



departamento de informática
FACULDADE DE CIÊNCIAS E TECNOLOGIA
UNIVERSIDADE NOVA DE LISBOA

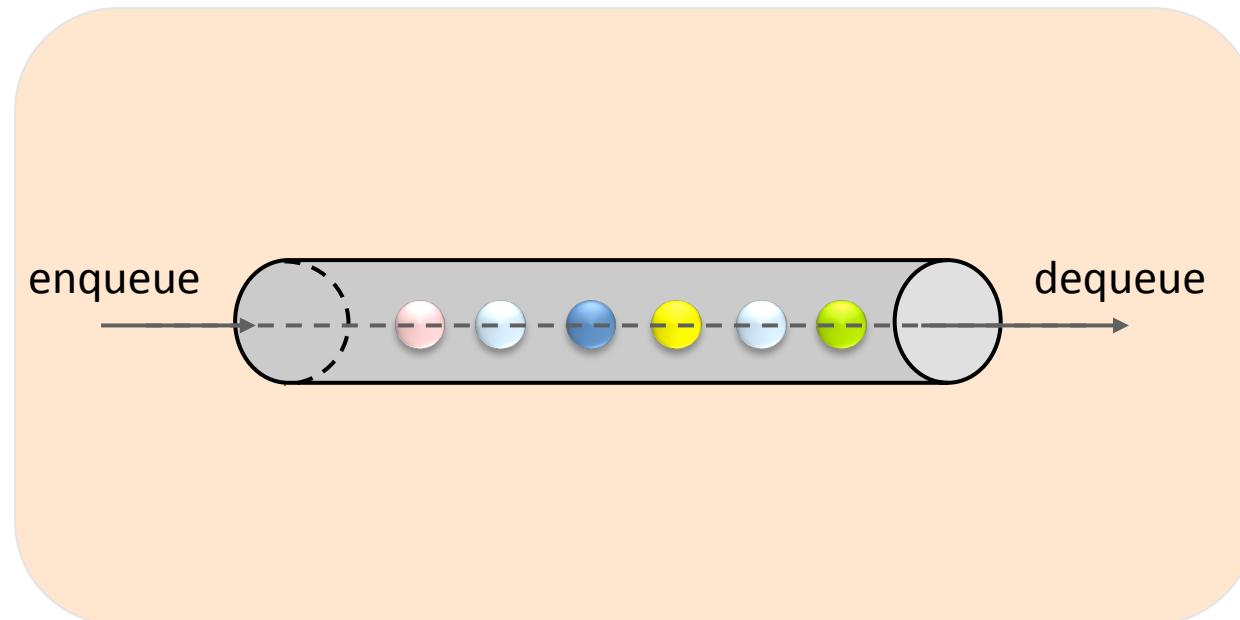
Concurrency and Parallelism **(Concorrência e Paralelismo – CP 11158)**



Lecture 4 — Lock-free Data Structures —

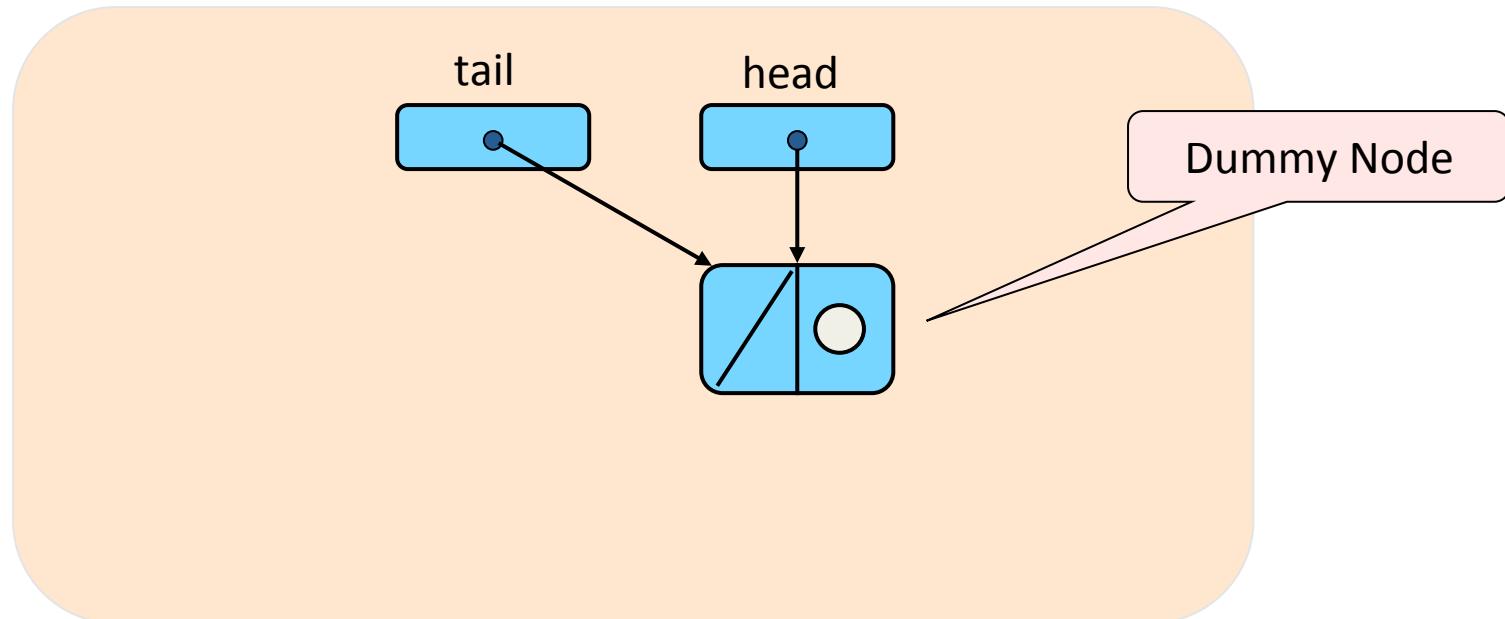
*Slides based in material from:
Gadi Taubenfeld (<http://www.faculty.idc.ac.il/gadi/book.htm>)*

