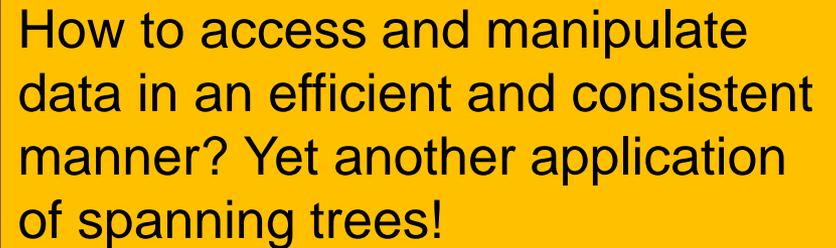


GIAN Course on Distributed Network Algorithms

Distributed Objects

GIAN Course on Distributed Network Algorithms

Distributed Objects

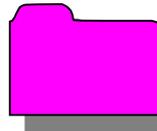
A yellow callout box with a black border and a pointer pointing towards the title 'Distributed Objects'.

How to access and manipulate data in an efficient and consistent manner? Yet another application of spanning trees!

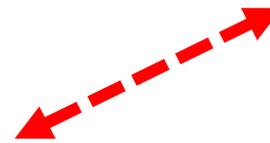
A fundamental challenge in distributed systems: coordinated access to a shared resource / object.



