CMPSCI 677 Operating Systems	Spring 2023

<div align="center">

Lecture 15: April 3

</div>

Lecturer: Prashant Shenoy *Scribe:* **Y. Vayunandhan Reddy, Yamini Kashyap**

15.1 Overview

This section covers the following topics:

Distributed Snapshots: Distributed Snapshot Algorithm

Leader Election: Bully Algorithm, Ring Algorithm

Distributed Locks: Centralized, Decentralized, Distributed algorithms

15.1.1 Distributed Snapshot

A simple technique to capture is to assume a global synchronized clock and freeze all machines at the same time and capture whatever that are in all memories at the exactly same time and capture everything that is in transit. But there is no global clock synchronization. Distributed snapshotting is a mechanism that gives a consistent global state without a global clock synchronization. This is done so that in case one or more components of the distributed system fail, the system as a whole can restart from a snapshot of all the components rather than starting from scratch as that would waste all the CPU cycles that have already been consumed in the processing before the system failed.

15.1.2 Distributed Snapshot Algorithm

Assume each process communicates with another process using undirectional point-to-point channels (e.g., TCP connections). Any process can initiate the snapshot algorithm. When a process initiates the algorithm, it will first checkpoint its local state, and then send a marker on every outgoing channel. When a process receives a marker, it will have different actions depending on whether it has already seen the marker for this snapshot. If the process sees a marker at the first time, it will checkpoint its state, send markers out and start saving messages on all its other channels. If the process sees a marker for the snapshot at the second time, it will stop saving messages for the corresponding channel from which the marker comes. A process will finish a snapshot until it receives a marker on each incoming channel and processes them all.

For example, in Figure 15.1, process 1, 2 and 3 communicates with each other through the duplex TCP connections. When process 1 initiate a snapshot, it will first save its memory contents (state) into the disk, and then send a marker (in blue) to 2 and 3. It will also start to save the incoming messages(in red) (TPC buffers) from 2 and 3. When 2 receives a marker from 1, it will start to checkpoint state, send a marker out to 1 and 3, and start saving messages from 3. Assume 3 sees the marker sent from 1 first, it will checkpoint its state, sends out a marker to 1 and 2, and start saving messages from 2. 1 will stop saving messages for 2 or 3 until a marker from 2 or 3 arrives. 2 will stop saving messages for 3 until a marker from 2 arrives. And P3 will stop saving messages for 2 until a marker from 3 arrives. Each process will finish this snapshot when it sees a marker from every incoming channel. Thus, a distributed snapshot captured.

Figure 15.1: Distributed Snapshot Example

Figure 15.2: Snapshot algorithm example.

15.1.3   Snapshot Algorithm Example

Q decides to take a snapshot. Consider Figure 15.2. b) It receives a marker for the first time and records its local state; c) Q records all incoming message; d) Q receives a marker for its incoming channel and finishes recording the state of the incoming channel. First incoming marker says start recording and all other incoming messages say stop recording.

Question: Why is it required to save messages on the incoming channels?
Answer: The process will save all the messages that were in transit and replay them when the system restarts as if those messages have just been received. If these messages are not saved, the information in them will get lost when the system restarts.

Question: After a process has saved its state but has not received the second marker, can it continue sending messages?
Answer: Yes, the process continues execution normally. The process of taking snapshots does not interfere with the regular execution of the system. Snapshotting is only there to determine at a certain point of time, which states and messages to store in the disk.

Question: Are there different types of markers, such as a START marker and a STOP marker?

Answer: There are no differences in the markers sent between processes. However, in terms of interpretation, the first marker received on any channel is both a START and a STOP marker for the channel it comes from and all other markers received across all channels are STOP markers for corresponding channels.

Question: What if there are some messages in the channel before the marker?
Answer: The messages before the marker are not part of the checkpoint.

Question: How are messages replayed?
Answer: All messages in transit between receiving the first and second marker get stored on disk. And when the process is reinitialized, the corresponding socket buffers are populated with the messages that are part of the snapshot.

Question: What happens if another checkpoint starts when a checkpoint is already going on?
Answer: You can have as many checkpoints going on simultaneously as you want. However, each checkpoint will have a unique Id associated with it so that all artifacts belonging to that checkpoint can be distinguished. All marker messages exchanged between processes are required to contain the checkpoint Id along with them.

