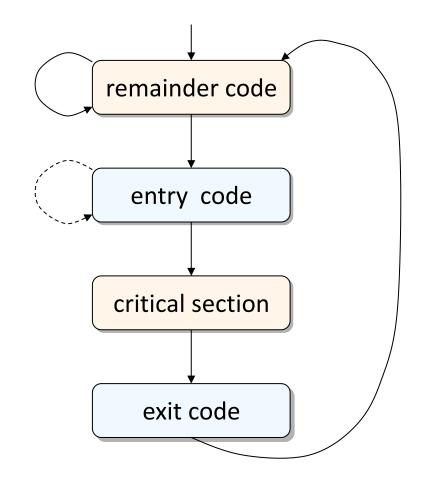
Mutual exclusion algorithms with constant RMR complexity and wait-free exit code

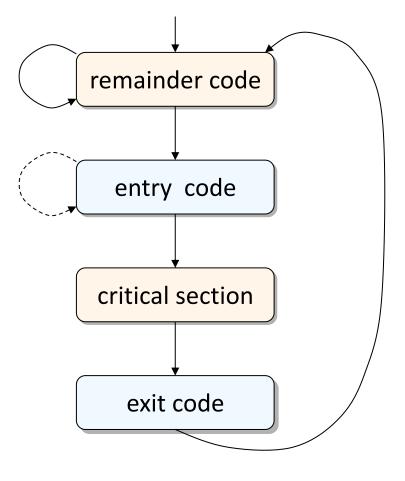
Rotem Dvir

Joint work with Gadi Taubenfeld

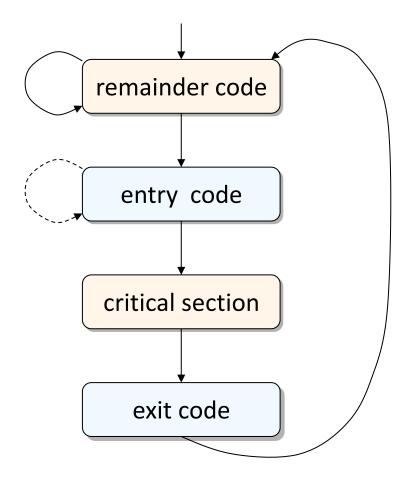
The Problem: design the entry and exit code in a way that guarantees that the mutual exclusion and deadlock-freedom properties are satisfied.



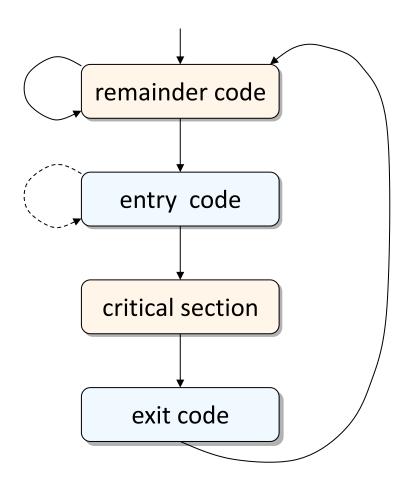
• <u>Mutual Exclusion</u>: No two processes are in their critical sections at the same time.



- <u>Mutual Exclusion</u>: No two processes are in their critical sections at the same time.
- <u>Deadlock-Freedom</u>: If a process is trying to enter its critical section, then some process, not necessarily the same one, eventually enters its critical section.

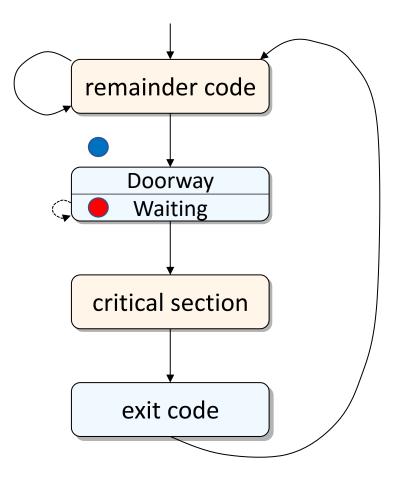


- <u>Mutual Exclusion</u>: No two processes are in their critical sections at the same time.
- <u>Deadlock-Freedom</u>: If a process is trying to enter its critical section, then some process, not necessarily the same one, eventually enters its critical section.
- <u>Starvation-Freedom</u>: If a process is trying to enter its critical section, then this process must eventually enter its critical section.



FIFO

FIFO (First-in-first-out): A beginning process cannot execute its critical section before a waiting process executes its critical section.



Memory References

- <u>Remote reference</u>: slow
- Local reference: fast



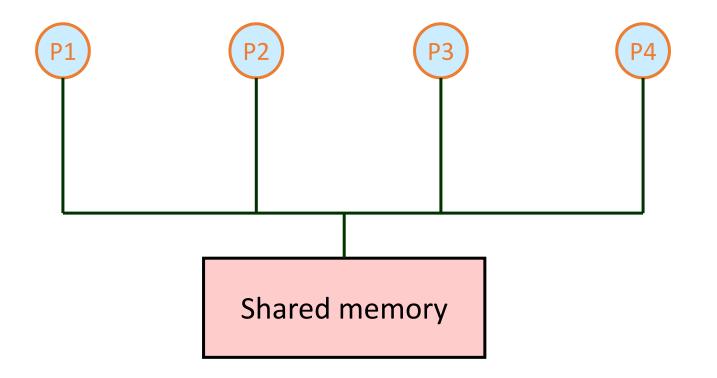
Memory References

- <u>Remote reference</u>: slow
- Local reference: fast

Depends on the memory model



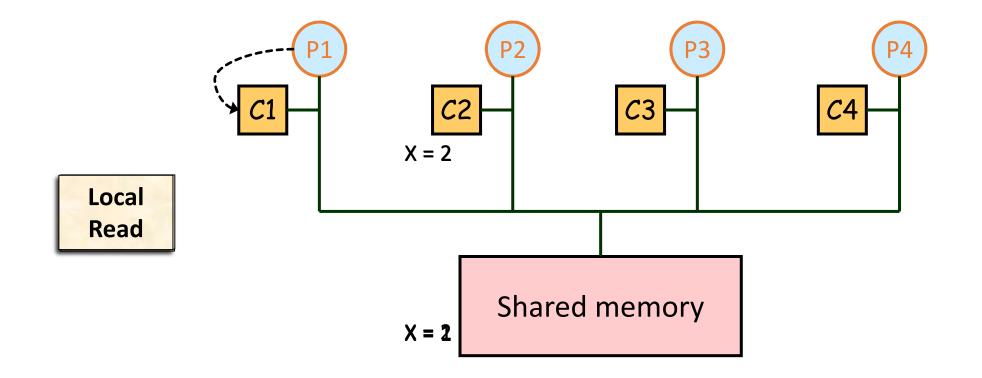
Memory Models: Simple Shared Memory



All memory accesses are <u>remote</u>

Rotem Dvir

Memory Models: <u>Coherent</u> <u>Caching</u>

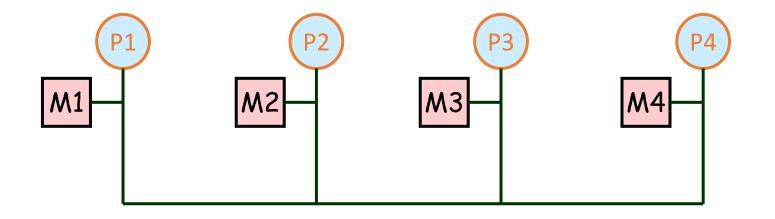


Local and remote memory accesses

Rotem Dvir

OPODIS 2017

Memory Models: <u>Distributed</u> <u>Shared</u> <u>Memory</u>



Local and remote memory accesses

Rotem Dvir

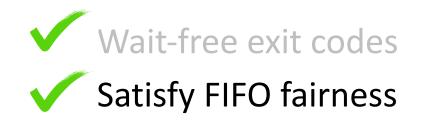
OPODIS 2017

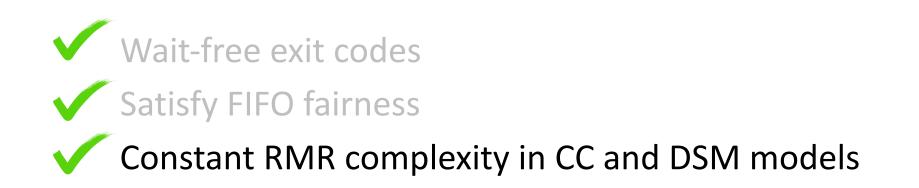
<u>**Remote Memory Reference Complexity</u>**</u>

<u>RMR Complexity:</u> The maximum number of remote memory references that a process may need to perform in its entry and exit sections.