Implementing a lock-free Queue

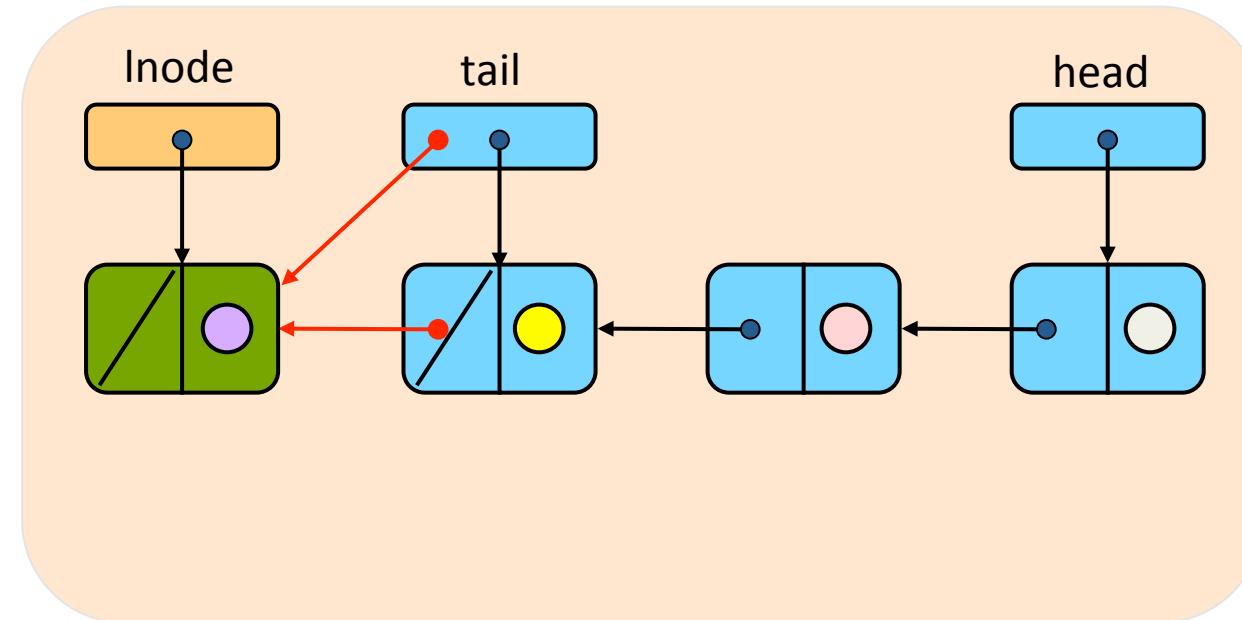


Implementing a lock-free Queue

Empty queue

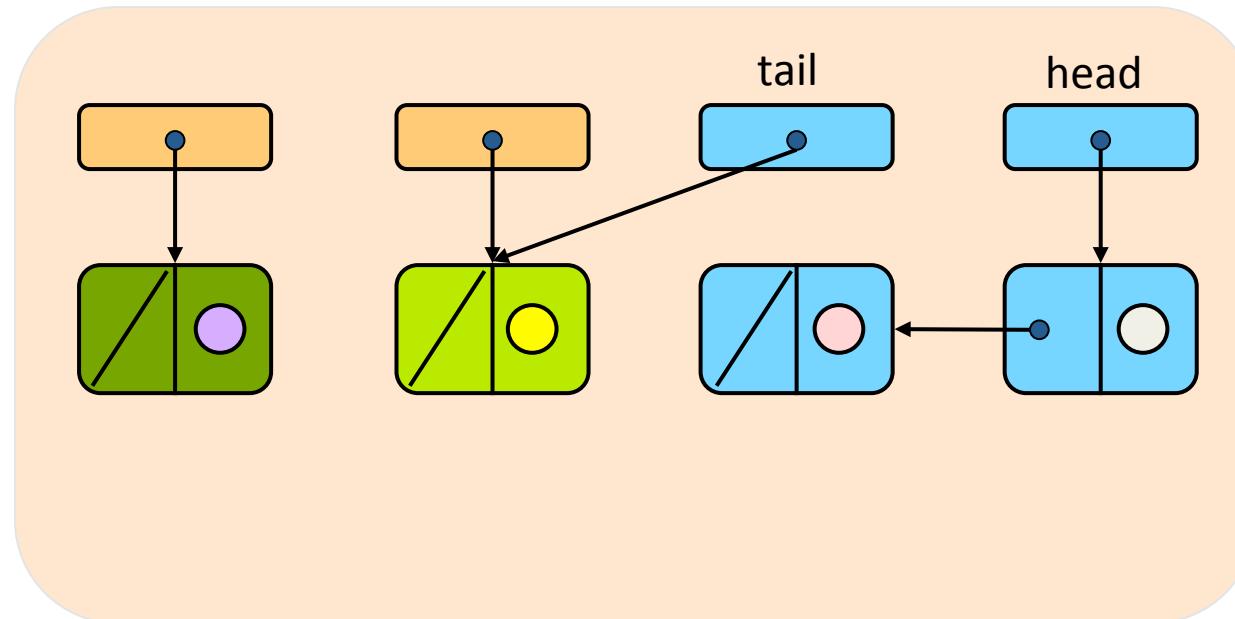


