

Leader Election

CS60002: Distributed Systems

Pallab Dasgupta
Professor,
Dept. of Computer Sc. & Engg.,
Indian Institute of Technology Kharagpur



Leader Election in Rings

- **Models**
 - Synchronous or Asynchronous
 - Anonymous (no unique id) or Non-anonymous (unique ids)
 - Uniform (no knowledge of N , the number of processes) or non-uniform (knows N)
- **Known Impossibility Result:**
 - There is no synchronous, non-uniform leader election protocol for anonymous rings

Election in Asynchronous Rings

- **LeLann's and Chang-Robert's Algorithms**
 - send own id to node on left
 - if an id received from right, forward id to left node only if received id greater than own id, else ignore
 - if own id received, declares itself “leader”
- Works on unidirectional rings
- Message complexity = $O(n^2)$

Hirschberg-Sinclair Algorithm

- Operates in phases, requires bidirectional ring
- In the k^{th} phase, send own id to 2^k processes on both sides of yourself (directly send only to next processes with id and k in it)
- If id received, forward if received id greater than own id, else ignore
- Last process in the chain sends a reply to originator if its id less than received id
- Replies are always forwarded
- A process goes to $(k+1)^{\text{th}}$ phase only if it receives a reply from both sides in k^{th} phase
- Process receiving its own id – declare itself “leader”. At most $\lg n$ rounds

Features: Hirschberg-Sinclair

- Message Complexity: $O(n \lg n)$
- Lots of other algorithms exist for rings
- Lower Bound Result:
 - *Any comparison-based* leader election algorithm in a ring requires $\Omega(n \lg n)$ messages

The Echo Algorithm – a wave algorithm

```
var  $rec_p$            : integer           init 0;           // Counts no of recvd mesgs  
     $father_p$        : process           init undef;
```

For the initiator

```
begin forall  $q \in Neigh_p$  do send  $\langle tok \rangle$  to  $q$  ;  
    while  $rec_p < \#Neigh_p$  do  
        begin receive  $\langle tok \rangle$  ;  $rec_p = rec_p + 1$  end ;  
        decide  
    end
```

For non-initiators

```
begin receive  $\langle tok \rangle$  from neighbor  $q$  ;  $father_p = q$  ;  $rec_p = rec_p + 1$  ;  
    forall  $q \in Neigh_p, q \neq father_p$  do send  $\langle tok \rangle$  to  $q$  ;  
    while  $rec_p < \#Neigh_p$  do  
        begin receive  $\langle tok \rangle$  ;  $rec_p = rec_p + 1$  end ;  
        send  $\langle tok \rangle$  to  $father_p$   
    end
```

Extinction on The Echo Algorithm

```
var  $caw_p$       : process   init undef ;    // Currently active wave
 $rec_p$          : integer   init 0 ;          // No of  $\langle tok, caw_p \rangle$  received
 $father_p$       : process   init undef ;    // Father in wave  $caw_p$ 
 $lrec_p$         : integer   init 0 ;          // No of  $\langle ldr, . \rangle$  received
 $win_p$          : process   init undef ;    // Identity of leader
```

begin if p is initiator then

 begin $caw_p = p$;

 forall $q \in Neigh_p$ do send $\langle tok, p \rangle$ to q ;

 end ;

while $lrec_p < \#Neigh_p$ do

 begin receive *msg* from q ;

 if $msg = \langle ldr, r \rangle$ then

 begin if $lrec_p = 0$ then

 forall $q \in Neigh_p$ do send $\langle ldr, r \rangle$ to q ;

$lrec_p = lrec_p + 1$; $win_p = r$;

 end ;

Extinction on Echo Algorithm contd..

```
else                                     // mesg is a  $\langle \text{tok}, r \rangle$  message
  begin if  $r < \text{caw}_p$  then           // Reinitialize the algorithm
    begin  $\text{caw}_p = r$  ;  $\text{rec}_p = 0$  ;  $\text{father}_p = q$  ;
      forall  $s \in \text{Neigh}_p, s \neq q$  do send  $\langle \text{tok}, r \rangle$  to  $s$ 
    end ;
    if  $r = \text{caw}_p$  then
      begin  $\text{rec}_p = \text{rec}_p + 1$  ;
        if  $\text{rec}_p = \#\text{Neigh}_p$  then
          if  $\text{caw}_p = p$ 
            then forall  $s \in \text{Neigh}_p$  do send  $\langle \text{ldr}, p \rangle$  to  $s$ 
            else send  $\langle \text{tok}, \text{caw}_p \rangle$  to  $\text{father}_p$ 
          end ;
        // If  $r > \text{caw}_p$  then the message is ignored – extinction
      end
    end ;
  end ;
  if  $\text{win}_p = p$  then  $\text{state}_p = \text{leader}$  else  $\text{state}_p = \text{lost}$ 
end
```


Features

- *If A is a centralized wave algorithm using M messages per wave, the algorithm $Ex(A)$ elects a leader using at most NM messages*