

Lecture 17: March 31

*Lecturer: Prashant Shenoy**Scribe: Nicholas Monath*

17.1 Replication & Consistency Continued

17.1.1 Primary-based Protocols

Primary-based protocols assume that there is a coordinator/primary replica for each data item, which is in charge of implementing the consistency method.

17.1.2 Replicated-Write Protocols

Here, there is no notion of a coordinator/primary replica, and a write can take place at any replica. All replicas will take part in the consistency method. Any replica is allowed to update and so consistency is more complicated to achieve.

17.1.3 Remote-Write Protocols

This is a primary-based protocol. Each data item has one node assigned to be the primary replica for that data item. Clients send write requests for a particular data item to their local replicas, which forward the request to the node which is the coordinator for that data item. The coordinator then enforces consistency, using any of the previously discussed techniques. Reads, assuming caching is not used, are implemented in the same way as writes. File Systems use this approach often.

17.1.4 Local-Write Protocols

This is another primary-based protocol. The read or write requests of clients are sent to a local replica. That replica now becomes the primary replica for the data item of the request. It is essentially a performance optimization of Remote-Write Protocols. Despite the performance optimization of subsequent requests taking place on a near-by replica, we have the limitation of needing to keep track of the primary replica for each data item and the primary replica might change often.

17.1.5 Quorum-Based Protocol

A replicated-write protocol that uses *voting* to request and acquire permissions from replicas. We define two quorums, a *read quorum*, notated N_r , and a *write quorum*, notated N_w , where N is the number of servers. The quorums represent the number of nodes that agree on the value for a read/write request. We make two constraints on the quorums:

1. $N_r + N_w > N$
2. $N_w > \frac{N}{2}$

A read or write request can only be processed if the above constraints are satisfied. This is also known as **Gifford's Quorum-Based Protocol**.

A special case of this protocol is known as *read once, write all* (ROWA) where $N_w = N$.

17.1.6 Replica Management

Issues such as Geo-locations of the replicas, number of replicas etc. have to be taken care of. This will be discussed in further lectures.

17.2 Fault Tolerance

There are many reasons for a system or service to crash. In a distributed system if a machine fails, other machines in the system can share the responsibility of the crashed machine. Some nodes can fail without the whole system going down. Fault tolerance deals with the semantics and the mechanisms of dealing with failed machines. As the number of nodes in the system increases (approaches infinity), the probability of a failure approaches 1. Fault tolerance system should provide services despite faults, such as transient faults, intermittent faults, and permanent faults.

17.2.1 A Perspective

Systems are inherently unreliable. OS crashes frequently, buggy software, unreliable hardware, software/hardware incompatibilities. With more “novice” users of computer system, we want to make computing systems more reliable. We will do so by adding an additional layer of reliability.

17.2.2 Basic Concepts

Requirements for dependable systems:

- **Availability:** A system as a whole should be available for use at any given time. Measured by the percentage of time a system is available. 99.999% (“five 9’s”) is desirable.
- **Reliability:** System should run continuously without failure
- **Safety:** temporary failures should not result in catastrophic failures.

- Maintainability: A failed system should be easy to repair

There are three kinds of fault tolerance:

- Transient faults: A “one time” fault, appears once and never appears again. No noticeable reason – a bit flipped in RAM etc.
- Intermittent fault: Appears every once and a while and then goes away. Difficult to reproduce or track
- Permanent fault: Faults such as system crashes.

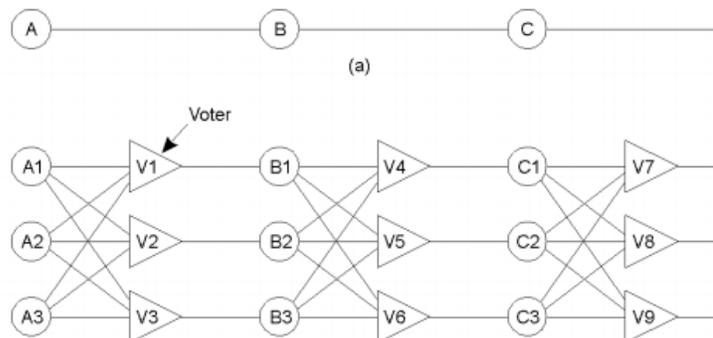
17.2.3 Failure Models

There are four types of failures:

- Crash failure: A server halts but is working correctly before the halt
- Omission failure (Receive/Send Omission): Server fails to respond to incoming requests or fails to send back a response
- Timing failure: The latency of server request causes a problem, i.e. it causes other processes to timeout
- Response failure (value failure, state transition failure): Server sends a response but the value of the response is incorrect or the server deviates from the correct flow of control.
- Arbitrary failure: A server may produce arbitrary responses at arbitrary times. Sometimes it works as expected, other times it does not. No way to predict when it is behaving correctly. Related to Byzantine faults. The most difficult kind of failure to deal with.

17.2.4 Failure Masking by Redundancy

Triple modular redundancy (shown below) is a way to design hardware that is resilient to crash failure and arbitrary failure. The top circuit in the diagram is not resilient to either kind of failure. In the lower diagram, the voter nodes pass along the majority of the responses of the incoming nodes. The voter is replicated as it too might be faulty.



17.3 Reliable One-One Communication

In lecture 3, we discussed RPC communication in the presence of crash failures. Now, we discuss how to achieve reliable communication when packets can be dropped by the network. To do so nodes send acknowledgements back to the sender of the message. However this does not scale well when multicasting a message to many nodes; the sender becomes a bottleneck. Negative acknowledgements can be used in multicasting. Each packet is given a sequence number. A node can determine when it has missed a packet if it sees a gap in the sequence numbers that it receives.