



SAPIENZA
UNIVERSITÀ DI ROMA
DIPARTIMENTO DI INFORMATICA

CONCURRENT SYSTEMS LECTURE 2

Prof. Daniele Gorla



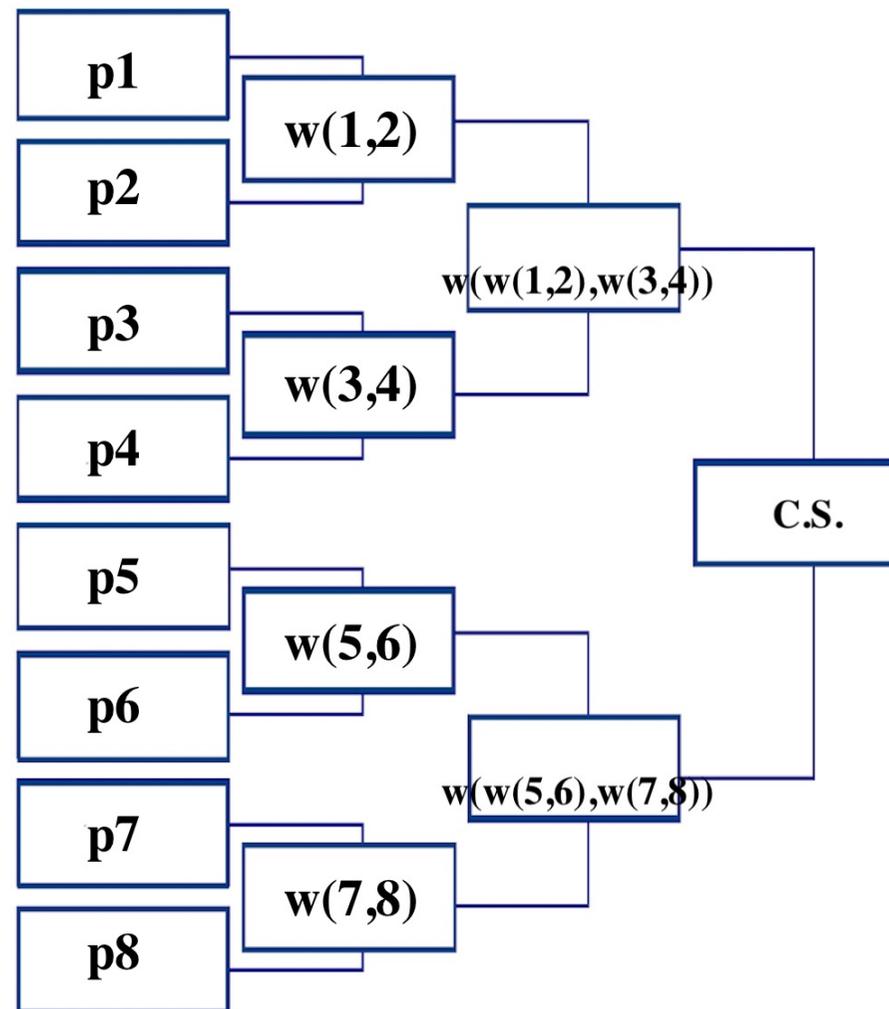
Tournament-based algorithm

Even without contention, Peterson's algorithm costs $O(n^2)$

A first way to reduce this cost is by using a tournament of MUTEX between pairs of processes:

By using Peterson's algorithm for 2 proc, a process wins after $\lceil \log_2 n \rceil$ competitions, each of constant cost.

→ $O(\log n)$





A constant-time algorithm (for n processes)

The cost can be further reduced to $O(1)$.

To begin, consider the following idea:

Initialize Y at \perp , X at any value (e.g., 0)

lock(i) :=	unlock(i) :=
$X \leftarrow i$	$Y \leftarrow \perp$
if $Y \neq \perp$ then FAIL	return
else $Y \leftarrow i$	
if $X = i$ then return	
else FAIL	

Without contention, this requires 4 accesses to the registers for entering the CS

Problem:

- we don't want the FAIL (that forces the process to invoke lock again and again), but an implementation of lock that keeps the process inside this primitive until it wins
- It is possible to have an execution where nobody accesses its CS
 → if repeated for ever, enatils a deadlock





Fast MUTEX algorithm (by Lamport)