Question: How to ensure that the ID of a checkpoint is unique?
Answer: The initiator process can create ID by using its IP address, or process ID or any other information it likes. The ID is not necessarily a number, it can be string as well so there are multiple options to choose unique snapshot IDs.

15.2   Termination Detection

This involves detecting the end of a distributed computation. We need to detect this since no process can just exit after completing its tasks if there is atleast one process that is still running. Let sender be the predecessor, and receiver be the successor. There are two types of markers: Done and Continue. After finishing its part of the snapshot, process Q sends a Done or a Continue to its predecessor. Q can send a Done only when:

- All of Q's successors(and children processes) send a Done

- Q has not received any message since it check-pointed its local state and received a marker on all incoming channels

- Else send a Continue

Computation has terminated if the initiator receives Done messages from everyone.

Question: Isn't it inefficient for all the processes to keep running even though only a subset of them have not finished?
Answer: It is necessary for all remaining processes to be alive as well because the unfinished process(es) might send them a message to obtain a result without which they cannot terminate. Since we don't know beforehand what is remaining to be done, we cannot exit other processes in a distributed system.

15.3   Leader Election

Many tasks in distributed systems require one of the processes to act as the *coordinator*. Election algorithms are techniques for a distributed system of N processes to elect a coordinator (leader). An example of this is the Berkeley algorithm for clock synchronization, in which the coordinator has to initiate the synchronization and tell the processes their offsets. A coordinator can be chosen amongst all processes through leader election.

15.3.1   Bully Algorithm

The bully algorithm is a simple algorithm, in which we enumerate all the processes running in the system and pick the one with the highest ID as the coordinator. In this algorithm, each process has a unique ID and every process knows the corresponding ID and IP address of every other process. A process initiates an election if it just recovered from failure or if the coordinator failed. Any process in the system can initiate this algorithm for leader election. Thus, we can have concurrent ongoing elections. There are three types of messages for this algorithm: *election*, *OK* and *I won*. The algorithm is as follows:

1. A process with ID i initiates the election.

2. It sends *election* messages to all process with ID $> i$.

3. Any process upon receiving the election message returns an OK to its predecessor and starts an election of its own by sending *election* to higher ID processes.

4. If it receives no OK messages, it knows it is the highest ID process in the system. It thus sends *I won* messages to all other processes.

5. If it received OK messages, it knows it is no longer in contention and simply drops out and waits for an *I won* message from some other process.

6. Any process that receives *I won* message treats the sender of that message as coordinator.

An example of Bully algorithm is given in Figure 15.3. Communication is assumed to be reliable during leader election. If the communication is unreliable, it may happen that the elected coordinator goes down after it being elected, or a higher ID node comes up after the election process. In the former case, any node might start an election process after gauging that the coordinator isn't responding. In the latter case, the higher ID process asks its neighbors who is the coordinator. It can then either accept the current coordinator as its own coordinator and continue, or it can start a new election (in which case it will probably be elected as the new coordinator). This algorithm runs in $O(n^2)$ time in the worst case when lowest ID process initiates the election. The name bully is given to the algorithm because the higher ID processes are bullying the lower ID processes to drop out of the election.

Question: Can multiple processes elect their own leaders? How do you decide who is the leader finally?
Answer: There are no subsets or groups that we are dealing with here. All processes are in the same group and they have a consensus that the process with the highest ID will be the leader. If there are multiple groups, we are only worried about the leader in that group.

Question: Can the leader be a single machine?
Answer: We don't actually care about number of machines in the system. Here, it is a process that is chosen as the leader, regardless of number of machines and number of processes running on each machine.

Figure 15.3: Depiction of Bully Algorithm

Question: Are we assuming here that every node knows about every other node in the system and that they can communicate with each other without a middle-man?
Answer: Yes, here we assume that all processes are aware of the presence of all other processes and can communicate with each other.

Question: Does the chosen leader also need to know about the presence of all other processes in the system?
Answer: That depends on the type of work that is expected out of the leader. For instance, if it is a time-server, then it only needs to answer time queries and not anything pertaining to other processes but if the leader's job is distributed lock synchronization, then it needs to know the state and other information of all the other processes. Responsibilities of the leader are not a part of the leader election mechanism.

Question: How do you tell whether a process(leader in this case) has really failed/stopped or is just slow?
Answer: There is no way to differentiate between a slow and a stopped process as such. So all these algorithms work based on time-outs. The process is declared to be dead if it is slow enough to respond after the time-out. Essentially, a failed/stopped process is equivalent to an infinitely slow process. However, when the old leader(slow process) discovers that another leader has been chosen and their process ID is smaller, the old leader can reinitiate the leader election process.