Wait-free exit codes
 Satisfy FIFO fairness
 Constant RMR complexity in CC and DSM models
 Not assumed that the number of processes is priori known

Wait-free exit codes
 Satisfy FIFO fairness
 Constant RMR complexity in CC and DSM models
 Not assumed that the number of processes is priori known
 Use only O(n) shared memory locations

Wait-free exit codes
Satisfy FIFO fairness
Constant RMR complexity in CC and DSM models
Not assumed that the number of processes is priori known
Use only O(n) shared memory locations
Make no assumptions on what and how memory is allocated

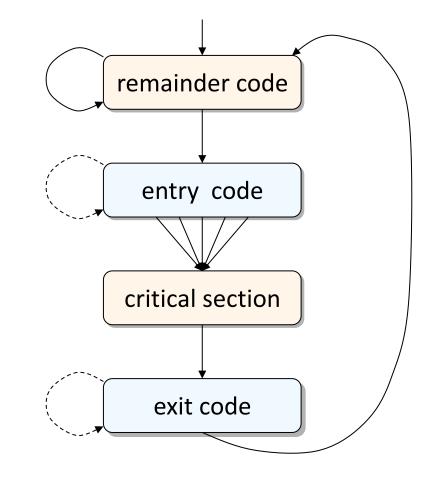
Wait-free exit codes Satisfy FIFO fairness Constant RMR complexity in CC and DSM models Not assumed that the number of processes is priori known ✓ Use only O(n) shared memory locations Make no assumptions on what and how memory is allocated Instruction set: Read/Write, Fetch-And-Store, Compare-And-Swap

MCS Algorithm Properties

X Wait-free exit code Satisfy FIFO fairness Constant RMR complexity in CC and DSM models Not assumed that the number of processes is priori known Use only O(n) shared memory locations Make no assumptions on what and how memory is allocated Instruction set: Read/Write, Fetch-And-Store, Compare-And-Swap

Wait-Free Vs. Not Wait-Free

- Exiting process may need to wait for entering process
- Particularly problematic in high contention scenarios



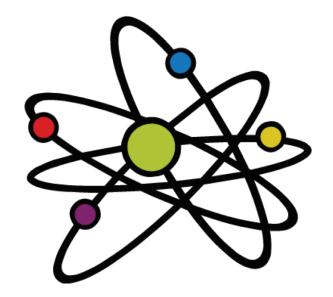
Space Complexity When Using L Locks

- <u>First algorithm</u>: O(Ln) shared memory locations
- <u>Second algorithm</u>: O(L+n) shared memory locations



Computational Model: Atomic Instruction Set

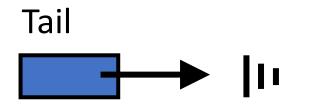
- Atomic read/write
- Compare-And-Swap (Destination, New, Old)
- Fetch-And-Store (Destination, Value)