Initialize Y at \perp , X at any value (e.g., 0)

lock(i) :=

* FLAG[i] \leftarrow up

X \leftarrow i

if Y \neq \perp then FLAG[i] \leftarrow down

wait Y = \perp

goto *

else Y \leftarrow i

if X = i then return

else FLAG[i] \leftarrow down

$\forall j$.wait FLAG[j] = down

if Y = i then return

unlock(i) :=

Y \leftarrow \perp

FLAG[i] \leftarrow down

return

else wait Y = \perp

goto *

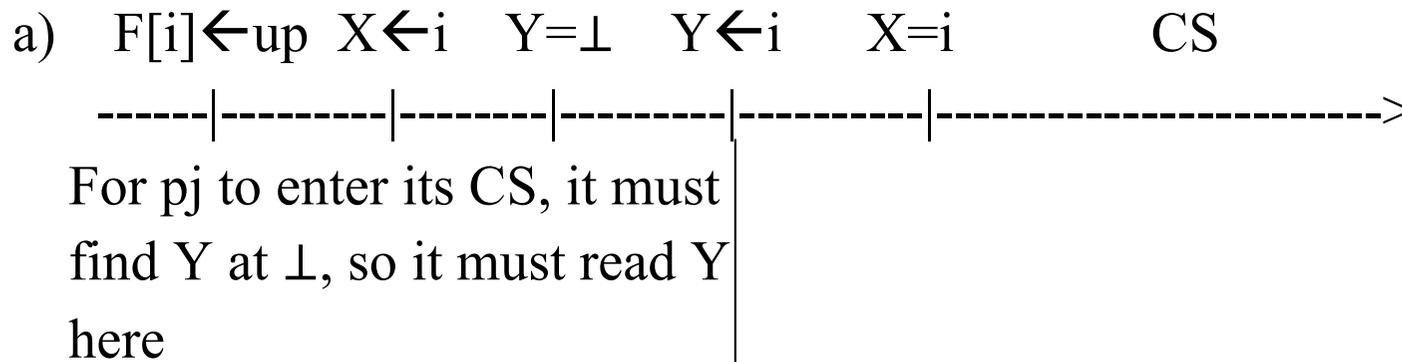




MUTEX: if p_i is in CS, then no other p_j can simultaneously be in CS

Proof:

How can p_i enter its CS?



Where did p_j write X ?

So, it must have written X here. | not here, otherwise p_i would not have read i in X

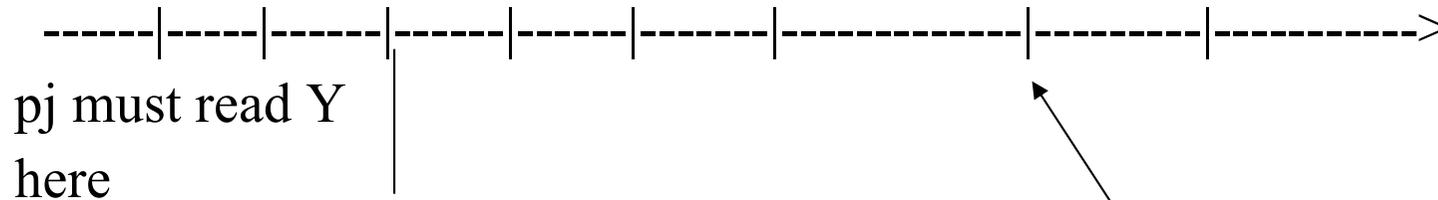
Hence, when p_j reads X , it finds it different from j

- it must wait for p_i 's unlock before starting again
- p_j cannot be in CS while p_i is





b) $F[i] \leftarrow \text{up } X \leftarrow i \quad Y = \perp \quad Y \leftarrow i \quad X \neq i \quad F[i] \leftarrow \text{down} \quad \forall k. F[k] = \text{down} \quad Y = i \quad \text{CS}$



Where did p_j write X ?

If here,
like (a)

So, let p_j write X here

- if p_j reads X at $j \rightarrow$ it enters its CS, that however must be finished before
 - $\rightarrow p_j$'s CS doesn't overlap with p_i 's one
- otherwise, p_j must have written Y before p_i (since p_i finds Y at i)
 - $\rightarrow p_j$ is blocked until p_i unlocks





Deadlock freedom: let p_i invoke lock

- If it eventually wins $\rightarrow \checkmark$
- If it is blocked for ever, where can it be blocked?
 1. In the second wait $Y = \perp$
 - \rightarrow in this case, it read a value in Y different from i
 - \rightarrow there is a p_h that wrote Y after p_i
 - \rightarrow let us consider the last of such p_h 's \rightarrow it will eventually win $\rightarrow \checkmark$
 2. In the $\forall j.$ wait $FLAG[j]=down$
 - \rightarrow this wait cannot block a process for ever
 - if p_j doesn't lock, its flag is down
 - if p_j doesn't find Y at \perp , it puts its flag down
 - if p_j doesn't find X at j , it puts its flag down
 - otherwise p_j enters its CS and eventually unlocks (flag down)





3. In the first wait $Y = \perp$

→ since p_j read a value different from \perp , there is at least one p_k that wrote Y before (but has not yet unlocked)

→ if p_k eventually enters its CS → \checkmark

otherwise, it must be blocked for ever as well. Where?