Enqueue

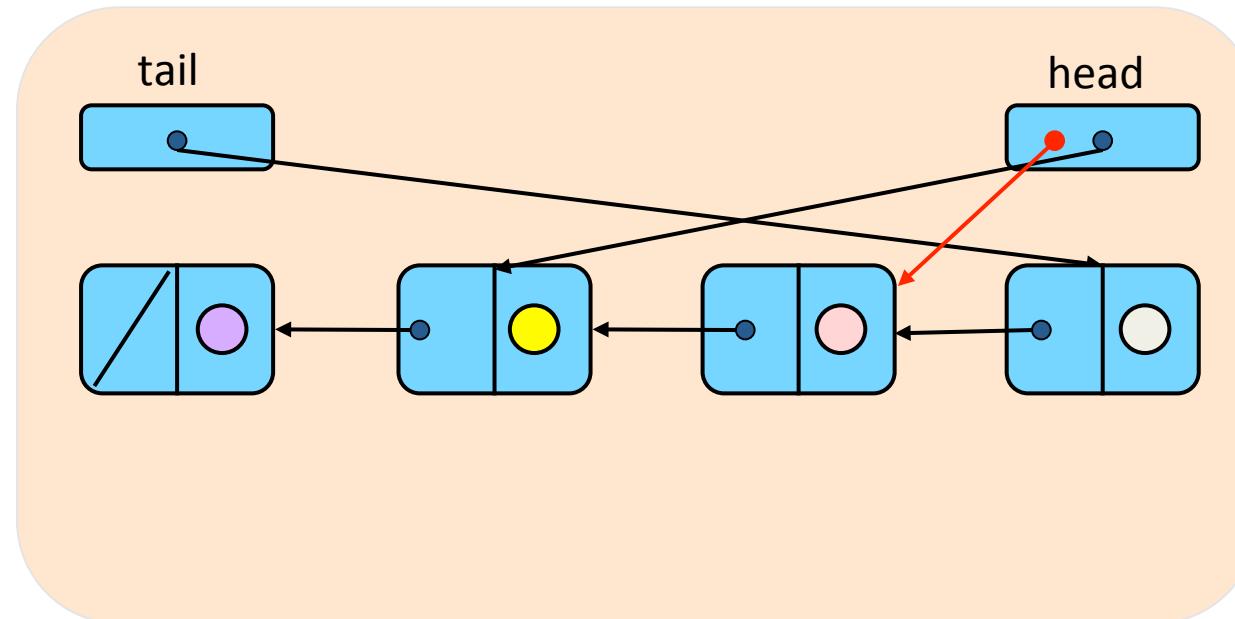
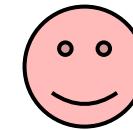


2 concurrent enqueue operations

What can go wrong when using registers?

