Distributed Systems
Exam 2 Review: Fall 2012

Paul Krzyzanowski

pxk@cs.rutgers.edu

Question 1

Reference count based distributed garbage collection is a more efficient use of network resources than lease-based garbage collection. Explain the advantages of lease-based garbage collection and why it won over the reference counting approach

It's more fault tolerant. With reference counting, you still had to take care of the case of a client abnormally terminating (or improper counting).

Question 2

Identify two reasons why you might want to use a higher replication factor for files in GFS.

1. High availability

2. Load balancing

3. Distribution for geographic proximity

Question 3

What compromise must be made in a distributed system with replicated data if you must have high availability and partition tolerance?

Consistency

Brewer's CAP theorem states that you can have at most two out of three of:

consistency + availability + partition tolerance

Question 4

Contrast the client caching approaches of NFS, AFS, and SMB with oplocks

NFS:
   Short-term caching, periodic checking with the server, ambiguous semantics

AFS:
   Long-term caching. Server sends invalidations when changes occur on the server. Session semantics.

SMB/oplocks:
   Oplocks are tokens that control client cache behavior.
   Long-term caching & read-ahead if nobody is modifying the data.
   Write-behind OK if nobody else is reading the data. Otherwise no caching.

Part II: 5-7

5.   ONC (Sun) RPC provides the ability to:
(a)   Use XML as a transport.
(b)   Start up the server process on demand.
(c)   Perform distributed garbage collection.
(d)   Have multiple versions of a function at the server.

6.   A multi-canonical marshaling format
(a)   Provides greater efficiency because both sides usually won't have to convert data.
(b)   Is a more compact way of representing data over a network.
(c)   Encodes data concurrently into both binary as well as text formats.
(d)   Allows one message to be sent to multiple servers.

7.   For RPC, a DCE cell directory server allows:
(a)   A client to find out on what server an interface is available.
(b)   A client to find the port number of a service on a specific machine.
(c)   A server to send callbacks to clients.
(d)   An object to be distributed among multiple servers.

Part II: 8-10

8. Java's Serializable interface:
(a) Allows an object's data to be converted to a sequence of bytes.
(b) Creates a remote reference for an object.
(c) Enforces concurrency control to ensure that concurrent accesses to an object are serialized.
(d) Creates client and server stubs for an object.

9. Compared with SOAP, REST:
(a) Is based on remote method calls.
(b) Identifies resources in the URL of an HTTP command.
(c) Uses XML for creating a message within the HTTP message.
(d) Is not tied to a single language.

10. Which distributed mutual exclusion algorithm does not require a participant to know anything about the composition of the group?
(a) Centralized
(b) Lamport
(c) Ricart and Agrawala
(d) Token Ring

Part II: 11-13

11. Which distributed mutual exclusion algorithm does not result in a higher number of requests (and hence network traffic and system load) when many processes want a resource at the same time?
(a) Centralized
(b) Lamport
(c) Ricart and Agrawala
(d) Token Ring

12. Which mutual exclusion algorithm creates replicated request queues on each process?
(a) Centralized
(b) Lamport
(c) Ricart & Agrawala
(d) Token Ring

13. With DCE and Microsoft RPC, the Unique Universal Identifier (UUID) is used to uniquely identify:
(a) A client.
(b) An interface to a set of procedures.
(c) A communication session.
(d) A server machine.

Part II: 14-16

14. Chubby presents itself to clients as this service:
(a) Centralized mutual exclusion
(b) Hierarchical mutual exclusion
(c) Token-based mutual exclusion
(d) Contention-based mutual exclusion.

15. Differing from a token-based algorithm, a contention-based mutual exclusion algorithm relies on:
(a) Reliable message delivery
(b) Unique Lamport timestamps in request messages
(c) A coordinator process
(d) Constructing a logical ring of processes.

16. The Chang & Roberts algorithm optimizes the ring algorithm by:
(a) Using UDP instead of TCP for message delivery .
(b) Testing higher-numbered processes first
(c) Diving the ring into sub-rings and using a divide-and-conquer approach
(d) Stopping multiple election messages from circulating.

17. A group of 10 processes (P0..P9) uses the bully algorithm to pick a leader with the highest numbered process ID. Process 6 detects the death of process 9 and holds an election. How many election messages are sent in the system as a whole (include failed messages to process 9)?

(a) 3
(b) 6
(c) 10
(d) 45

3 messages from P_6

Part II: 17

17. A group of 10 processes (P0..P9) uses the bully algorithm to pick a leader with the highest numbered process ID. Process 6 detects the death of process 9 and holds an election. How many election messages are sent in the system as a whole (include failed messages to process 9)?

