

## UNIT-2

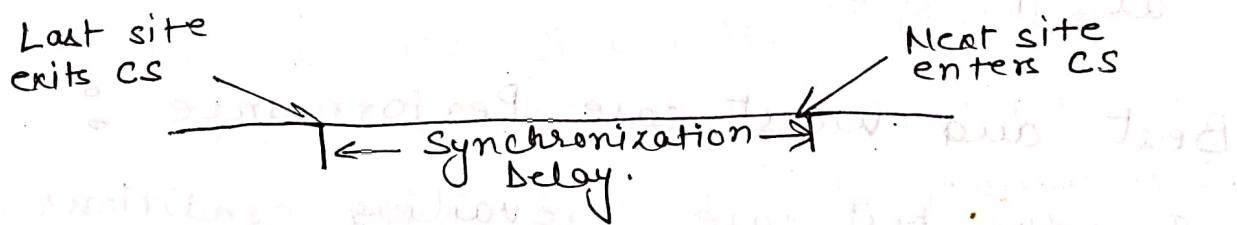
### Distributed Mutual Exclusion

- Requirements of Mutual Exclusion Algorithm

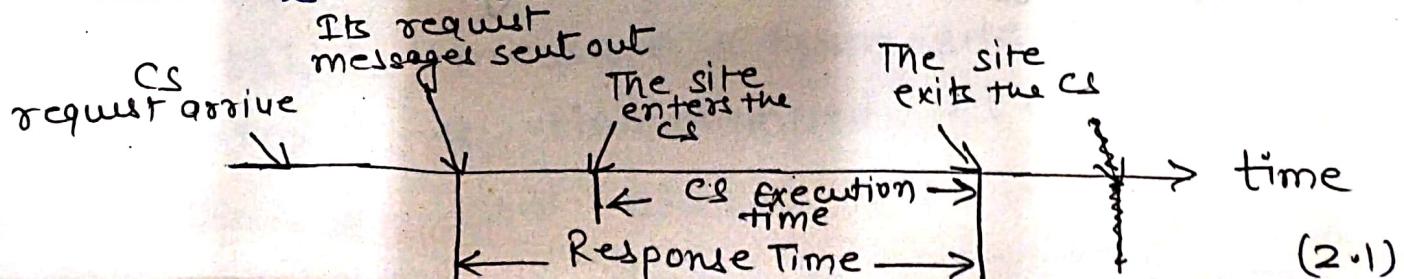
1. Freedom from Deadlock.
2. freedom from starvation.
3. Fairness
4. Fault Tolerance.

- How to measure the Performance

1. Number of messages
2. Synchronization delay : It is the time required after a site leaves the critical section and before the next site enters the critical section.



3. Response time : It is the time interval a request waits for its CS execution to be over after its request messages have been sent out.



4. System Throughput, which is the rate at which the system executes requests for token. If 'sd' is the synchronization delay and E is the average critical section execution time, then the throughput is given by the following equation

$$\text{System throughput} = 1/(sd + E)$$

Low and High load performance :

Under low load condition, there is seldom more than one request for mutual exclusion simultaneously in the system.

Under high load conditions, there is always a pending request for mutual exclusion at a site.

Best and Worst case Performance :

In the best case, prevailing conditions are such that a performance metric attains the best possible value. Often for mutual exclusion algorithm, the best and worst cases coincide with low and high loads, respectively.

## Token-Based and Non-Token Based Algo.

### Non-Token Based Algorithms :

- \* It uses timestamps to order requests for the CS and to resolve conflicts between simultaneous requests for the CS.
- \* Each request for the CS gets a timestamp, and smaller timestamp requests have priority over larger timestamp requests.

#### (A) Lamport's Algorithm :

Every site  $S_i$  keeps a queue, request-queue, which contains mutual exclusion requests ordered by their timestamps.

#### The Algorithm :

##### (a) Requesting the critical section :

- when a site  $S_i$  wants to enter the CS, it sends a REQUEST  $(ts_i, i)$  message to all the sites in its request set  $R_i$  and places the request on request-queue; where  $(ts_i, i)$  is the timestamp of the request.

- when a site  $S_j$  receives the REQUEST  $(ts_i, i)$  message from site  $S_i$ , it returns a timestamped REPLY message to  $S_i$  and places site  $S_i$ 's request

(2.3)

on request queuej.

(b) Executing the critical section :

Site  $s_i$  enters the critical section when the two following conditions hold :

- (i)  $s_i$  has received a message with timestamp larger than  $(ts_i, i)$  from all other sites.
- (ii)  $s_i$ 's request is at the top of request queuei.

(c) Releasing the critical section :

- (i) Site  $s_i$ , upon releasing the CS, removes its request from the top of its request queue and sends a timestamped RELEASE message to all the sites in its request set.
- (ii) When a site  $s_j$  receives a RELEASE message from site  $s_i$ , it removes  $s_i$ 's request from its request queue.