read-write

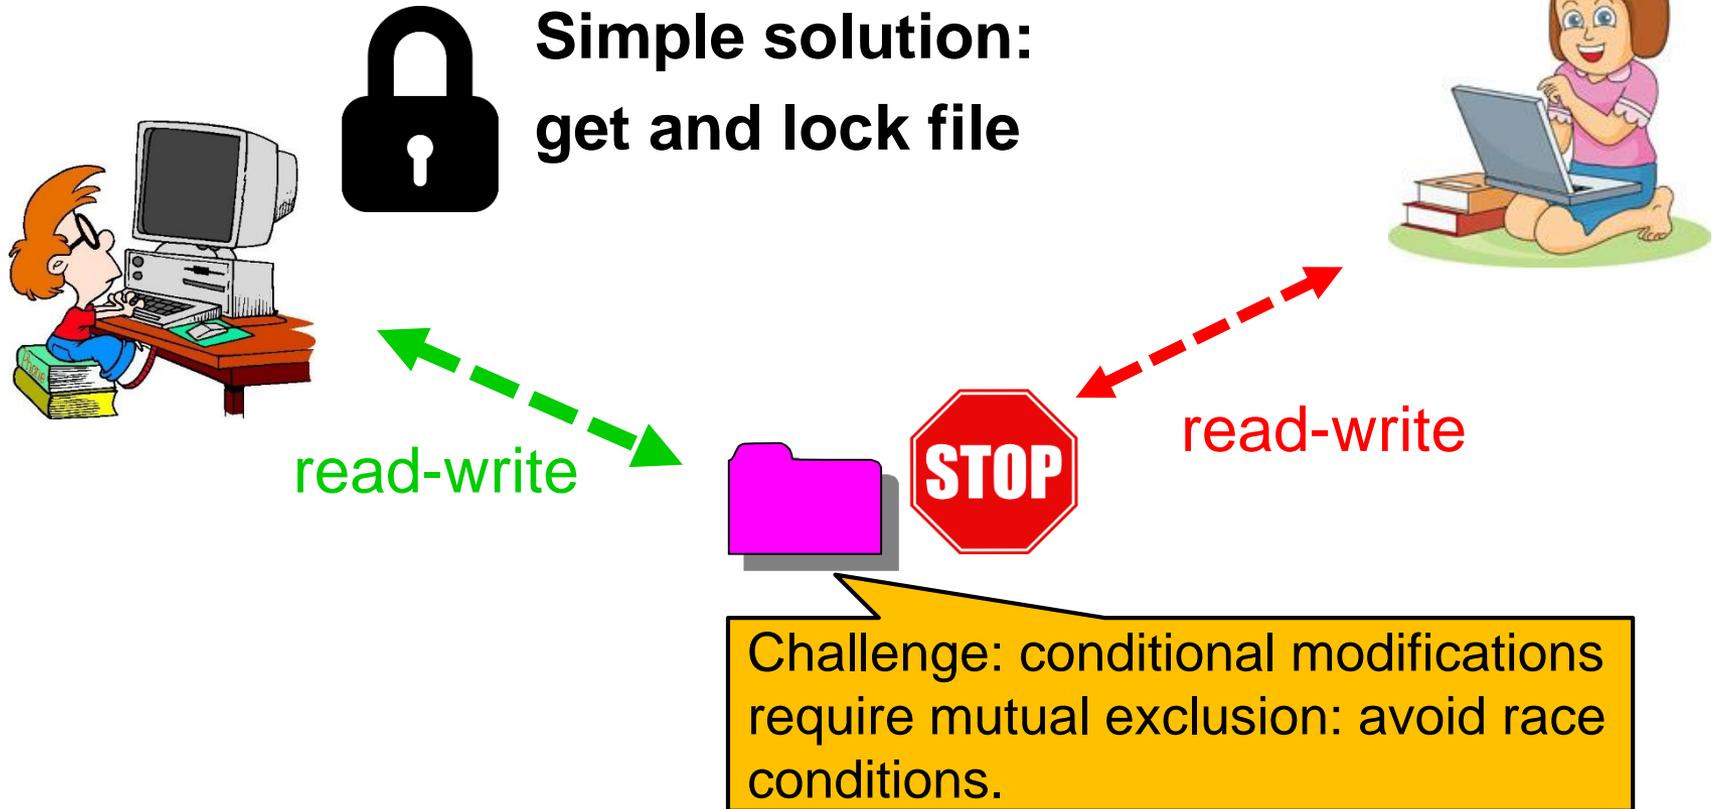


read-write

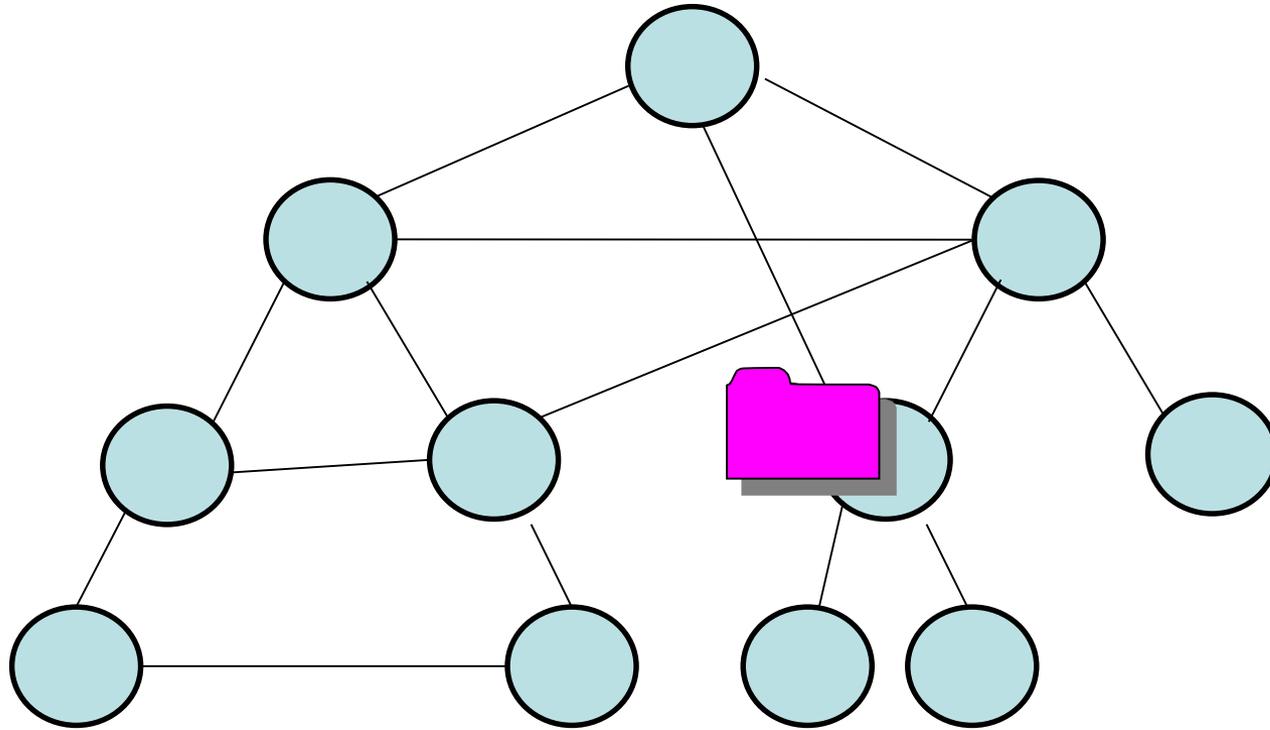


Challenge: conditional modifications require mutual exclusion: avoid race conditions.