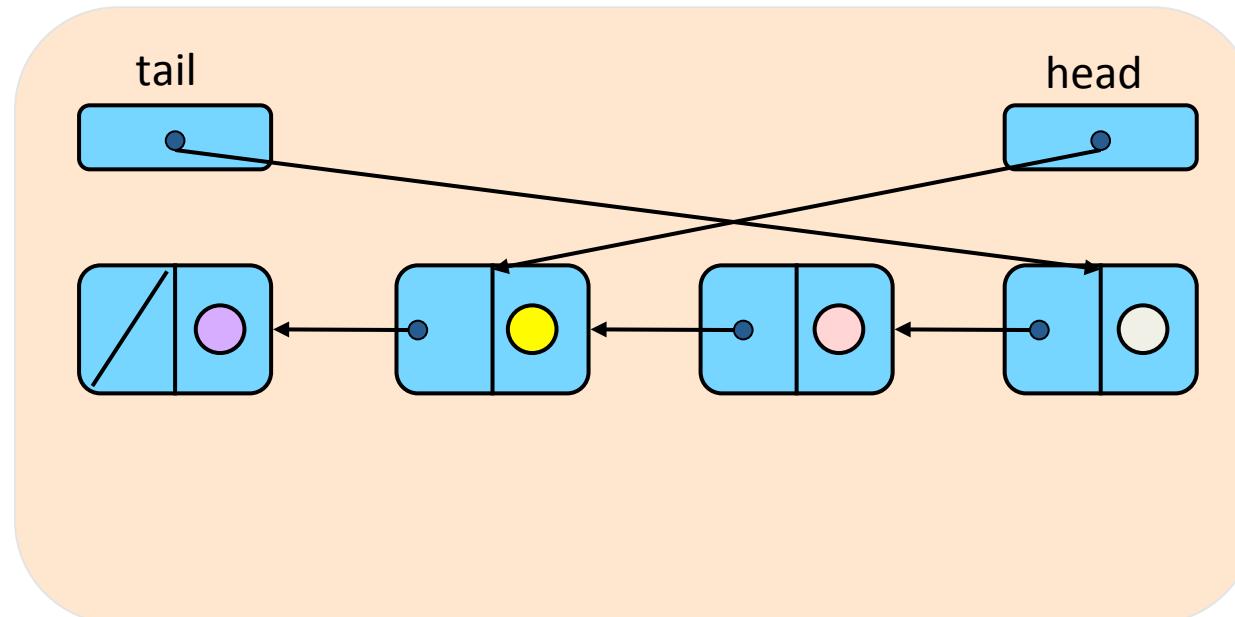
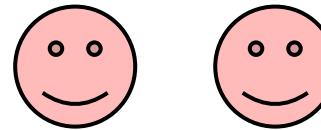


Dequeue

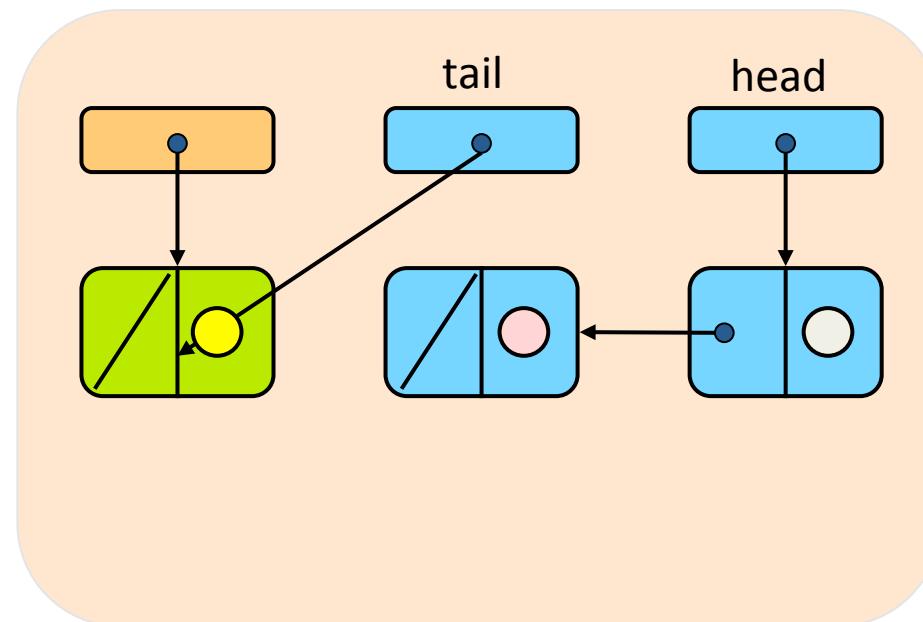
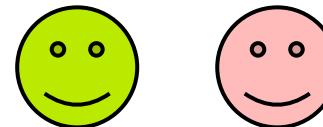