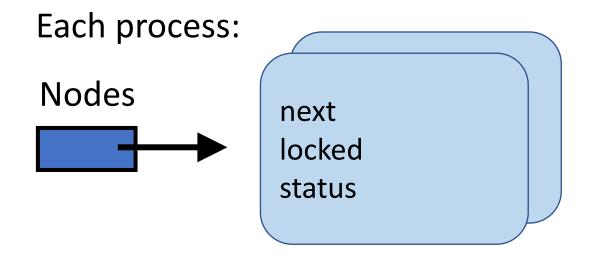


The First Algorithm



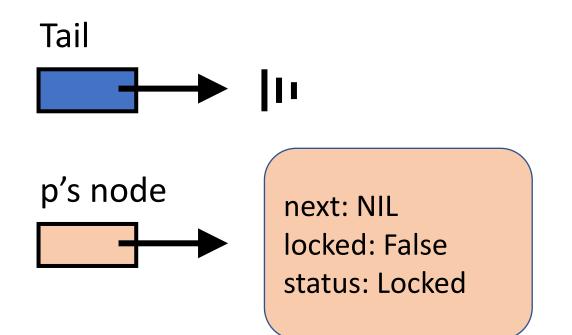
Initially





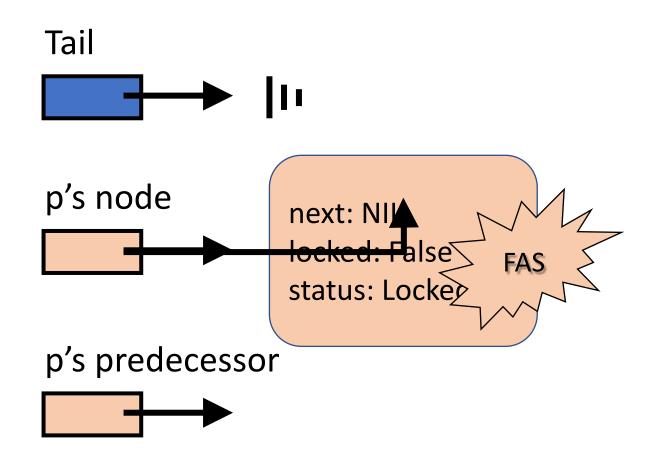
Rotem Dvir

OPODIS 2017



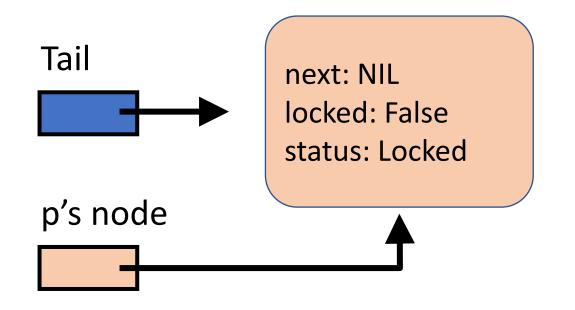


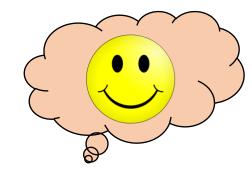
Initializing





Ththeochiette the queue

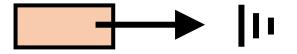