A fundamental challenge in distributed systems: coordinated access to a shared resource / object.

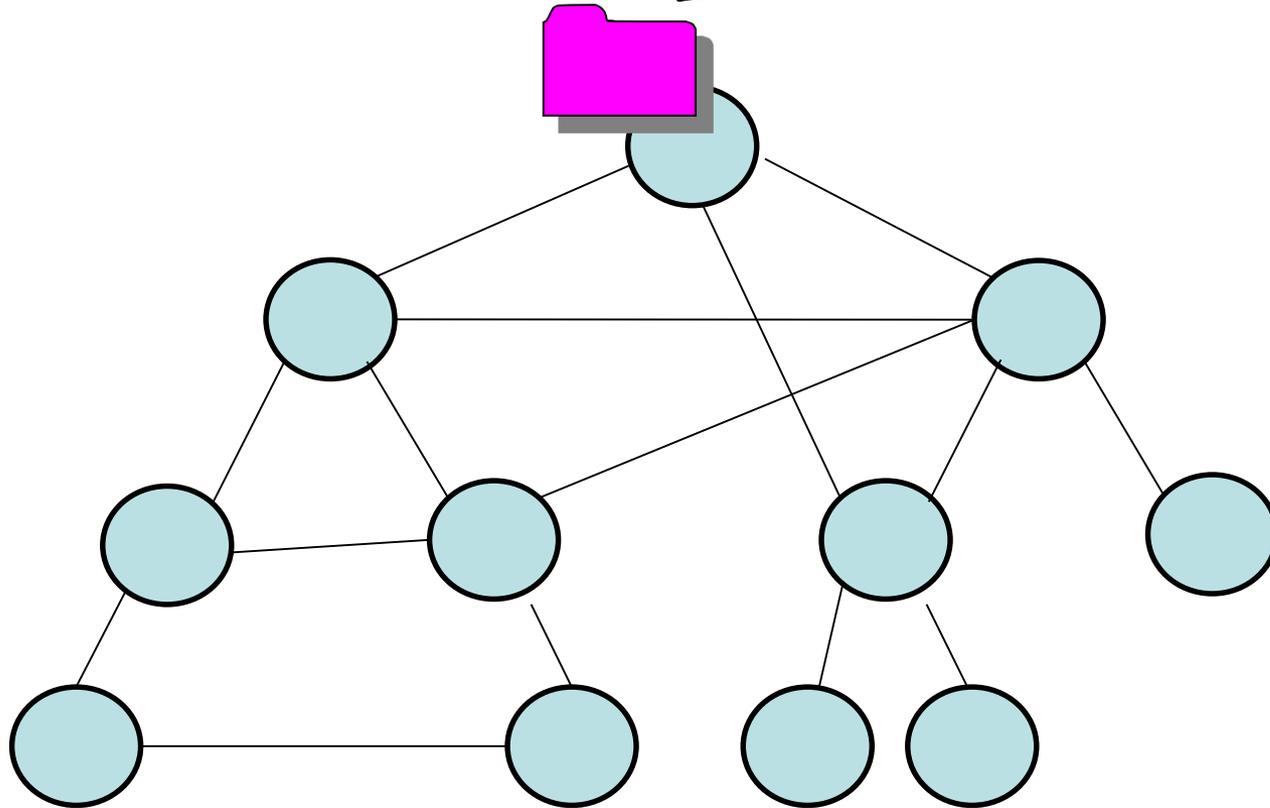


How to implement consistent read/write access *in a network*?



