caching.
- Disconnected operation.

- Coda supports disconnected operation (when the AVSG = ∅)
- At that time any file updates are logged in the client modification log (CML)
 - When the system is re-connected to the network
 - The CML is played back to propagate changes to the server
- (d) is correct in that the CML is used during disconnected operation.
 - However, the CML is not needed to access files while disconnected
 - (a) is the better answer since reintegration is the specific part of disconnected operation that the CML is designed to handle

November 2, 2016

2014-2016 Paul Krzyzanowski

25

Fall 2015 – Question 22

Credit-based flow control in SMB:

- Allows the server to control the load from each client.
- Provides a billing system for clients that sue a server.
- Improves the level of concurrency in file system operations.
- Enables the client to cache more data.

- The server gives each client a number of "credits"
- Each client request consumes a credit
- When the client runs out, it waits to get more credits from the server

November 2, 2016

2014-2016 Paul Krzyzanowski

26

Fall 2015 – Question 23

Which file system was designed to support concurrent appends to a file?.

- NFS.
- AFS.
- SMB.
- Coda.

- AFS & Coda: session semantics; no concurrent file modifications of any kind
- NFS: stateless design doesn't provide capability for appends

November 2, 2016

2014-2016 Paul Krzyzanowski

27

Fall 2015 – Question 24

An SMB oplock gives clients:

- The possibility of caching file attributes.
- The ability to lock a remote file.
- The ability to check out a file for disconnected use.
- The ability to synchronize local copies of remote files.

- Oplocks (similar to DFS tokens) give clients permission to cache file data
 - Examples
 - Level 1: exclusive access – cache attributes, read-ahead, write-behind, cache lock data
 - Level 2: multiple readers, one writer – writer can cache file attributes & read-ahead

November 2, 2016

2014-2016 Paul Krzyzanowski

28

Fall 2015 – Question 25

Which file systems place file metadata on separate servers from file data?

- GFS and Chubby.
- Dropbox and GFS.
- Chubby, Dropbox, and GFS.
- AFS and GFS.

- Each AFS, NFS, SMB, DFS server stores complete file systems: data & metadata
- Chubby is a single-server file system that happens to be replicated
 - Each server contains all data
- GFS
 - Separate **master** server stores names and chunk IDs associated with each name
 - Actual data (which may be massive) distributed among **chunkservers**

November 2, 2016

2014-2016 Paul Krzyzanowski

29

Fall 2015 – Question 26

Why did Dropbox add notification servers to their architecture?.

- To provide a mechanism to alert administrators when problems arise.
- Dropbox servers needed to be notified when the file has been uploaded to Amazon servers.
- To ensure that files are consistent among multiple clients.
- To reduce load from clients polling the servers..

- Notification servers send notifications of changes to clients
- Before that, each client would poll the server

November 2, 2016

2014-2016 Paul Krzyzanowski

30

Fall 2015 – Question 27

Chubby does not support:

- a) **Partial file reads.**
- b) Event notifications.
- c) Whole file uploads.
- d) File locking.

- Chubby was designed to store small files (e.g., configuration data), provide locking, and notifications
- Since files are small, file I/O is only entire file reads & writes

November 2, 2016

2014-2016 Paul Krzyzanowski

31

Fall 2015 – Question 28

AS GFS master:

- a) Identifies the addresses of all the name servers that keep track of file names and data.
- b) Accepts every client request and routes it to the appropriate server.
- c) **Stores all the names in the file system along with the location of their data**
- d) Receives file write operations from clients that are then propagated to replicas.

- The GFS master stores metadata
- File content is distributed – and replicated – among chunkservers

November 2, 2016

2014-2016 Paul Krzyzanowski

32

Fall 2015 – Question 29

How does Bigtable manage the growth of a table?

- a) Individual table cells are distributed among a large set of servers.
- b) An entire column of a table can be migrated to a different server.
- c) Each new row is allocated to one server in a large set of servers based on a hash of its key.
- d) **A table is split along rows into sub-tables.**

- As a table gets bigger, it is split along a row into sub-tables (**tablets**)
 - Rows in bigtable are always kept sorted by the row's key
- (a) No. All columns of a row are kept together & adjacent rows are kept together
 - (b) No. Columns aren't split out and don't move.
 - (c) No. Rows are not distributed among servers; they stay together

November 2, 2016

2014-2016 Paul Krzyzanowski

33

The End

November 2, 2016

2014-2016 Paul Krzyzanowski

34