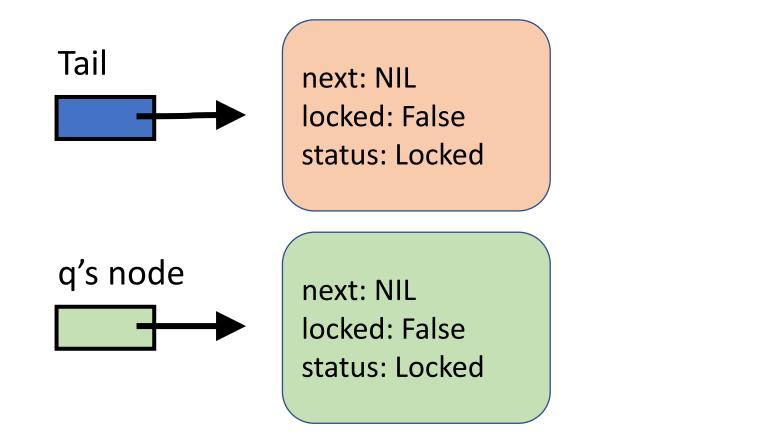




Continues to critical section

p's predecessor



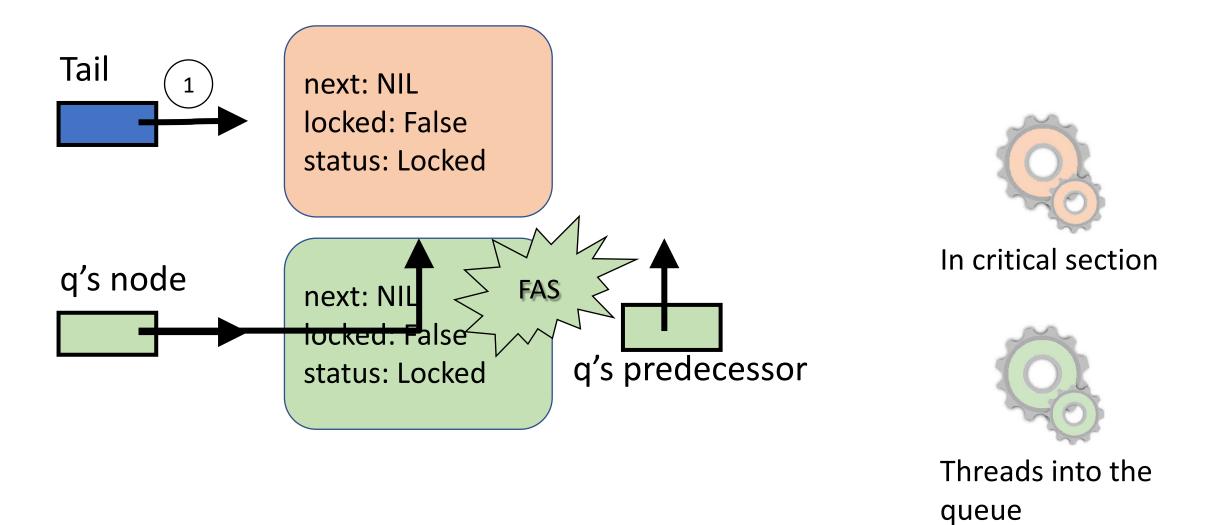


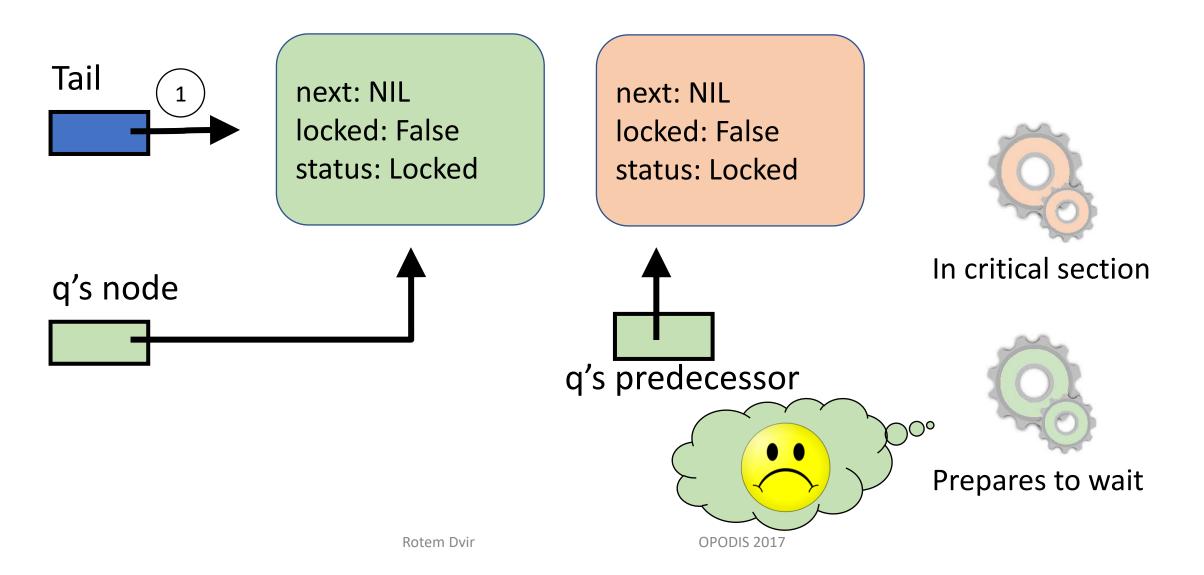


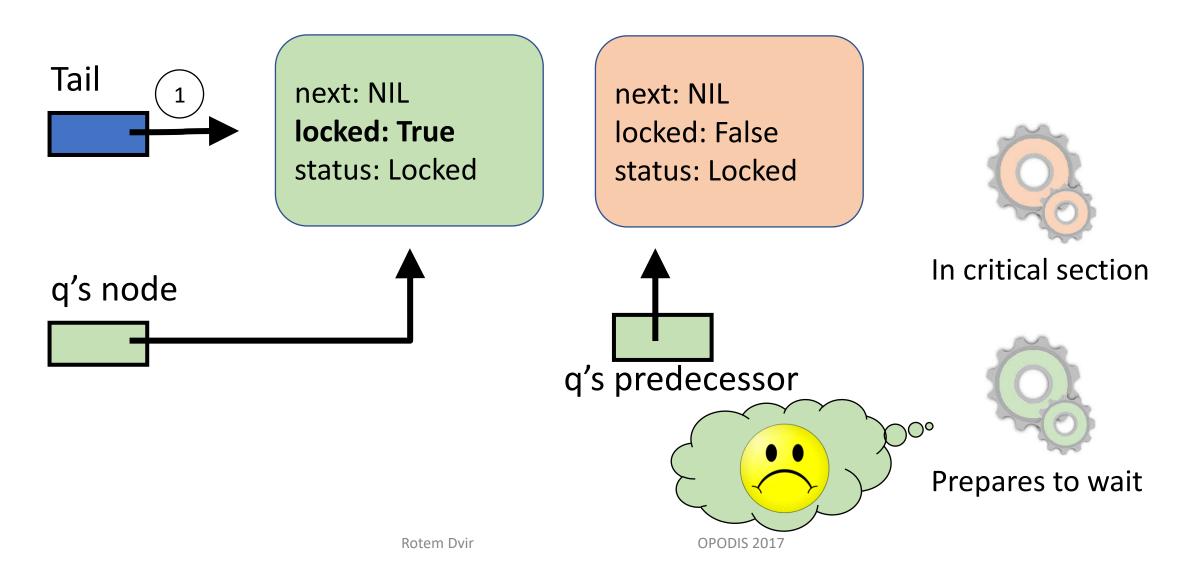
In critical section

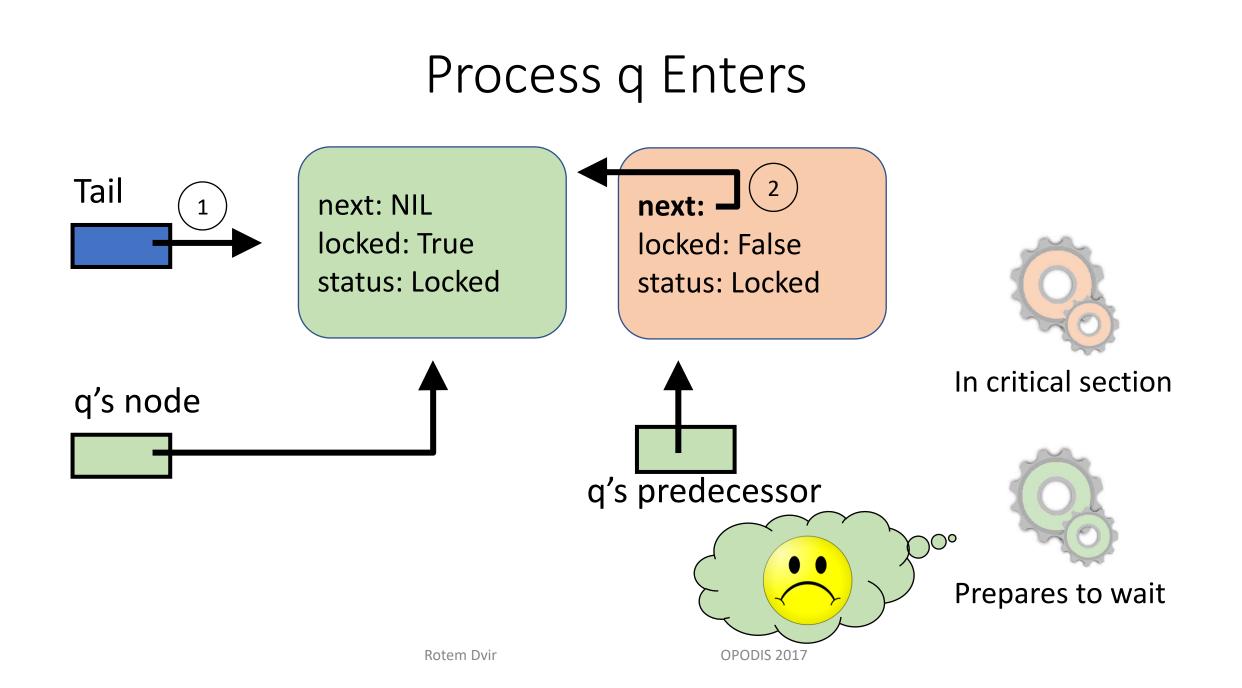


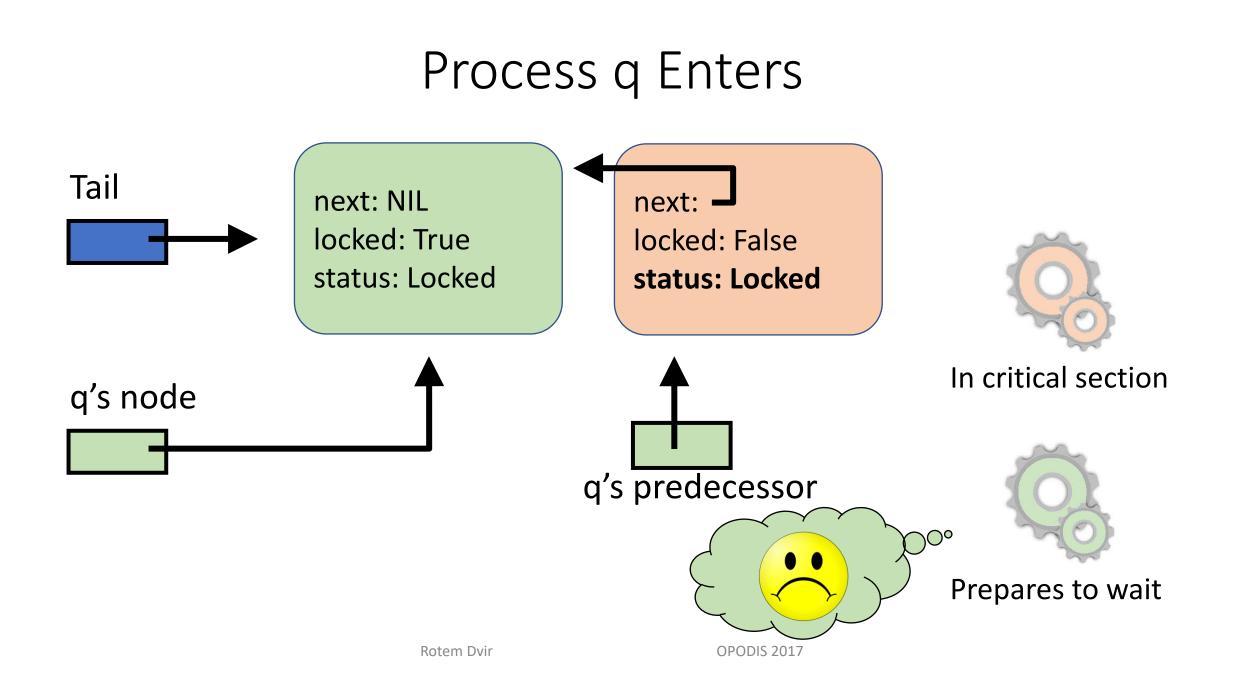
Initializing

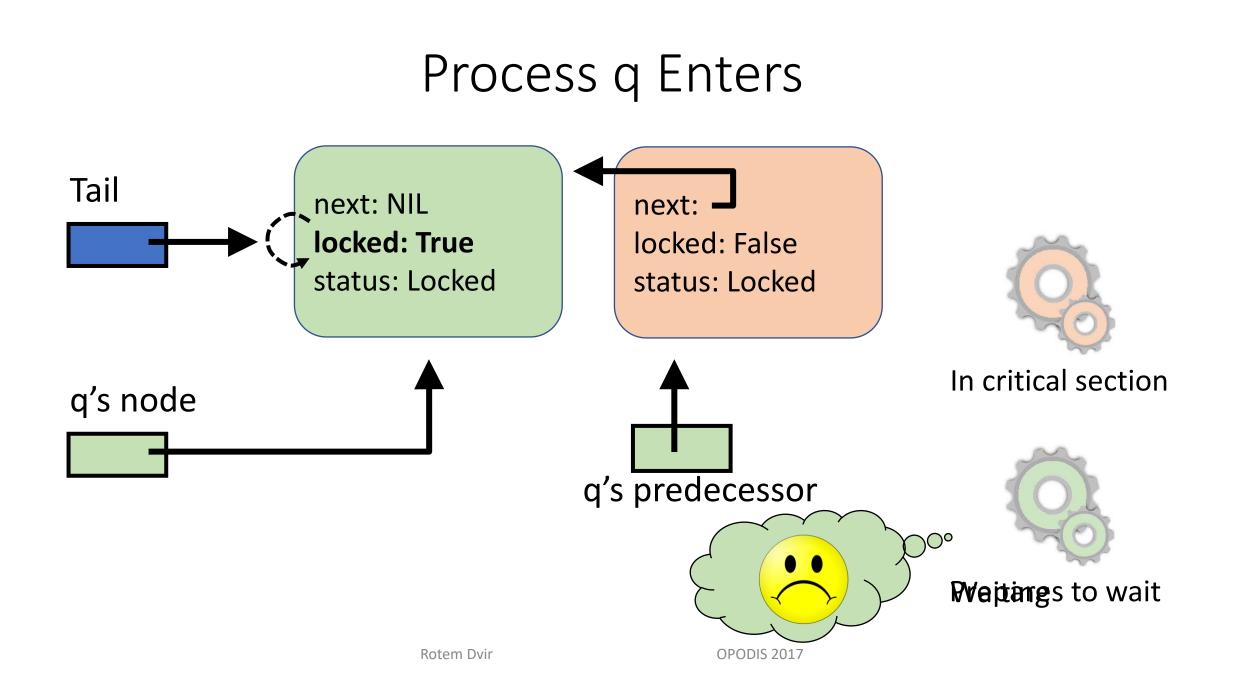