15.3.2 Ring-based Election Algorithm

The ring algorithm is similar to the bully algorithm in the sense that we assume the processes are already ranked through some metric from 1 to n. However, here a process i only needs to know the IP addresses of its two neighbors (i+1 and i-1). We want to select the node with the highest id. The algorithm works as follows:

- Any node can start circulating the election message. Say process i does so. We can choose to go clockwise or counter-clockwise on the ring. Say we choose clockwise where i+1 occurs after i.

- Process i then sends an election message to process i+1.

- Anytime a process j ≠ i receives an election message, it piggybacks its own ID (thus declaring that it is not down) before calling the election message on its successor (j+1).

- Once the message circulates through the ring and comes back to the initiator i, process i knows the list of all nodes that are alive. It simply scans the list and chooses the highest id.

- It lets all other nodes about the new coordinator.

Note that a process might have to jump over its neighbor and contact the neighbor's neighbor in case the neighbor has failed, otherwise the circulation of the election message across the ring will never finish.
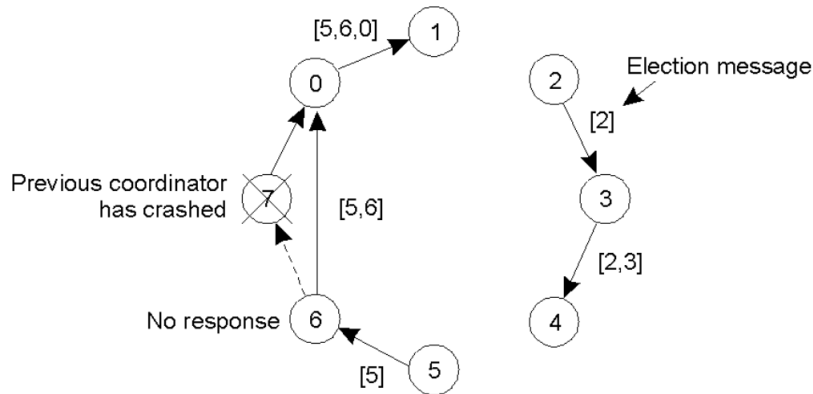

Figure 15.4: Depiction of Ring Algorithm

An example of Ring algorithm is given in Figure 15.4. If the neighbor of a process is down, it sequentially polls each successor (neighbor of neighbor) until it finds a live node. For example, in the figure, when 7 is down, 6 passes the election message to 0. Another thing to note is that this requires us to enforce a logical ring topology on the underlying application, i.e. we need to construct a ring topology on top of the whole system just for leader election.

Question: What initiates the election?
Answer: Any process can initiate the election. Say the leader L was a time-server and a process X discovers that X is not responding anymore, then X can initiate the election process.

Question: Is the ring topology applied only for elections?
Answer: Yes, as this is a distributed system, communication between any pair of processes can take place. But for the purpose of leader election, a process of ID i will only communicate with the processes having id $i-1$ and $i+1$, i.e. as a ring topology because this brings the complexity of the leader election down to O(n).

Question: How many members of the ring does a process need to know?
Answer: In the worst case, a process will have to communicate with all members of the system during the election process, i.e. in a scenario when most of the processes have failed and stopped responding, the remaining processes will have to jump over multiple processes in between in order to elect the next leader.

Question: If every node knows the process ID of all other nodes in the system, why do we need to perform leader election?
Answer: This is because the process starting the election does not know the highest ID process that is alive at that point of time.

Question: If two elections are going on simultaneously, what happens when the first one concludes?
Answer: The second election continues until its completion.

Question: Does election have an ID?
Answer: Yes, so as to facilitate multiple elections occurring simultaneously.

Question: What if the leader crashes even before it is announced?
Answer: You will still announce the chosen process as the leader and it will continue to remain the leader until some process notices that it is no more alive and initiates another election algorithm.

Question: How do you annouce the leader? Is it a message broadcast?
Answer: The leader information is circulated along the ring itself by the initiator of the election. It is not a one-to-all broadcast.

15.3.3   Time Complexity

The bully algorithm runs in

- $O(n^2)$ in the worst case (this occurs when the node with lowest ID initiates the election)

- $O(n-2)$ in the best case (this occurs when the node with highest ID that is alive initiates the election)

The ring algorithm always takes 2(n-1) messages to execute. The first (n-1) is during the election query and second time to announce the results of the election. It is easy to extend ring algorithm for other metrics like load, etc.

