Distributed Systems

2016 Exam 1 Review

Paul Krzyzanowski

Rutgers University

Fall 2016

Question 1

Why does it not make sense to use TCP (Transmission Control Protocol) for the Network Time Protocol (NTP)?

TCP offers reliable delivery **but** via *retransmission*.

TCP also may delay the transmission of data.

These factors may lead to jitter – variations in the delay, which will make the assumption that the timestamp is generated in the middle invalid

Bad answers:

- *TCP has longer latency*
- *TCP has high overhead*

Question 2

What is a benefit of lease-based garbage collection over reference count based garbage collection?

It's not fault tolerant.

If a client process dies or exits without properly decrementing reference counts, the object would not get deleted.

Question 3

(a) Explain the role of an interface definition language in remote procedure calls

Describes the programming interface for remote (functions, data types, parameters, return values) so that stub functions can be generated.

Bad answer: creates stubs

(b) Explain the purpose of marshaling in remote procedure calls.

Convert a list of parameters into a sequence of bytes (a serialized format).

Question 4

A client requests time from a server. It takes 80 ms (0.080 seconds) to get the response, which contains a timestamp of 5:23:30.000. Using Cristian's algorithm, to what time does the client set its clock?

$T_{new} = T_s + ½$ delay

5:23:30.000 + (0.080÷2) = 5:23:30.000 + 0.040 = 5:23:30.040

Question 5a

The diagram on the right shows three multicasts to a group of three processes. The sending of each multicast message is a single event.

(a) Assign vector timestamps to each event using a sequencing of (P_0, P_1, P_2). The first event on each process gets a sequence number of 1.

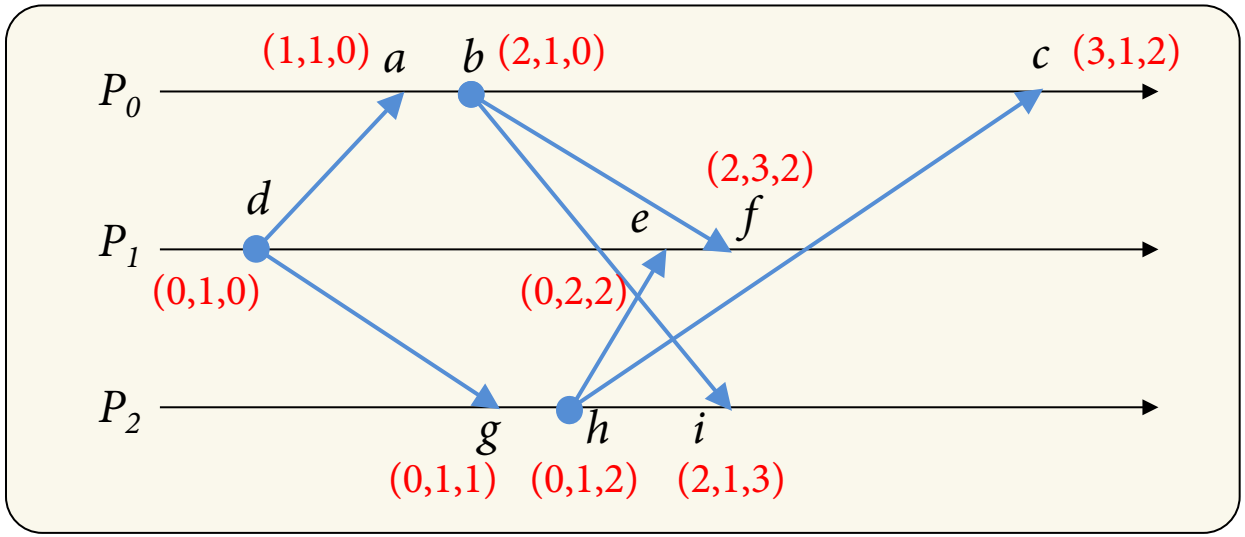

Question 5b

(b) Based on the vector clock values, which events on concurrent with event b?