When a site removes a request from its request queue, its own request may come at the top of the queue, enabling it to enter the CS. The algorithm executes CS requests in the increasing order of timestamps.

## The Ricart-Agrawala Algorithm :

It is an optimization of Lamport's algorithm that dispenses with RELEASE messages by cleverly merging them with REPLY message.

### The Algorithm

#### (i) Requesting the critical section :

1. When a site  $S_i$  wants to enter the CS, it sends a timestamped REQUEST message to all the sites in the request set.
2. When site  $S_j$  receives a REQUEST message from site  $S_i$ , it sends a REPLY message to site  $S_i$  if site  $S_j$  is neither requesting nor executing the CS or if site  $S_j$  is requesting and  $S_i$ 's request's timestamp is smaller than site  $S_j$ 's own request's timestamp. The request is deferred otherwise.

#### (ii) Executing the critical section :

3. Site  $S_i$  enters the CS after it has received REPLY messages from all the sites in its request set.

#### (iii) Releasing the critical section :

4. When site  $S_i$  exits the CS, it sends REPLY messages to all the deferred requests.

(2.5)

A site's REPLY messages are blocked by sites that are requesting the CS with higher priority (i.e. a smaller timestamp). THE  
MAEKAWA

Thus, when a site sends out REPLY message to all the deferred requests, the site with the next highest priority request receives the last needed REPLY message and enters the CS.

The execution of CS requests in the algorithm is always in the order of their timestamp.

### (C) Maekawa's Algorithm :

first, a site does not request permission from every other site, but only from a subsets of the sites.

Second, in Maekawa's algorithm a site can send out only one REPLY message at a time i.e. after it has received a RELEASE message for the previous REPLY message. Therefore, a site si locks all the site in R<sub>i</sub> in exclusive mode before executing its CS.

~~EX~~

## MAEKAWA'S ALGORITHM.

### The Algorithm

#### Requesting the critical section:

1. A site  $s_i$  requests access to the CS by sending REQUEST( $i$ ) message to all the sites in its request set  $R_i$ .
2. When a site  $s_j$  receives the REQUEST( $i$ ) message, it sends a REPLY( $j$ ) message to  $s_i$  provided it hasn't sent a REPLY message to a site from the time it received the last RELEASE message. Otherwise, it queues up the REQUEST for later consideration.

#### Executing the critical section :

3. Site  $s_i$  accessed the CS only after receiving REPLY messages from all the sites in  $R_i$ .

#### Releasing the critical section :

4. After the execution of CS is over, site

(2.7)

$s_i$  sends RELEASE(i) message to the sites in  $R_i$ .

5. When a site  $s_j$  receives a RELEASE message from site  $s_i$ , it sends a REPLY message to the next site waiting in the queue and deletes the entry from the queue. If the queue is empty, then the site updates its state to reflect that the site has not sent out any REPLY message.

## Token-Based Algorithm :

- \* A unique token is shared among all sites. A site is allowed to enter its CS if it possesses the token.
- \* It uses sequence numbers instead of timestamps. A site increments its sequence number counter every time it makes a request for the token.

### (A) Suzuki-Kasami's Broadcast Algorithms :

#### (i) Requesting the critical section :

1. If the requesting site  $s_i$  does not have the token, then it increments its sequence number,  $RN_i[i]$  and sends a REQUEST( $i, Sn$ ) message to all other sites. ( $Sn$  is the updated value of  $RN_i[i]$ .)
2. When a site  $s_j$  receives this message, it sets  $RN_j[i]$  to  $\max(RN_j[i], Sn)$ . If  $s_j$  has the idle token, then it sends the token to  $s_i$  if  $RN_j[i] = LN[i] + 1$ .

#### (ii) Executing the critical section :

3. Site  $s_i$  executes the CS when it has received the token.

- (iii) Releasing the critical section : Having finished the execution of the CS, site  $s_i$  takes the following actions :
4. It sets  $LNC[i]$  elements of the token array equal to  $RNi[i]$ .
  5. For every site  $s_j$  whose ID is not in the token queue, it appends its ID to the token queue if  $RNi[j] = LN[j] + 1$ .
  6. If token queue is nonempty after the above update, then it deletes the top site ID from the queue and sends the token to the site indicated by the ID.

Thus, after having executed its CS, a site gives priority to other sites with outstanding requests for the CS.

- \* The Suzuki-Kasami algorithm is not symmetric because a site retains the token even if it does not have a request for the CS, which is contrary to the spirit of Ricart and Agrawala's definition of a symmetric algorithm : "no site possesses the right to access its CS when it has not been requested."

(2.10)

- (tokens going to sites)*
- In this algo, each site maintains information about the state of other sites in the system and uses it to select a set of sites that are likely to have the token.
  - The site requests the token only from these sites, reducing the number of messages required to execute the cl. It is called a heuristic algorithm because site are heuristically selected for sending token request message.