2 concurrent dequeue operations

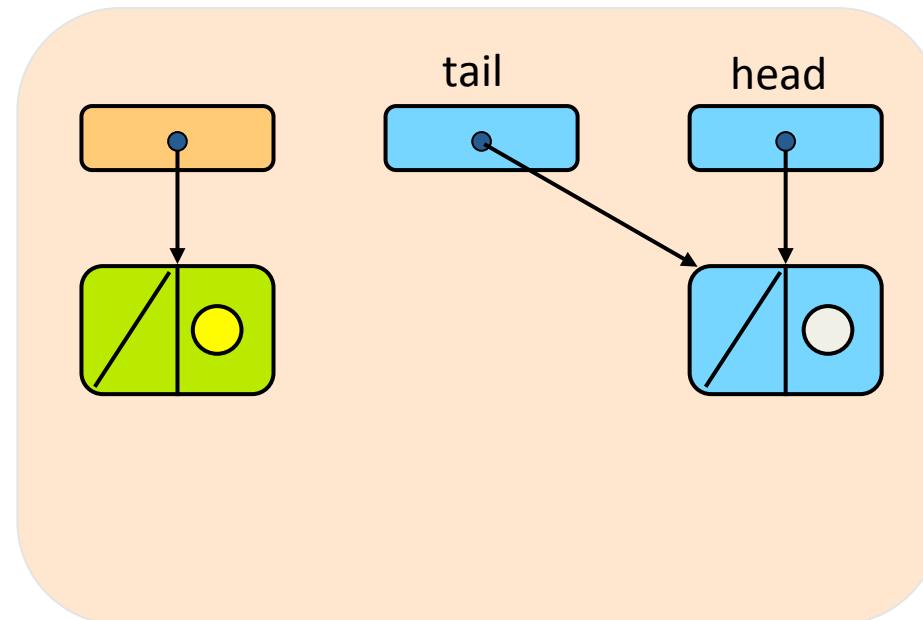
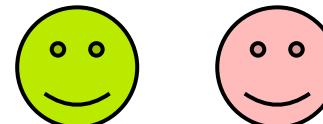
What can go wrong when using registers?



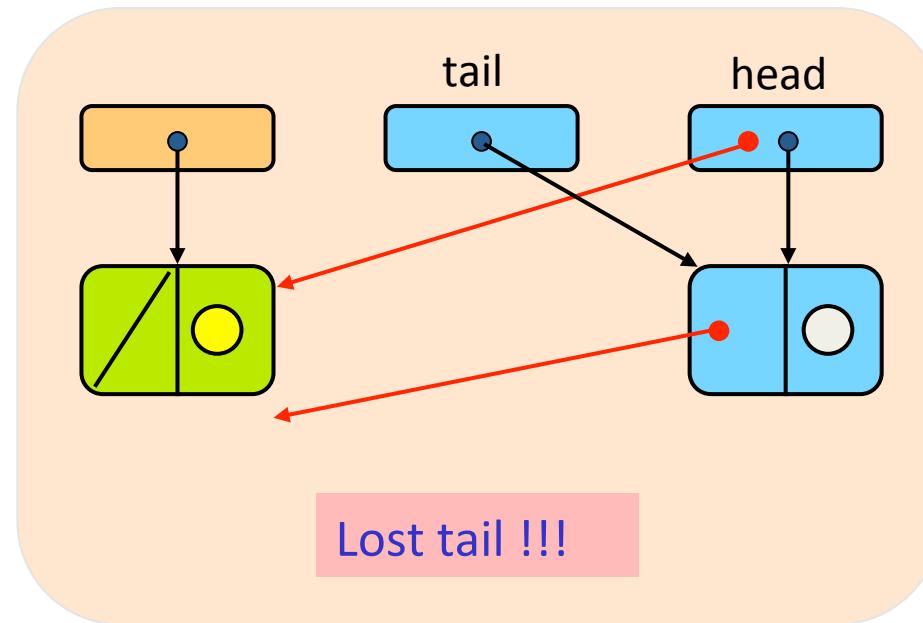
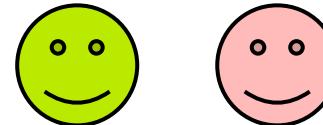
Enqueue + Dequeue (Non-empty queue)



Enqueue + Dequeue (Empty queue)

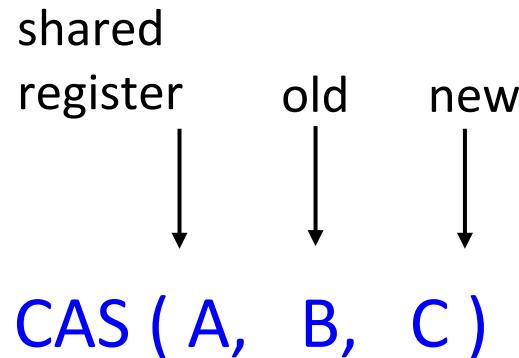


Enqueue + Dequeue (Empty queue)



The dequeuer must advance tail before redirecting head

Compare & swap (CAS)

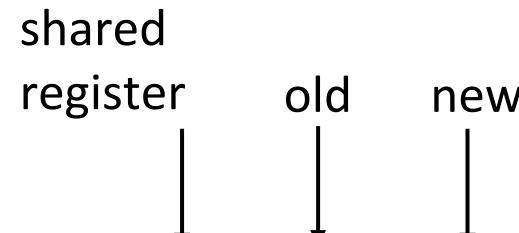


```
if A=B then A:=C; return(true)  
else return(false)
```

Supported by Sun, Intel, AMD, ...