- In the second wait $Y = \perp$: but then there exists a p_h that eventually enters its CS (see point 1 above) → \checkmark
- In the $\forall j.$ wait $FLAG[j]=down$: this wait cannot block a process for ever (see point 2 above)





Fast MUTEX algorithm (by Lamport)

Without contention, this algorithm requires 5 accesses to the shared registers

It can be proved to satisfy MUTEX and deadlock freedom (you can easily build a scenario where a process is starved)

→ we will see that every deadlock-free algorithm can be turned into a bounded bypass one (but with a quadratic bound...)

To sum up: with atomic R/W registers, we have

- With 2 processes, a $O(1)$ algorithm that satisfies bounded bypass (with bound 1)
- With n processes:
 - a $O(n^2)$ algorithm that satisfies starvation freedom
 - a $O(\log n)$ algorithm that satisfies bounded bypass (with bound $\lceil \log_2 n \rceil$)
 - a $O(1)$ algorithm that satisfies deadlock freedom





From deadlock freedom to bounded bypass

Let DLF be a deadlock free protocol for MUTEX.

We now want to turn it into a bounded bypass protocol for MUTEX

Round Robin algorithm

→ the name comes from a middle age habit for signing petitions, called *Ruban Rond* (that means «round ribbon»)

→ a circular way of signing, to hide the identity of the initiator

Initialize FLAG[i] to down ($\forall i$) and TURN to any proc.id.

```
lock(i) :=  
  FLAG[i] ← up  
  wait (TURN = i OR  
        FLAG[TURN] = down)  
  DLF.lock(i)  
  return
```

```
unlock(i) :=  
  FLAG[i] ← down  
  if FLAG[TURN] = down then  
    TURN ← (TURN+1) mod n  
  DLF.unlock(i)  
  return
```



MUTEX for RR algorithm follows from the assumed MUTEX of DLF

Deadlock freedom of RR: if at least one process invokes RR.lock, then at least one process enters the CS.

Proof:

Since DLF enjoys deadlock freedom, it suffices to prove that at least one process invokes DLF.lock (i.e., at least one proc exists from its wait)

If $TURN=k$ and p_k invoked lock, then it finds $TURN = k$ and exits its wait

Otherwise, any other process finds $FLAG[TURN]=down$ and exits from its wait



Lemma 1: If $\text{TURN} = i$ and $\text{FLAG}[i] = \text{up}$, then p_i enters the CS in at most $(n-1)$ iterations

Proof:

OBS1: TURN changes only when $\text{FLAG}[i]$ is down (i.e., after p_i has completed its CS)

OBS2: $\text{FLAG}[i]=\text{up} \rightarrow$ either p_i is in its CS $\rightarrow \checkmark$
or p_i is competing for its CS \rightarrow it eventually invokes
(or has already invoked)
DLF.lock

OBS3: if p_j invokes lock after that $\text{FLAG}[i]$ is set, p_j blocks in its wait

Let Y be the set of processes competing for the CS (i.e., suspended on DLF.lock)

- Because of OBS2, $i \in Y$
- Because of OBS3, once $\text{FLAG}[i]$ is set, Y cannot grow anymore



- Because DLF is deadlock free, eventually one $p_y \in Y$ wins

If $y=i \rightarrow \checkmark$

otherwise, Y shrinks by one (the p_y that entered the CS). Indeed:

because of OBS1, TURN (and FLAG[TURN]) don't change

$\rightarrow p_y$ cannot enter Y again

We can iterate this reasoning and eventually p_i will win

\rightarrow the worst case is when Y contains all proc's and p_i is the last winner

Lemma 2: If FLAG[i] = up, then TURN is set to i in at most $(n-1)^2$ iterations

Proof:

If TURN= i when FLAG[i] is set $\rightarrow \checkmark$

By Deadlock freedom of RR, at least one proc eventually unlocks

- If FLAG[TURN]=down, then TURN is increased; othw., by Lemma1 p_{TURN} wins in at most $(n-1)$ iterations (and increases TURN)
- If now TURN= i then \checkmark ; otherwise, we repeat the reasoning

The worst case is when TURN= $(i+1) \bmod n$ when FLAG[i] is set



Bounded bypass of RR: if a process invokes RR.lock, then it enters the CS in at most $n(n-1)$ iterations

Proof:

- p_i invokes lock \rightarrow FLAG[i] is set to up
- By lemma 2, in $(n-1)^2$ iterations TURN is set to i
- By lemma 1, in $(n-1)$ iterations p_i enters the CS
- $(n-1)^2 + (n-1) = n(n-1)$