15.4   Distributed Synchronization

Every time we wish to access a shared data structure or critical section in a distributed system, we need to guard it with a lock. A lock is acquired before the data structure is accessed, and once the transaction has completed, the lock is released. Consider the example below:

In this example, there are two clients sending a buy request to the Online Store Server. The store implements a thread-pool model. Initially the item count is 3. The correct item count should be 1 after two buy operations. If locks are not implemented there may be chance of race condition and item count can be 2. This is because the decrement is not an atomic operation. Each thread needs to read, update and write the item value. The second thread might read the value while first thread is updating the value (it will read 3) ans update it to 2 and save it, which is incorrect. This is an example of trivial race condition.

Figure 15.5: Example of a race condition in an online store.

15.4.1   Centralized Mutual Exclusion

In this case, locking and unlocking coordination are done by a master process. All processes are numbered 1 to n. We run leader election to pick the coordinator. Now, if any process in the system wants to acquire a lock, it has to first send a lock acquire request to the coordinator. Once it sends this request, it blocks execution and awaits reply until it acquires the lock. The coordinator maintains a queue for each data structure of lock requests. Upon receiving such a request, if the queue is empty, it grants the lock and sends the message, otherwise it adds the request to the queue. The requester process upon receiving the lock executes the transaction, and then sends a release message to the coordinator. The coordinator upon receipt of such a message removes the next request from the corresponding queue and grants that process the lock. This algorithm is fair and simple.
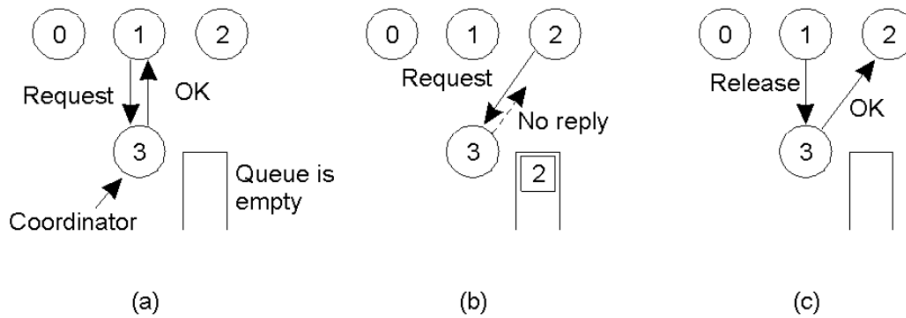

Figure 15.6: Depiction of centralized mutual exclusion algorithm.

An example of the algorithm is given in Figure 15.6. There are two major issues with this algorithm, related to failures. When coordinator process goes down while one of the processes is waiting on a response to a lock request, it leads to inconsistency. The new coordinator that is elected (or reboots) might not know that the earlier process is still waiting for a response. This issue can be tackled by maintaining persistent data on disk whenever a queue of the coordinator is altered. Even if the process crashes, we can read the file and persist the state of the locks on storage and recover the process.

The harder problem occurs when one of the client process crashes while it is holding the lock (during one of its transactions). In such a case, coordinator is just waiting for the lock to be released while the other process has gone down. We cannot use timeout in this case, because usually transactions take arbitrary amount of time to go through. All other processes that are waiting on that lock are also blocked forever. Even if the coordinator somehow knew that the client process has crashed, it may not always be advisable to take the lock forcibly back because the client process may eventually reboot and think it has the lock and continue its transaction. This causes inconsistency. This is a thorny problem which does not have any neat

solution. This limits the practicality of such an centralized algorithm.

Question: Can you check whether the process holding the lock is alive rather than forcing it to release the lock?
Answer: Yes, this is a feasible solution. When the lock holding time expires, the coordinator can just ask the process whether it wants to renew the lock for another time slice or give the lock up. You can grant locks in intervals of time and have the provision to extend that time so long as the process is alive and willing to hold the lock.

Question: What about consistency when the process holding the lock crashes?
Answer: This is an application-level problem. It needs to be handled by the application logic such that all transactions happening within the synchronized block adhere to the desired ACID properties.