(a) 3
(b) 6
(c) 10
(d) 45

3 messages from P_6
2 messages from P_7

Part II: 17

17. A group of 10 processes (P0..P9) uses the bully algorithm to pick a leader with the highest numbered process ID. Process 6 detects the death of process 9 and holds an election. How many election messages are sent in the system as a whole (include failed messages to process 9)?

(a) 3
(b) 6
(c) 10
(d) 45

3 messages from P_6
2 messages from P_7
1 message from P_8
6 total messages

Part II: 18-20

18. The two-army problem demonstrates that reliable communication with unreliable communication links:
(a) Can be achieved with n2 message exchanges for a system of n processes.
(b) Can be achieved with a simple message acknowledgement protocol.
(c) Requires a two-way acknowledgement.
(d) Cannot be achieved with 100% certainty.

19. Paxos reaches agreement when:
(a) All proposers agree on a value to send to the acceptors.
(b) All acceptors agree to a proposed value.
(c) The majority proposers agree on a value to send to the acceptors.
(d) The majority of acceptors agree to a proposed value.

20. A hierarchical lease:
(a) Allows clients to get both exclusive and shared leases.
(b) Allows multple clients to request leases for parts of an object.
(c) Allows a client that has a lease for an object to get a lock for that object.
(d) Uses an elected coordinator to manage a set of leases.

Part II: 21-23

21. The purpose of the first phase in a two-phase commit protocol is to:
(a) Tell all processes participating in the transaction to start working on the transaction.
(b) Wait for all processes to commit their transactions.
(c) Find out whether processes are still working on the transaction.
(d) Get consensus from all processes participating in the transaction on whether to commit.

22. A three-phase commit protocol:
(a) Improves the consistency of the two-phase protocol.
(b) Tells the coordinator of the final commit vs. abort outcome.
(c) Sets time limits for the protocol.
(d) Gives cohort processes the ability to authorize the commit.

23. Paxos avoids the "split brain" problem that can arise when a network is partitioned by:
(a) Placing proposers and acceptors on the same machine.
(b) Placing acceptors and learners on the same machine.
(c) Requiring over 50% of acceptors to be accessible.
(d) Using a two-phase commit protocol for each incoming request.

Part II: 24-26

24. Which condition is not necessary for deadlock?
(a) Mutual exclusion (a resource can be held by only one process).
(b) Hold and wait (processes holding resources can wait for another resource).
(c) Preemption (a resource can be taken away from a process).
(d) Circular wait (a cycle of resource holding and waiting exists).

25. False deadlock is caused by:
(a) Releasing one resource before waiting on another.
(b) Waiting on a resource before releasing one that is already held.
(c) Improper message ordering at the coordinator.
(d) Two processes competing to grab the same resource.

26. The wait-die algorithm is a technique of deadlock prevention that:
(a) Ensures that circular wait will not exist.
(b) Relaxes the use of locking to avoid waiting on resources.
(c) Introduces time-outs if a process cannot get a resource within a time limit.
(d) Schedules transactions in a serial order so that only one runs at a time.

Part II: 27-29

27. Compared with two-phase locking, strict two-phase locking:
(a) Guarantees that there is only one growing and one shrinking phase per transaction.
(b) Ensures that a transaction cannot access data written by an uncommitted transaction.
(c) Uses a two-phase commit protocol to get a lock.
(d) Makes the use of resource locks mandatory.

28. Optimistic concurrency control schemes usually allow multiple transactions to run concurrently and:
(a) Grab locks for resources they need.
(b) Avoid the use of locks.
(c) Use a distributed consensus algorithm to agree on a commit order.
(d) Replicate data for fault tolerance.

29. While NFS was originally designed to be stateless, state was first added to support:
(a) File locking.
(b) Coherent client-side caching.
(c) RPC-based remote file access.
(d) File replication.

Part II: 30-32

30. DFS tokens are most comparable to:
(a) *Shared locks* and *write locks* in concurrency control.
(b) The token in a token-ring mutual exclusion algorithm.
(c) Getting consensus in a Paxos leader election algorithm.
(d) A callback promise in AFS.

31. Commands sent to a Chubby cell:
(a) Are load balanced among the machines in the cell.
(b) Must be sent to and are processed by the current master.
(c) Are executed by whichever machine gets the request.
(d) Go to the master and are then forwarded to whichever Chubby replica holds the needed data.

32. Which of these operations is most efficiently implemented on a large-scale GFS system?
(a) Read one 1 TB file.
(b) Read 1 million 1 MB files.
(c) Write one 1 TB file.
(d) Write 1 million 1 MB files.

Part II: 33

33. HDFS (Hadoop File System) is closely patterned after GFS (Google File System) but does not support:
(a) Concurrent appends.
(b) Partial file reads.
(c) Redundancy for file storage.
(d) Distributing a file's contents across multiple storage servers.

The End

10/29/2015