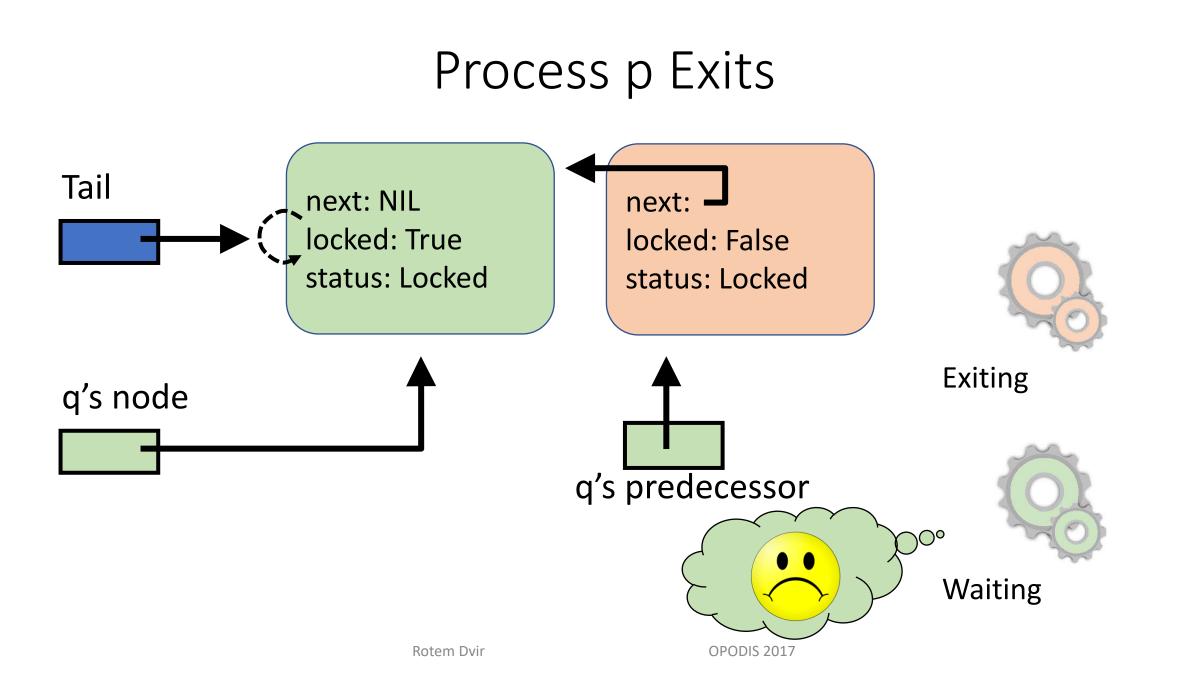


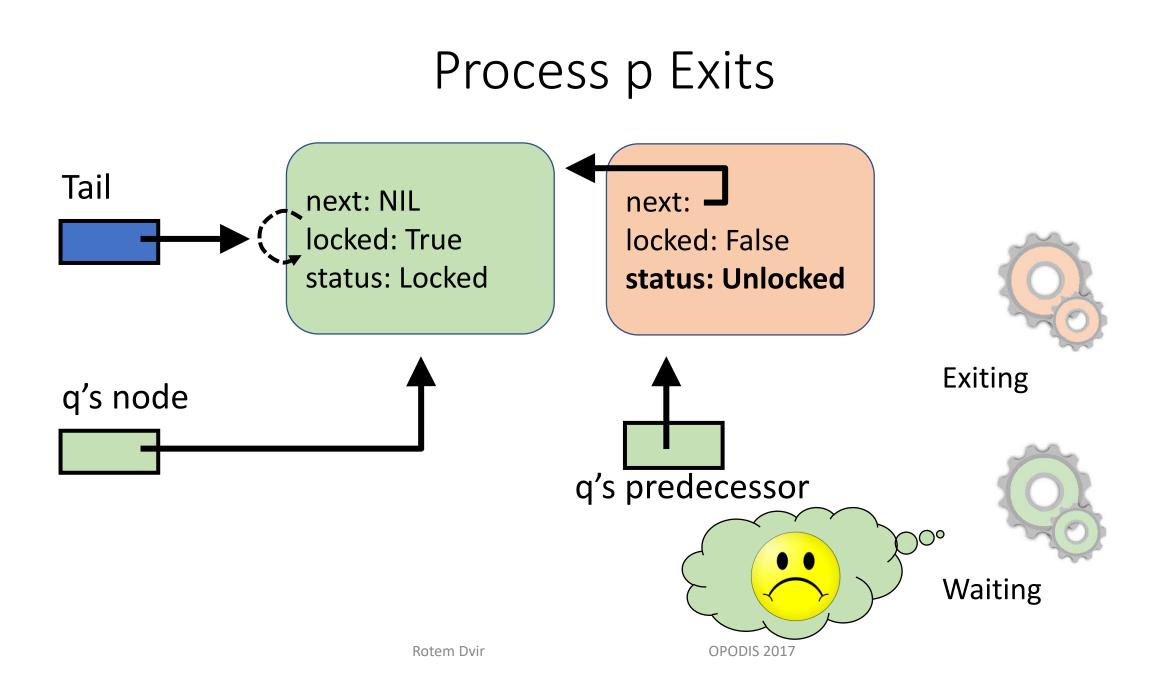


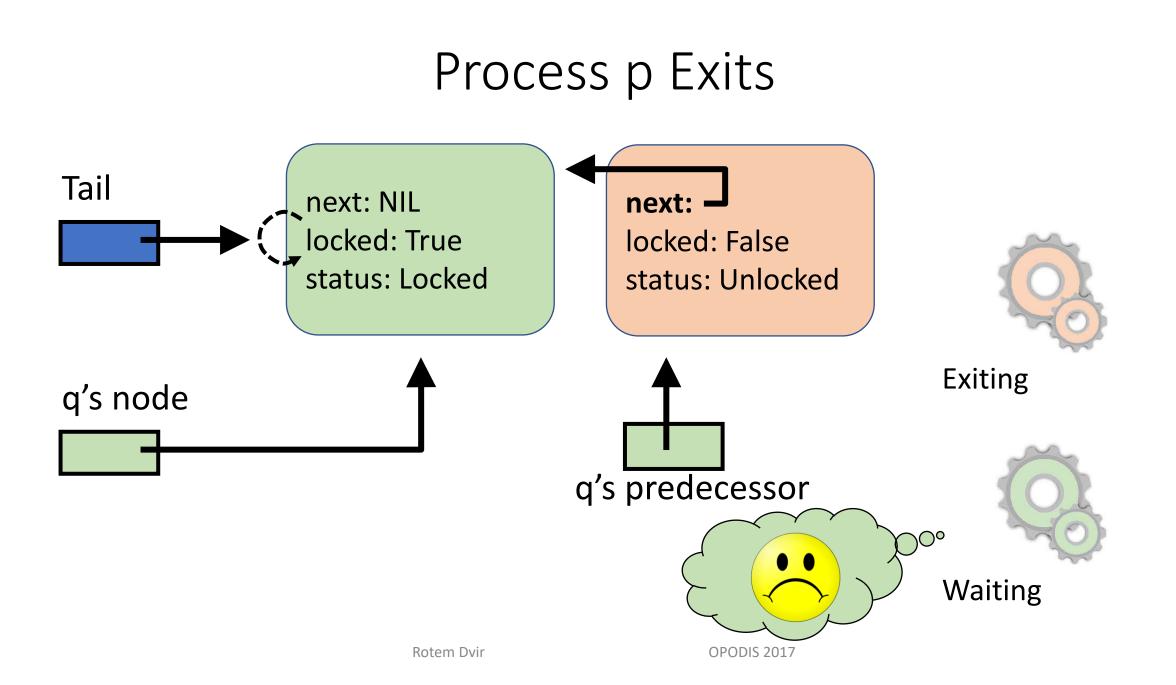


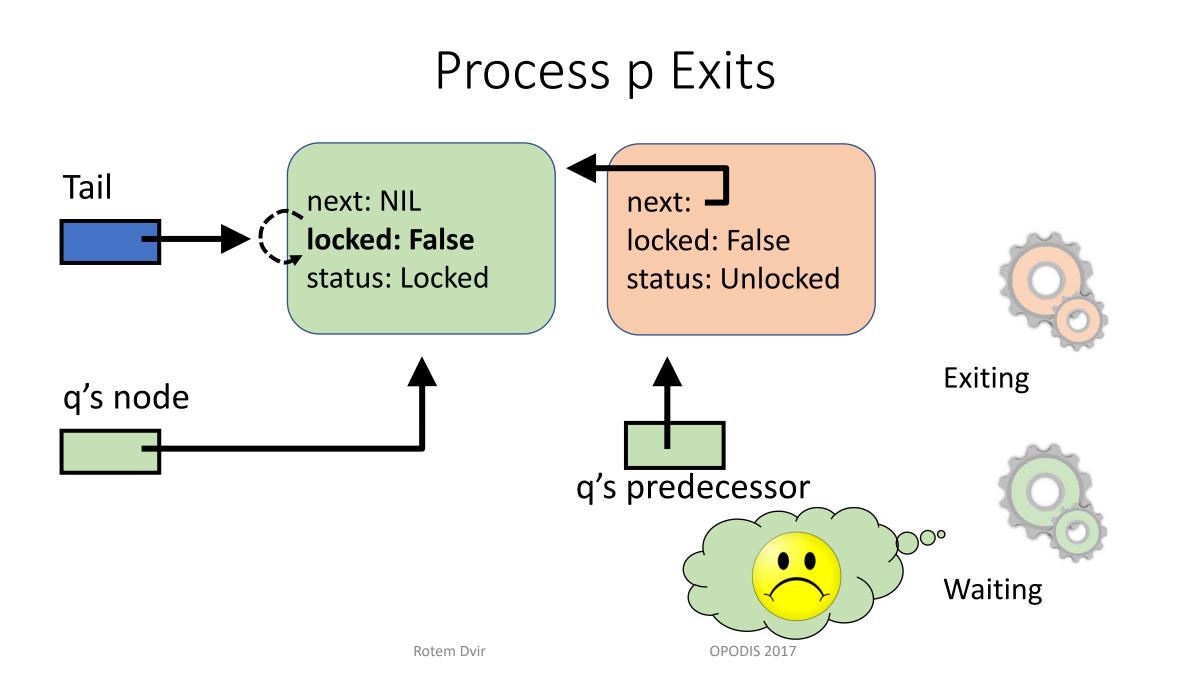


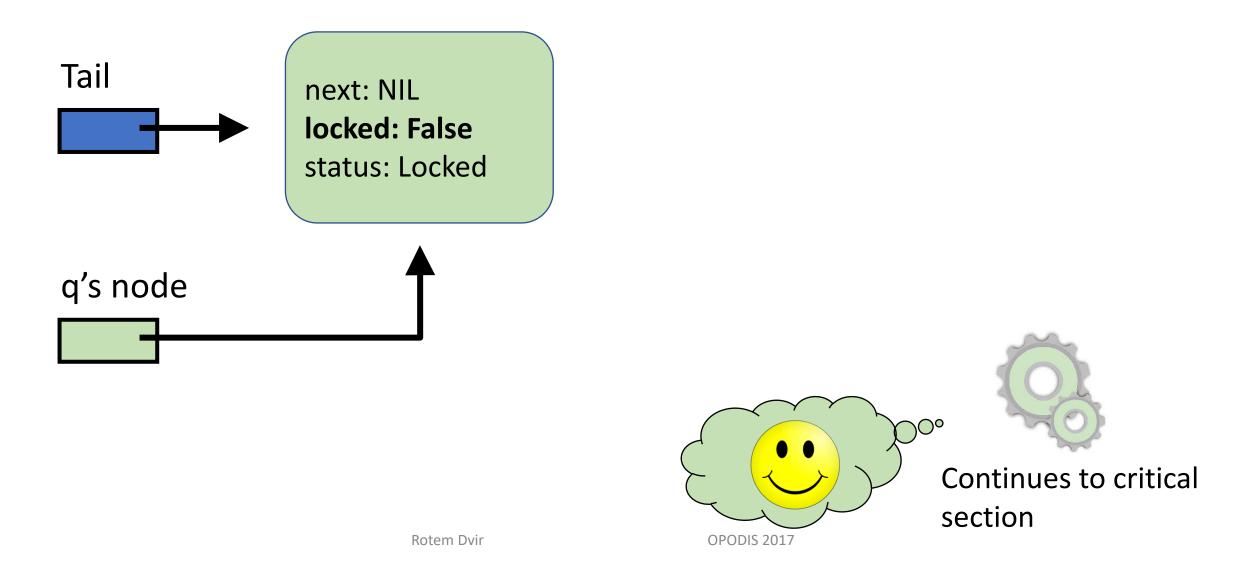




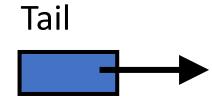








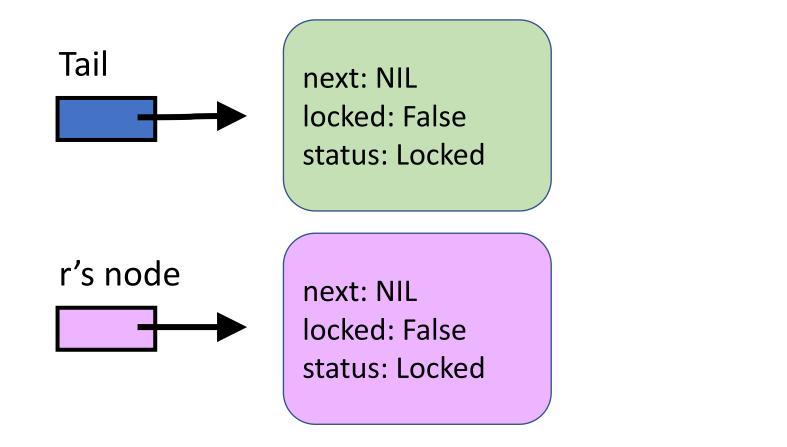
Process q In Critical Section



next: NIL locked: False status: Locked