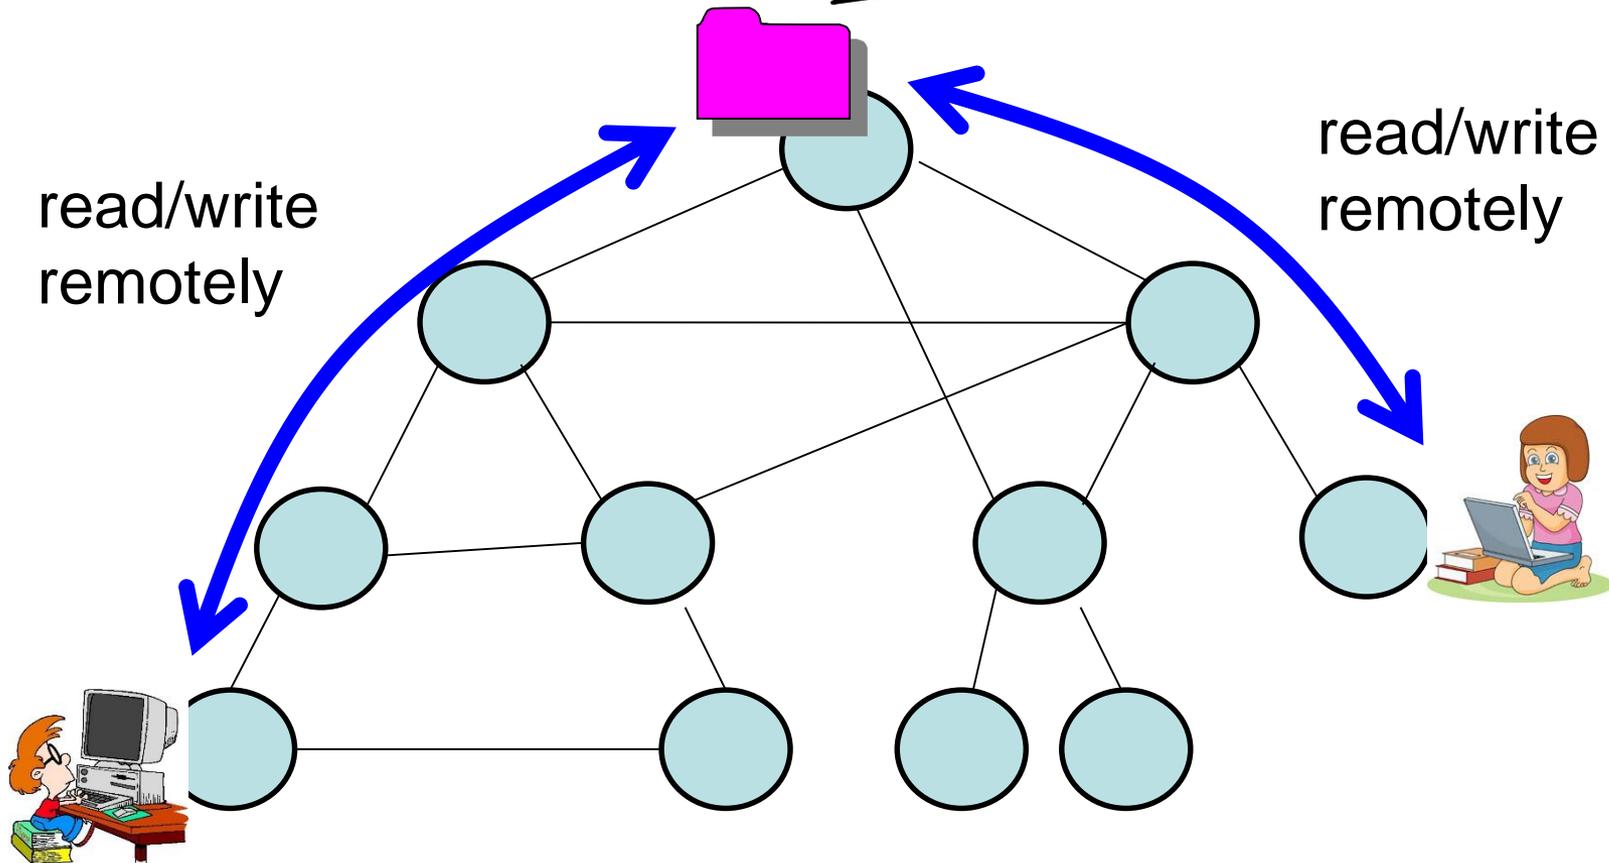
Naive Solution:

Idea: just keep a single copy at some *leader node*!