### The Algorithm :

#### Requesting the critical section :

1. If the requesting site  $s_i$  does not have the token, then it takes the following actions :
  - It sets  $SV_i[i] := R$ .
  - It increments  $SN_i[i] := SN_i[i] + 1$ .
  - It sends REQUEST( $i, sn$ ) message to all sites  $s_j$  for which  $SV_i[j] = R$ . ( $sn$  is the updated value of  $SN_i[i]$ .)
2. When a site  $s_j$  receives the REQUEST( $i, sn$ ) message, it discards the message if  $SN_j[i] \geq sn$  because the message is outdated.

(2.11)

Otherwise, it sets  $SN_j[i]$  to 'sn' and takes the following actions based on its own state :

- $SV_j[j] = N$ : set  $SV_j[i] := R$
- $SV_j[j] = R$ : If  $SV_j[i] \neq R$  then set  $SV_j[i] := R$  and send a REQUEST( $j, SN_j[i]$ ) message to  $s_i$  (else do nothing).
- $SV_j[j] = E$  : set  $SV_j[i] := R$
- $SV_j[j] = H$  : set  $SV_j[i] := R$ ,  $TSV[i] := R$ ,  $TSN[i] := sn$ ,  $SV_j[j] := N$  and send the token to site  $s_i$ .

#### Executing the critical Section

3.  $s_i$  executes the CS after it has received the token. Before entering the CS,  $s_i$  ~~gets~~ sets  $SV_i[i]$  to E.

#### Releasing the critical section

4. Having finished the execution of the CS, site  $s_i$  sets  $SV_i[i] := N$  and  $TSV[i] := N$ , and updates its local and token vectors in the following way :

For all  $s_j, j = 1$  to  $N$  do

if  $SN_i[j] > TSN[j]$

then

(\* update token information from local information \*)

{  $TSV[j] := SV_i[j]$ ;  $TSN[j] := SN_i[j]$  }

(2.12)

else

(\* update local information from token  
information \*)

{  $SV_i[j] := TSV[j]$ ;  $SN_i[j] := TSN[j]$  }

5. If ( $\forall j :: SV_i[j] = N$ ), then set  $SV_i[j] := H$ ,  
else send the token to a site  $S_j$  such  
that  $SV_i[j] = R$ .

### (c) Raymond's Tree-Based Algorithm

Sites are logically arranged as a directed tree such that the edges of the tree are assigned directions towards the site (root of the tree) that has the token.

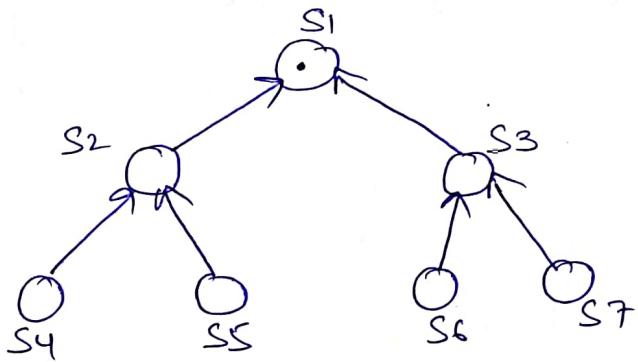
Every site has a local variable holder that points to an immediate neighbor node on a directed path to the root node.

Every site keeps a FIFO queue, called request-q, which stores the requests of those neighboring sites that have sent a request to this site, but have not yet been sent the token.

## The Algorithm :

Requesting the critical section

- When a site wants to enter the CS, it sends a REQUEST message to the node along the directed path to the root, provided it does not hold the token and its request-q is empty. It then adds its request to its request-q.



- When a site on the path receives this message, it places the REQUEST in its request-q and sends a REQUEST message along the directed path to the root provided it has not sent out a REQUEST message on its outgoing edge (for a previously received REQUEST on its request-q).
- When the root site receives a REQUEST message, it sends the token to the site from which it received the REQUEST message and sets its holder

(214)

variable to point at that site.

4. When a site receives the token, it deletes the top entry from its request-q, sends the token to the site indicated in this entry, and sets its holder variable to point at that site. If the request-q is nonempty at this point, then the site sends a REQUEST message to the site which is pointed at by holder variable.

Executing the critical section :

5. A site enters the CS when it receives the token and its own entry is at the top of its request-q. In this case, the site deletes the top entry from its request-q and enters the CS.

Releasing the critical section :

6. After a site has finished execution of the CS, it takes the following actions :
  1. If its request-q is non-empty, then it deletes the top entry from its request-q, sends the token to that site, and sets its holder variable to point at

(contd)

that site.

7. If the request<sub>q</sub> is nonempty at this point, then the site sends a REQUEST message to the site which is pointed at by the holder variable.

Performance :

The average message complexity of Raymond's algorithm is  $O(\log N)$ , because the average distance between any two nodes in the tree with  $N$  nodes is  $O(\log N)$ .