CAS: The ABA problem

shared register old new



CAS (A, B, C)

if $A=B$ then $A:=C$

local := A

CAS (A, local, 100)

local:

value	tag
-------	-----

A:

value	tag
-------	-----

CAS (A, local, <100,local.tag+1>)

In the following, we will ignore the ABA problem and assume that it is resolved using additional tag fields.

A Lock-Free (non-blocking) Queue

- The algorithm is due to M. M. Michael and M. L. Scott (1996).
- The algorithm is included in the Standard Java Concurrency Package.
- Every process must be prepared to encounter a half-finished enqueue operation, and help to finish it.
- Invariant: The tail refers to either the last node or to the node just before the last node.

Enqueue

repeat

$ltail \leftarrow tail$

$lnext \leftarrow ltail.next$

 if $ltail = tail$ then

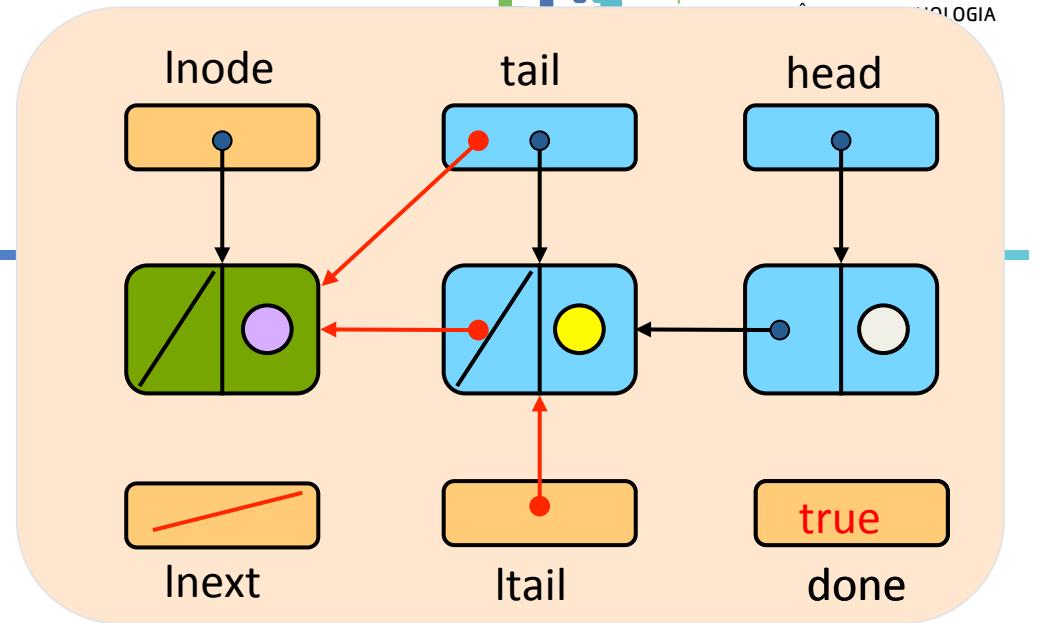
 if $lnext = NULL$ then

 if CAS ($ltail.next$, $lnext$, $Inode$) then $done \leftarrow true$ fi

 else CAS ($tail$, $ltail$, $lnext$) fi fi

 until $done = true$

 CAS ($tail$, $ltail$, $Inode$)