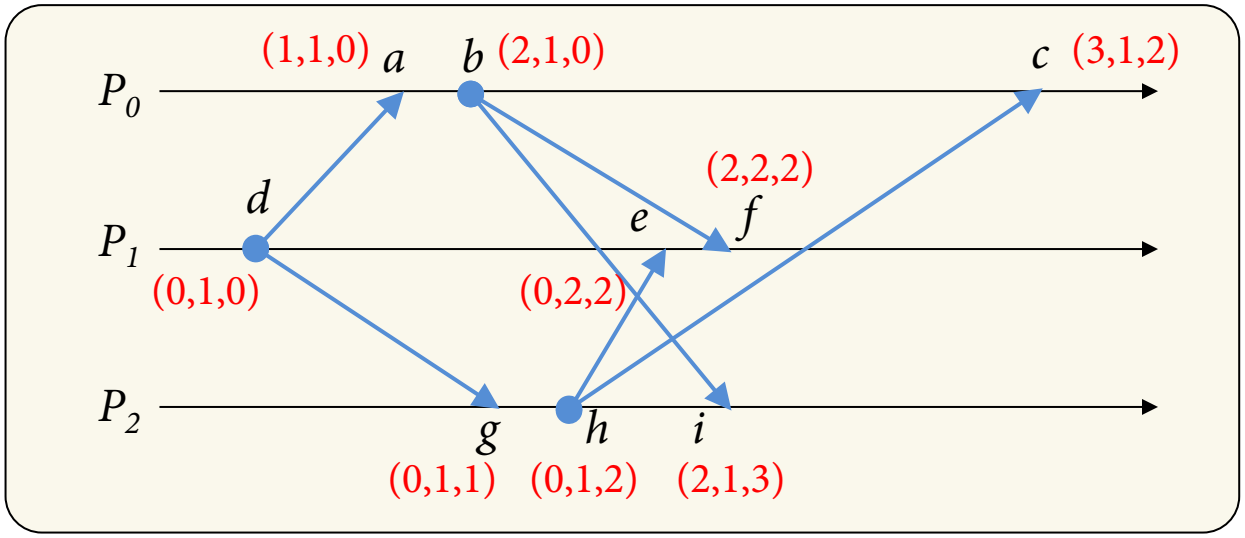

Concurrent with *b: e, g, h*

b: (2, 1, 0)	e: (0, 2, 2)	(2, 1, 0) $\not\leq$ (0, 2, 2)
	g: (0, 1, 1)	(2, 1, 0) $\not\leq$ (0, 1, 1)
	h: (0, 1, 2)	(2, 1, 0) $\not\leq$ (0, 1, 2)

Question 6

Metcalfe's Law, named after the inventor of Ethernet, tells us that the value of a network increases as:

a) The speed of the network increases.
b) The cost of networking hardware decreases.
c) The amount of data sent per person increases.
d) The number of connected users increases.

Metcalfe's law states that the value of a telecommunications network is proportional to the square of the number of connected users of the system

Question 7

A directory in a cache coherent NUMA architecture is:

a) A tree-structured mapping of processes to processors.
b) A per-processor table that keeps track of which other processors have cached copies of regions of memory.
c) A global structure that enables the operating system to switch the connections of processors to memory.
d) A listing of files that are accessible by each processor.

- Table for blocks of memory that records the caching state of each block: which CPU has the latest version

Question 8

A snoopy cache can speed up:

a) Memory reads.
b) Memory writes.
c) Both memory reads and writes.
d) Neither memory reads or writes.

- Writes are write-through and go to main memory (dealing with possible bus contention)
- Reads may be served directly from cache (if data is cached)
- Generally, reads outnumber writes, so we win

Question 9

Ethernet communication uses:

a) Time division multiplexing.
b) Frequency division multiplexing.
c) Channel access via token passing.
d) Random access.

Question 10

IP is designed to be implemented over:

a) Unreliable connectionless networks.
b) Reliable connectionless networks.
c) Unreliable connection-oriented networks.
d) Reliable connection-oriented networks.

Question 11

Port numbers are used in:

a) IP.
b) UDP only.
c) UDP & TCP.
d) TCP only.

- Port numbers are a transport-layer construct to identify socket endpoints.

Question 12

TCP cannot provide:

a) Reliable delivery.
b) In-order delivery.
c) Constant latency.
d) Congestion control.