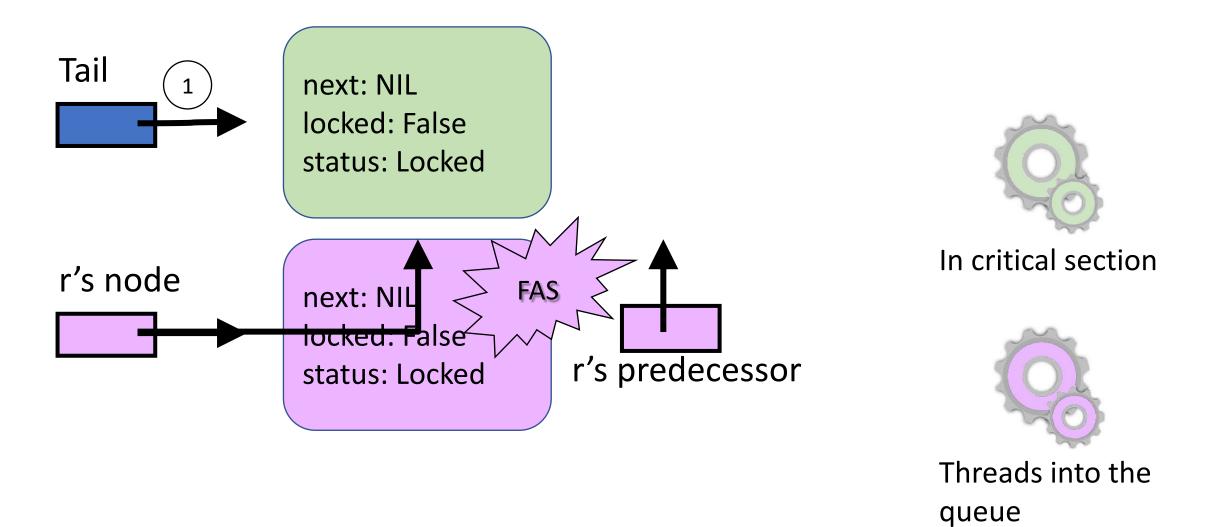
In critical section

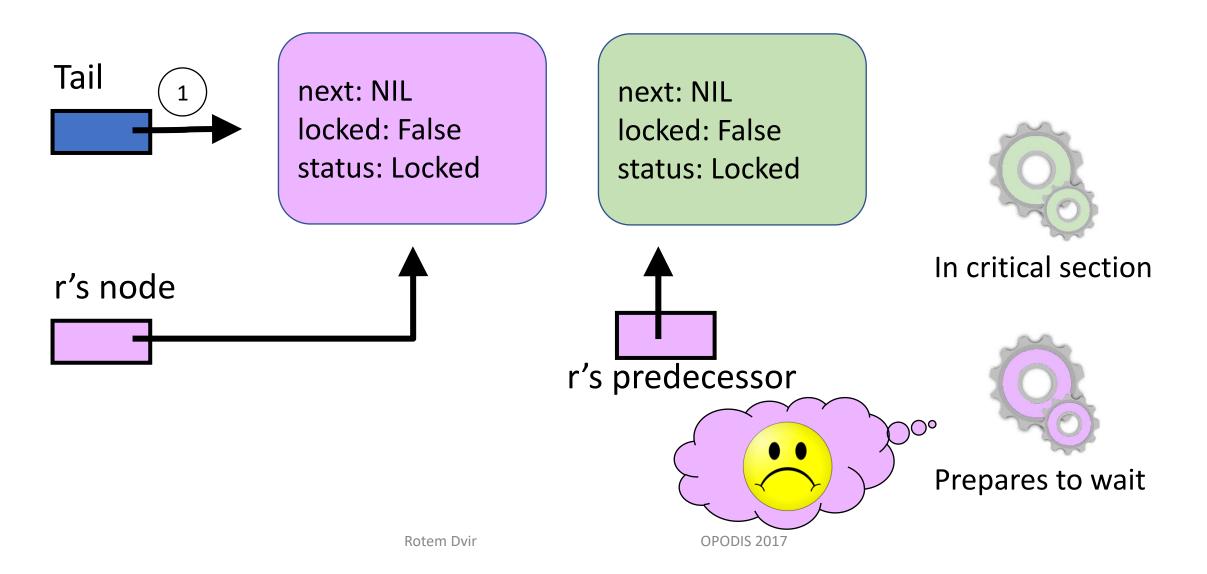


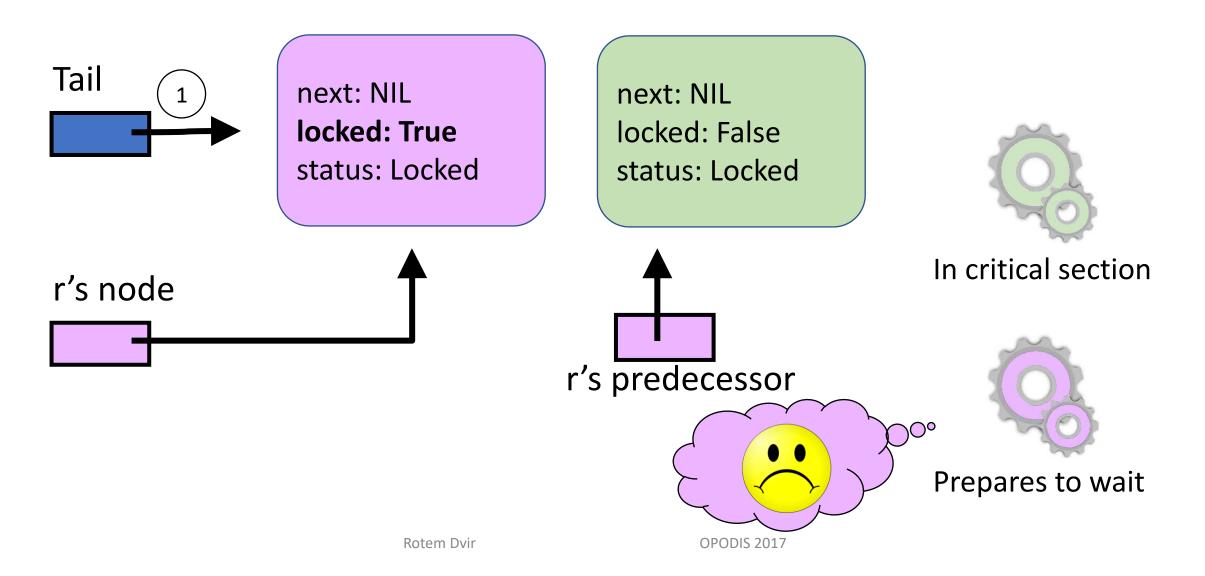


In critical section

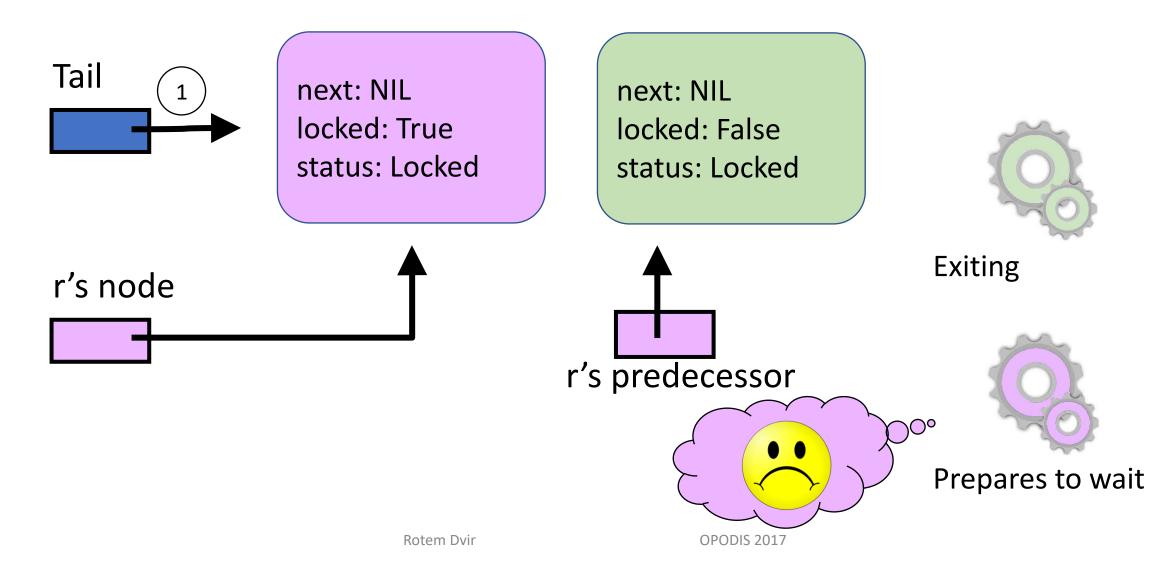




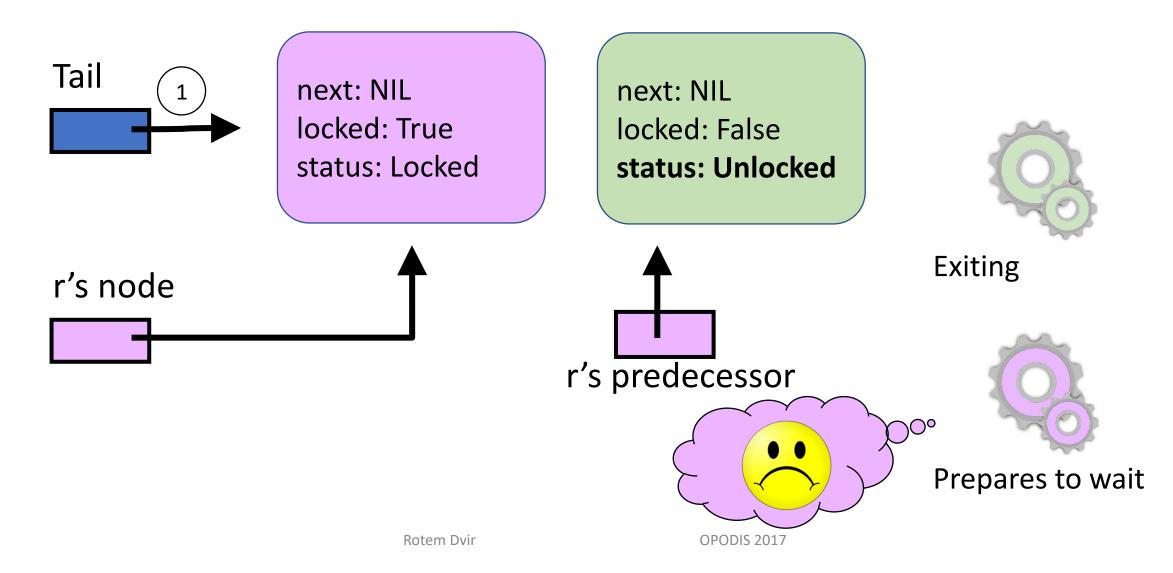




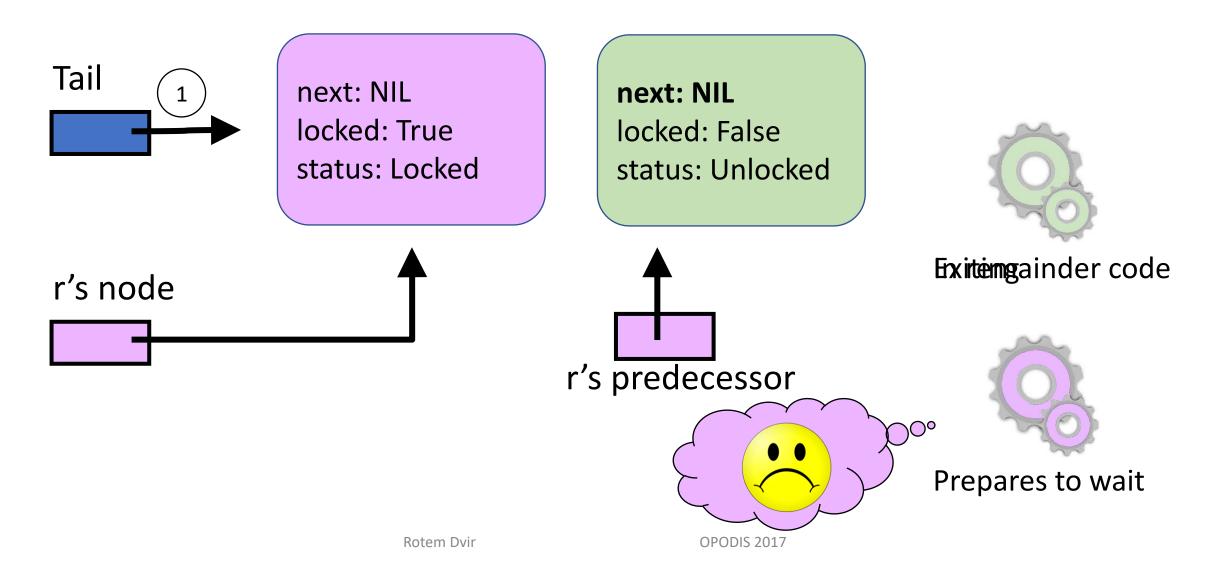
Process q Exits

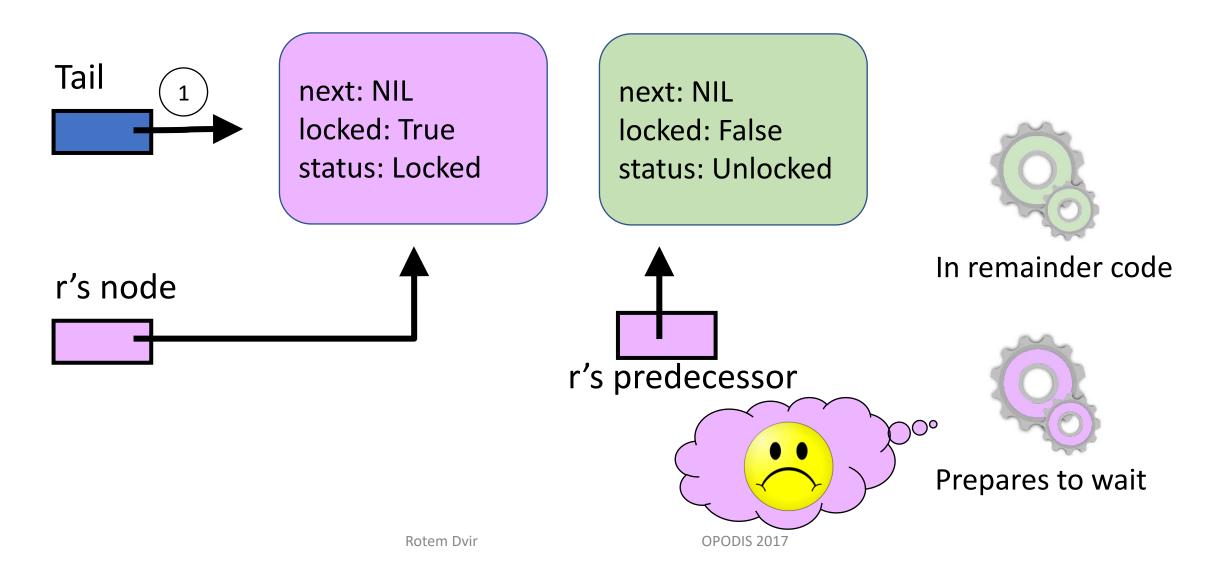


Process q Exits