Enqueue

repeat

$ltail \leftarrow tail$

$lnext \leftarrow ltail.next$

 if $ltail = tail$ then

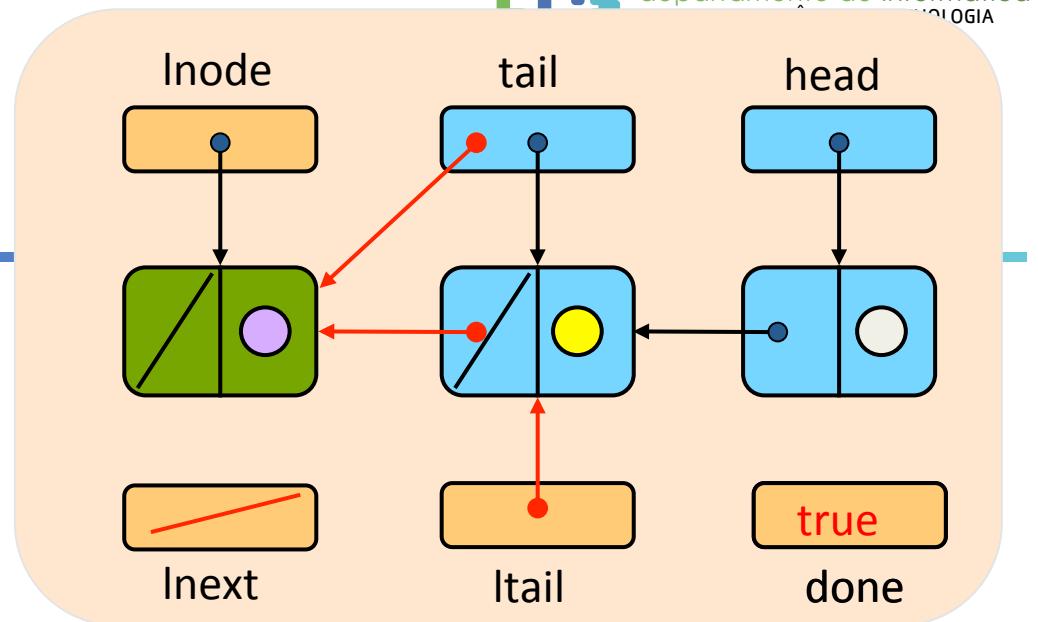
 if $lnext = \text{NULL}$ then

 if CAS ($ltail.next$, $lnext$, node) then $done \leftarrow \text{true}$ fi

 else CAS ($tail$, $ltail$, $lnext$) fi fi

 until $done = \text{true}$

CAS ($tail$, $ltail$, $inode$)



Question: Would the enqueue operation be correct if the last line is omitted?

Answer: Yes, but...

Enqueue

repeat

$ltail \leftarrow tail$

$lnext \leftarrow ltail.next$

 if $ltail = tail$ then

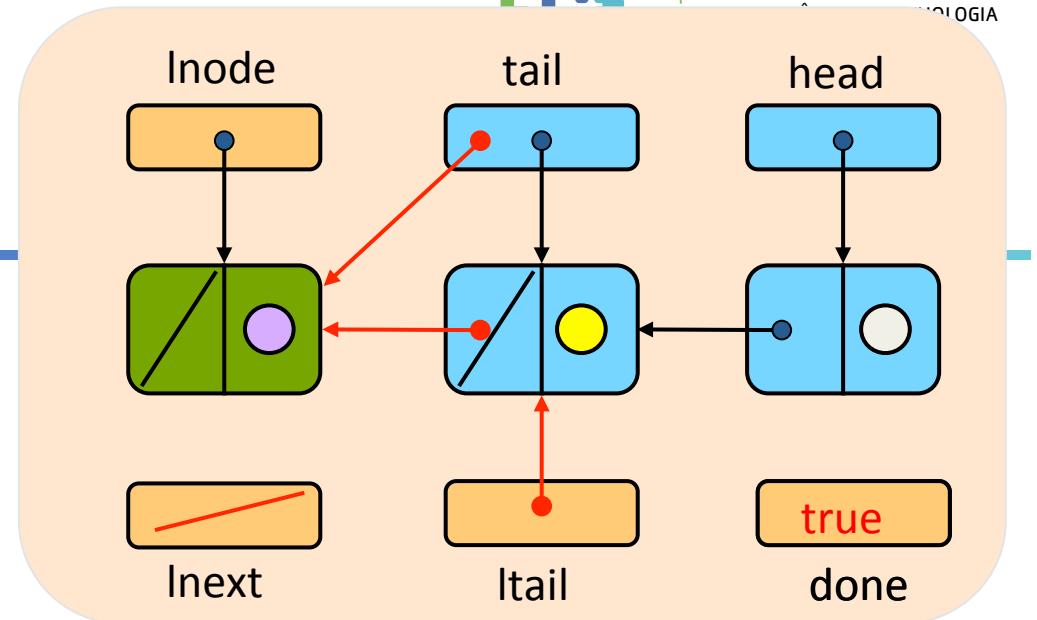
 if $lnext = \text{NULL}$ then

 if CAS ($ltail.next$, $lnext$, node) then $done \leftarrow \text{true}$ fi

 else CAS ($tail$, $ltail$, $lnext$) fi fi

 until $done = \text{true}$

CAS ($tail$, $ltail$, $Inode$)



Question: Would the enqueue operation be correct if this line is omitted?

Yes.