- Reliable delivery = retransmit lost or damaged data
- In-order delivery = each segment contains a sequence number
- Congestion control = reduce transmission rate (window size) if packet loss is detected

- TCP cannot control how long it takes to deliver a packet.

Question 13

To enable a TCP socket to receive incoming connections, the following system call should be used:

a) bind
b) listen
c) accept
d) recvfrom

- *listen* sets up a socket to be able to receive incoming TCP connections.
- *bind*: assigns an address and port # to a socket; used on both client & server
- *accept*: wait for an incoming connection on a listening socket.
- *recvfrom*: receive a UDP packet.

I will also accept ©

Question 14

An idempotent function:

a) Is any function that can accept parameters in a serialized format.
b) Can be stored in byte code format and downloaded by a client.
c) Is a remote function rather than a local function.
d) Can be called multiple times without side-effects.

Question 15

A key advantage of multi-canonical marshaling is that it:

a) Enables a set of data to be sent to multiple servers simultaneously.
b) Allows clients and servers to have different processor architectures.
c) Reduces the overall amount of data conversion that needs to be performed.
d) Allows clients to communicate directly with servers without routing messages through a proxy.

- Ideally, neither client nor server will have to convert data to a local format.

Question 16

DCE RPC improved ONC (Sun) RPC by adding:

a) An interface definition language.
b) A cell directory server to look up RPC services.
c) Support for distributed objects.
d) An RPC compiler (stub generator).

- ONC RPC had (a) and (d)
- DCE RPC did not have object support (c). Microsoft added this in an enhancement of DCE RPC, called MS-RPC.

Question 17

A surrogate process in Microsoft's COM+:

a) Runs on the client and loads client-side stub objects.
b) Runs on the client and receives requests if the server cannot be reached.
c) Runs on the server and starts RPC services at boot time.
d) Runs on the server and loads objects based on client requests.

Question 18

In a group of two computers, a client's local clock reads 6:27:10. Using the Berkeley clock synchronization algorithm, to what value does the client set its time if the server's clock reads 6:28:30? Ignore message transit times.

a) 6:27:50
b) 6:28:30
c) 6:29:10
d) 6:29:50

- There is no concept of a server that has the "true time"
- Server = master; client = slave
- Berkeley synchronization averages out all time values
- (6:27:10 + 6:28:30) / 2 = 6: 27: (10 + 90)÷2 = 6:27:(100÷2) = 6:27:50

Question 19

An NTP synchronization subnet is:

a) A high-speed network that is dedicated to clock synchronization.
b) The set of servers that offers clock synchronization services.
c) Reserved capacity dedicated to clock synchronization in an existing network.
d) Any network over which an NTP server continuously sends time broadcasts.

Question 20

With the use of Lamport timestamps, we can achieve:

a) Global ordering.
b) Total ordering.
c) Partial ordering.
d) Sync ordering.

- Lamport timestamps define partial ordering
- We can convert that to total ordering by adding a per-process ID as a secondary sorting value <*L, P*> (for example, a decimal value)
- The resultant messages have unique timestamps but do not enable you to identify causal relationships.

Question 21

Atomic multicast differs from reliable multicast because atomic multicast

a) Is much faster since it uses the network hardware to ensure reliability.
b) Only requires partial ordering.
c) Does not need to deliver messages reliably.
d) Accounts for system failures.

Question 22

Lamport's mutual exclusion algorithm is an example of a:

a) Centralized algorithm.
b) Token-based algorithm.
c) Contention-based algorithm.
d) Random-selection algorithm.

- Centralized = contact a central service
- Token-based = pass a token; access resource if you have the token
- Contention-based = everyone who wants a resource asks for it; resolve multiple requests by comparing timestamps
- Random selection = ???

Question 23

Which mutual exclusion algorithm does not require knowledge of group members?

a) Centralized.
b) Token ring.
c) Lamport.
d) Ricart and Agrawala..

(b) Requires contacting next process in a logical ring (and knowing your position in the ring)

(c) Requires sending a request to all group members

(d) Same as (c)

(a) Only need to know a mutex server – no need to know a single group member

Question 24

The Chang & Roberts ring algorithm improves the ring election algorithm by:

a) Stopping redundant elections when possible.
b) Using a centralized coordinator to decide on election results.
c) Using total ordering of election messages.
d) Using reliable message delivery.