15.4.2   Decentralized Algorithm

Decentralized algorithms use voting to figure out which lock requests to grant. In this scenario, each process has an extra thread called the coordinator thread which deals with all the incoming locking requests. Essentially, every process keeps a track of who has the lock, and for a new process to acquire a new lock, it has to be granted an *OK* or go ahead vote from the strict majority of the processes. Here, majority means more than half the total number of nodes (live or not) in the system. Thus, if any process wishes to acquire a lock, it requests it from all other processes and if the majority of them tell it to acquire the lock, it goes ahead and does so. The majority guarantees that a lock is not granted twice. Upon the receipt of the vote, the other processes are also told that a lock has been acquired and thus, the processes hold up any other lock request. Once a process is done with the transaction, it broadcasts to every other process that it has released the lock.

This solves the problem of coordinator failure because if some nodes go down, we can deal with it so long as the majority agrees that whether the lock is in use or not. Client crashes are still a problem here.

Question: If the minority of processes say NO to a lock request, will there be inconsistency?
Answer: Since we are taking a majority vote here, we will assume that the majority of the processes have the correct state of the lock. Normally, there wouldn't be any inconsistency issues. Only inconsistency comes when a process crashes and is not aware of the state of the lock after it restarts. Such a process will tend to give bad replies and this issue is overcome by relying on the majority to give the right replies.

Question: If we want to overcome getting bad replies from processes, can we use techniques like process start time?
Answer: Strictly speaking, the system can force a process to first recover its lock state fully before it can start responding to lock vote requests. However, this is not used in the method currently because we are anyway robust to the bad replies as long as they are in minority.

Question: Is it possible that nobody gets the majority votes?
Answer: If more than half the processes of the system crash, then this will happen as although the minority of processes that are alive have the correct state of the lock but the majority is sending garbage back during voting.

Question: Is there an associated RELEASE message for every REQUEST message?
Answer: Yes, once a process is done executing their critical section, they send a RELEASE request to all other

processes so that they can update their lock state and grant the lock to the next process in REQUEST queue.

Question: Is there an associated RELEASE message for every REQUEST message?
Answer: Yes, once a process is done executing their critical section, they send a RELEASE request to all other processes so that they can update their lock state and grant the lock to the next process in REQUEST queue.

Question: In the decentralized algorithm, what are we voting on?
Answer: The vote is simply a YES or NO. Every process responds with YES or NO based on whether it wants to allow the asking process to acquire the lock or not.

Question: If a process is not using the lock, will it say YES?
Answer: In thos technique, every process is running the lock manager code as well. So they are aware of the state of the lock and they will grant the lock based on that state.

15.4.3 Distributed Algorithm

This algorithm, developed by Ricart and Agrawala, needs $2(n-1)$ messages and is based on Lamport's clock and total ordering of events to decide on granting locks. After the clocks are synchronized, the process that asked for the lock first gets it. The initiator sends request messages to all $n-1$ processes stamped with its ID and the timestamp of its request. It then waits for replies from *all* other processes.

Any other process upon receiving such a request either sends reply if it does not want the lock for itself, or is already in the transaction phase (in which case it doesn't send any reply and the initiator has to wait), or it itself wants to acquire the same lock in which case it compares its own request timestamp with that of the incoming request. The one with the lower timestamp gets the lock first.

- Process k enters critical section as follows:
    - Generate new time stamp $TS_k = TS_{k+1}$
    - Send request(k,TS_k) all other n-1 processes
    - Wait until reply(j) received from all other processes
    - Enter critical section
- Upon receiving a request message, process j
    - Sends reply if no contention
    - If already in critical section, does not reply, queue request
    - If wants to enter, compare TS_j with TS_k and send reply if $TS_k < TS_j$, else queue (recall: total ordering based on multicast)

This approach is fully decentralized but there are n points of **failure**, which is worse than the centralized one.

15.4.4 Token Ring Algorithm

In the *token ring algorithm*, the actual topology is not a ring, but for locking purpose there is a logical ring and processes only talk to neighboring processes. A process with the token has the lock at that time. To

acquire the lock one needs to wait. The token is circulated through the logical ring. No method is there to request the token, the process needs to wait to get a lock. Once the process has token it can enter the critical section.

This was designed as part of a networking protocol Token Ring. In physical networking, only one node can transmit at a time. If multiple nodes transmit at a time there is a chance of collision. Ethernet handled it by detecting collisions. A node transmits and if there is a collision it backs off and will succeed eventually. In Token Ring this was handled using locks. Only one machine has lock at an instance in the network and it transmit at that particular time.

In this algorithm one problem is loss of token. Regenerating the token is non-trivial, as you can not use timeout strategy.