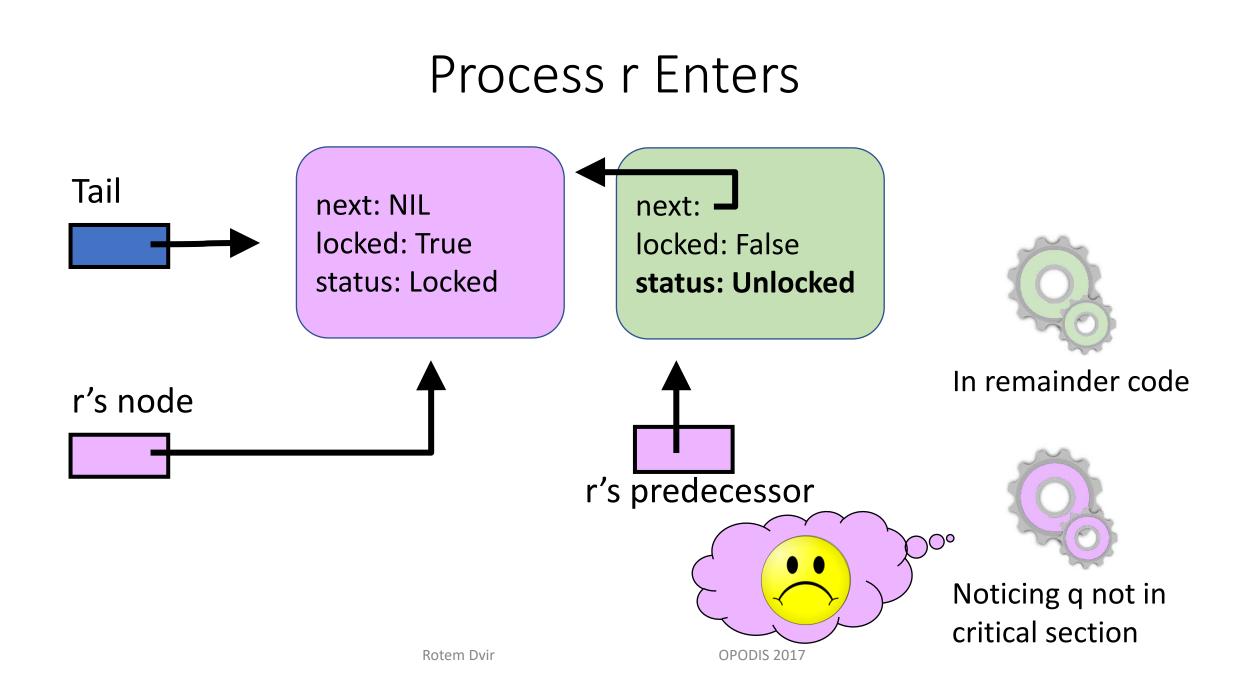


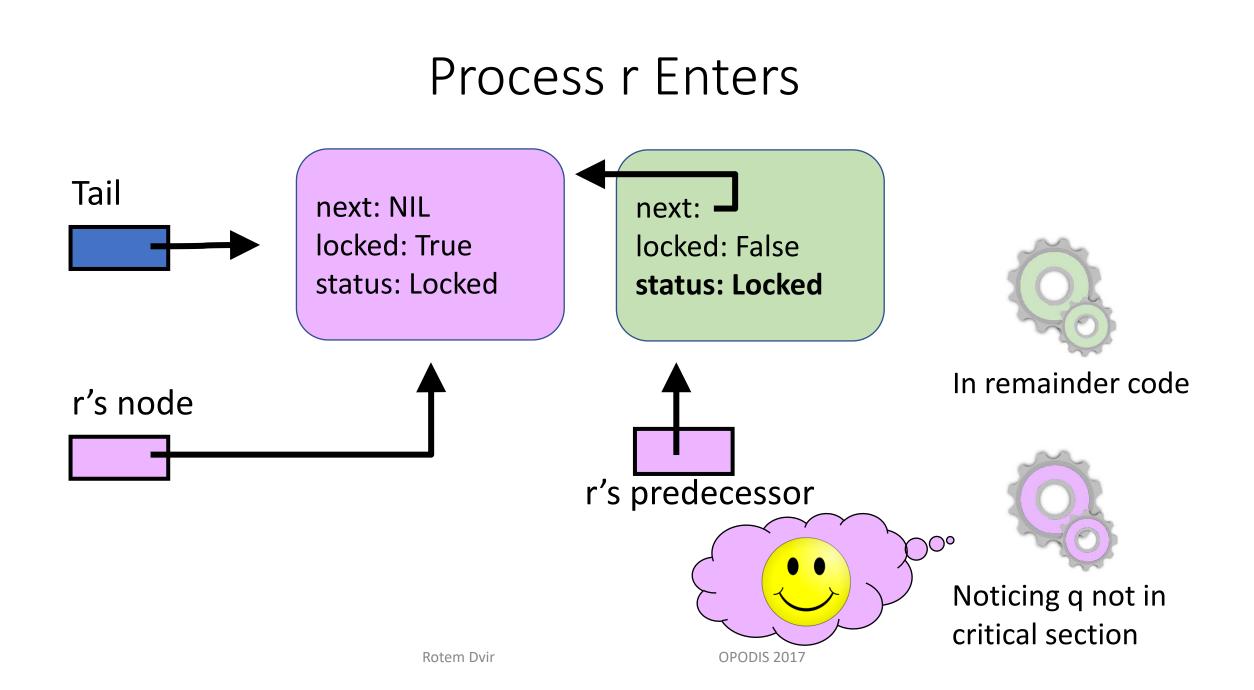
Process q Exits

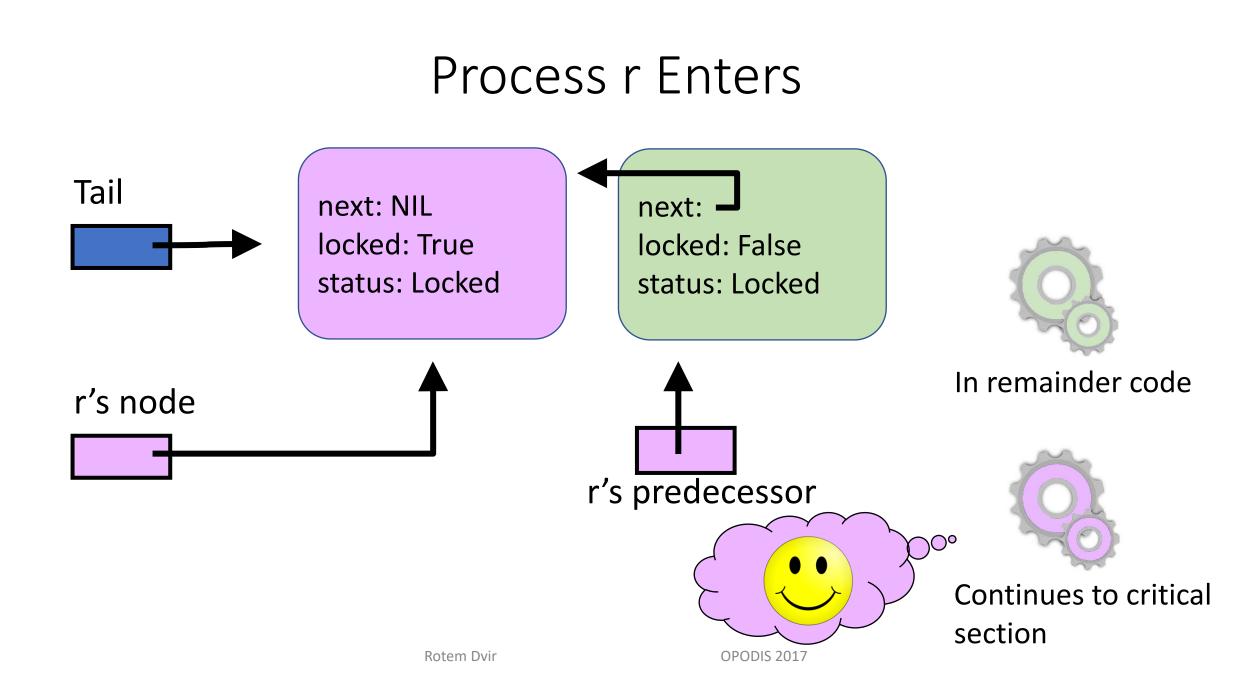




Process r Enters next: \square ² Tail next: NIL 1 locked: True locked: False status: Locked status: Unlocked In remainder code r's node r's predecessor Prepares to wait **OPODIS 2017** Rotem Dvir







Summary

Wait-free exit codes Satisfy FIFO fairness Constant RMR complexity in CC and DSM models Not assumed that the number of processes is priori known Use only O(n) shared memory locations Make no assumptions on what and how memory is allocated Instruction set: Read/Write, Fetch-And-Store, Compare-And-Swap

Open Questions

Algorithms extensions:

- Readers-writers lock
- Group mutual exclusion
- Abortable
- Recoverable