Chang & Roberts improves the ring algorithm in two ways:

1. Does not send a list but processes a vote at each member.
2. Kills off an election message if a receiving process is already participating in an election and this new message is from a smaller process ID than its own.

Question 25

The Bully election algorithm chooses:

a) The first process to notice a dead leader.
b) The highest living process ID.
c) The process that sends the most election messages during the election.
d) The process that gets the majority consensus from the group.

Question 26

On multiprocessor systems, processors share a system clock.

TRUE

By definition, multiprocessors have:

- Shared memory
- Shared clock
- All-or-nothing failure

Question 27

A switched network connection enables greater scalability than a bus-based one.

TRUE

Question 28

Ethernet transmission is unreliable.

TRUE

Question 29

TCP is a network-layer (layer 3) protocol.

FALSE

- IP is a network layer protocol
- TCP and UDP are transport layer (layer 4) protocols

Question 30

Explicit typing identifies the data elements in a message.

TRUE

Question 31

If the Lamport timestamp associated with event *A* is less than the Lamport timestamp associated with event *B*, we can conclude that event *A* happened before event *B.*

FALSE

With Lamport timestamps, if $A \rightarrow B$ then $L(A) < L(B)$
However, if $L(A) < L(B)$, you cannot tell if $A \rightarrow B$ since *A* & *B* may be concurrent events.

Question 33

A hold-back queue is used to resequence messages at the receiver.

TRUE

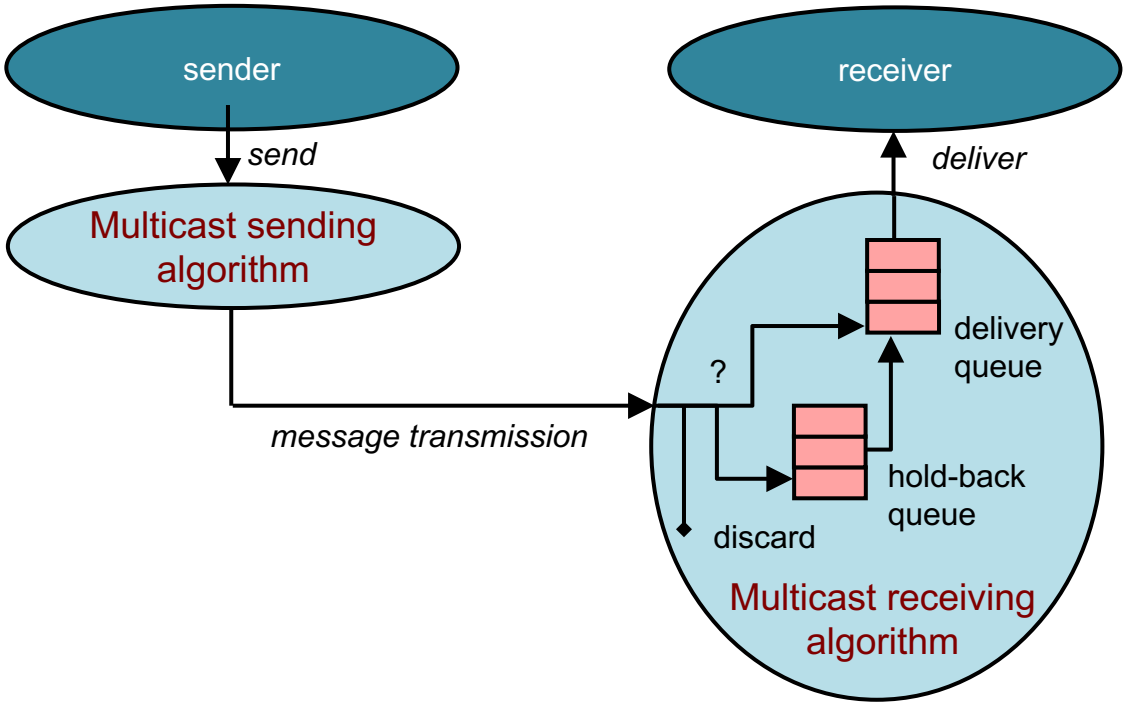

The End