Dequeue

repeat

 lhead \leftarrow head

 ltail \leftarrow tail

 lnext \leftarrow lhead.next

 if lhead = head then

 if lhead = ltail then /* empty or tail behind?

 if lnext = NULL then return(NULL) fi ; /* empty

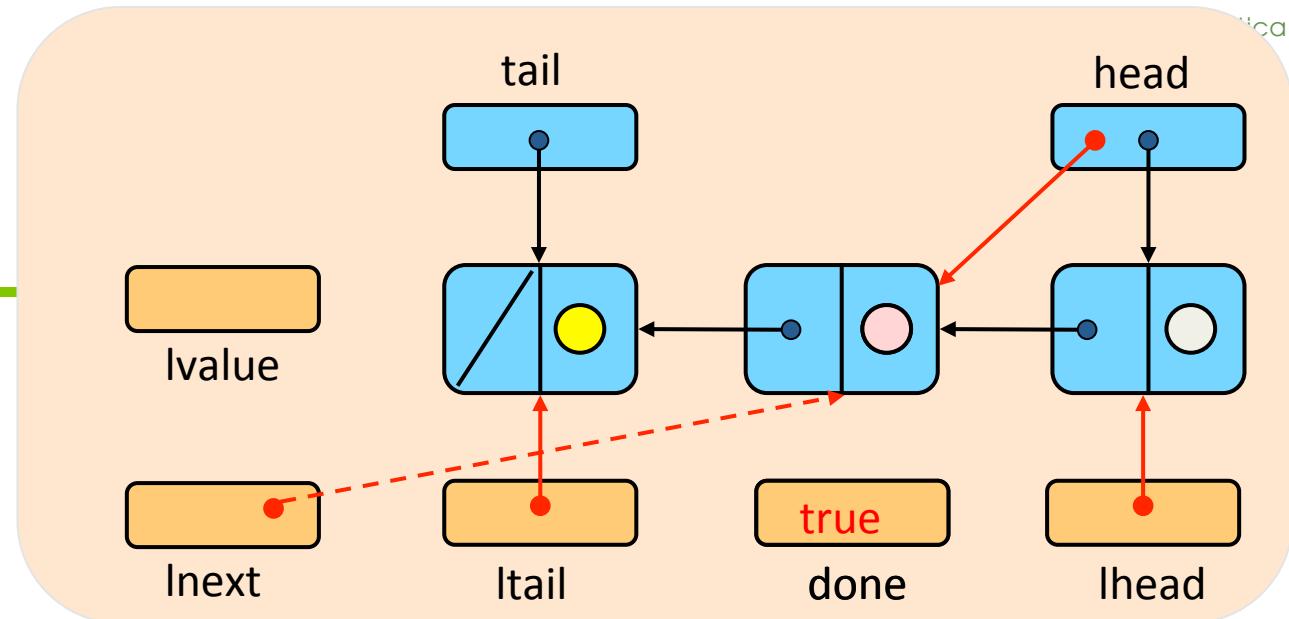
 CAS (tail, ltail, lnext) /* try to advance tail

 else lvalue \leftarrow lnext.value /* no need to deal with tail

 if CAS (head, lhead, lnext) then done \leftarrow true fi fi fi

 until done = true

 free(lhead) ; return(lvalue)



Dequeue

repeat

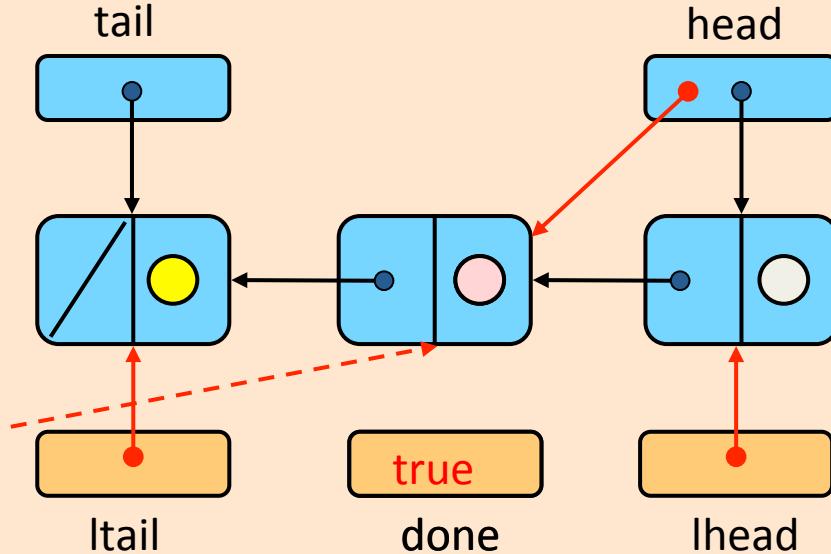
 lhead \leftarrow head

 ltail \leftarrow tail

 lnext \leftarrow lhead.next

 if lhead = head then

 if lhead = ltail then



Question: Would the dequeue operation be correct if this line is omitted?

 /* empty or tail behind?

 if Inext = NULL then return(NULL) fi ; /* empty

 CAS (tail, ltail, Inext) /* try to advance tail

 else lvalue \leftarrow Inext.value /* no need to deal with tail

 if CAS (head, lhead, Inext) then done \leftarrow true fi fi fi

until done = true

free(lhead) ; return(lvalue)

The End