Naive Solution:

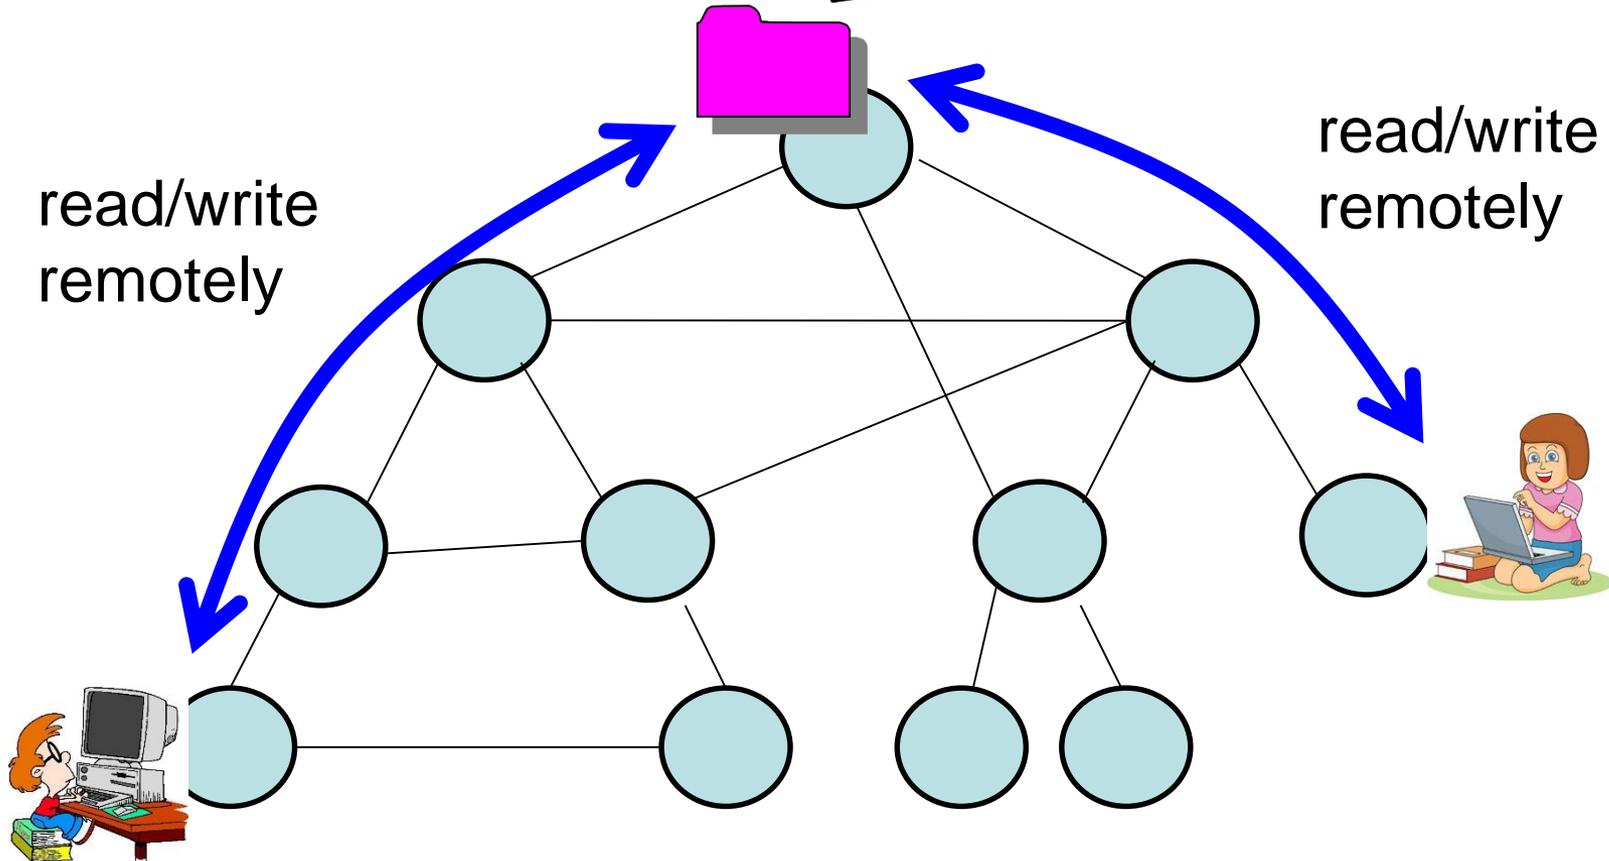
Idea: just keep a single copy at some *leader node*!



- Access:** 1. send message to leader,
2. leader processes request, 3. result sent back down the tree.

Problem 1: How to find leader?

Idea: just keep a single copy at some *leader node*!



Access: 1. send message to leader,
2. leader processes request, 3. result sent back down the tree.

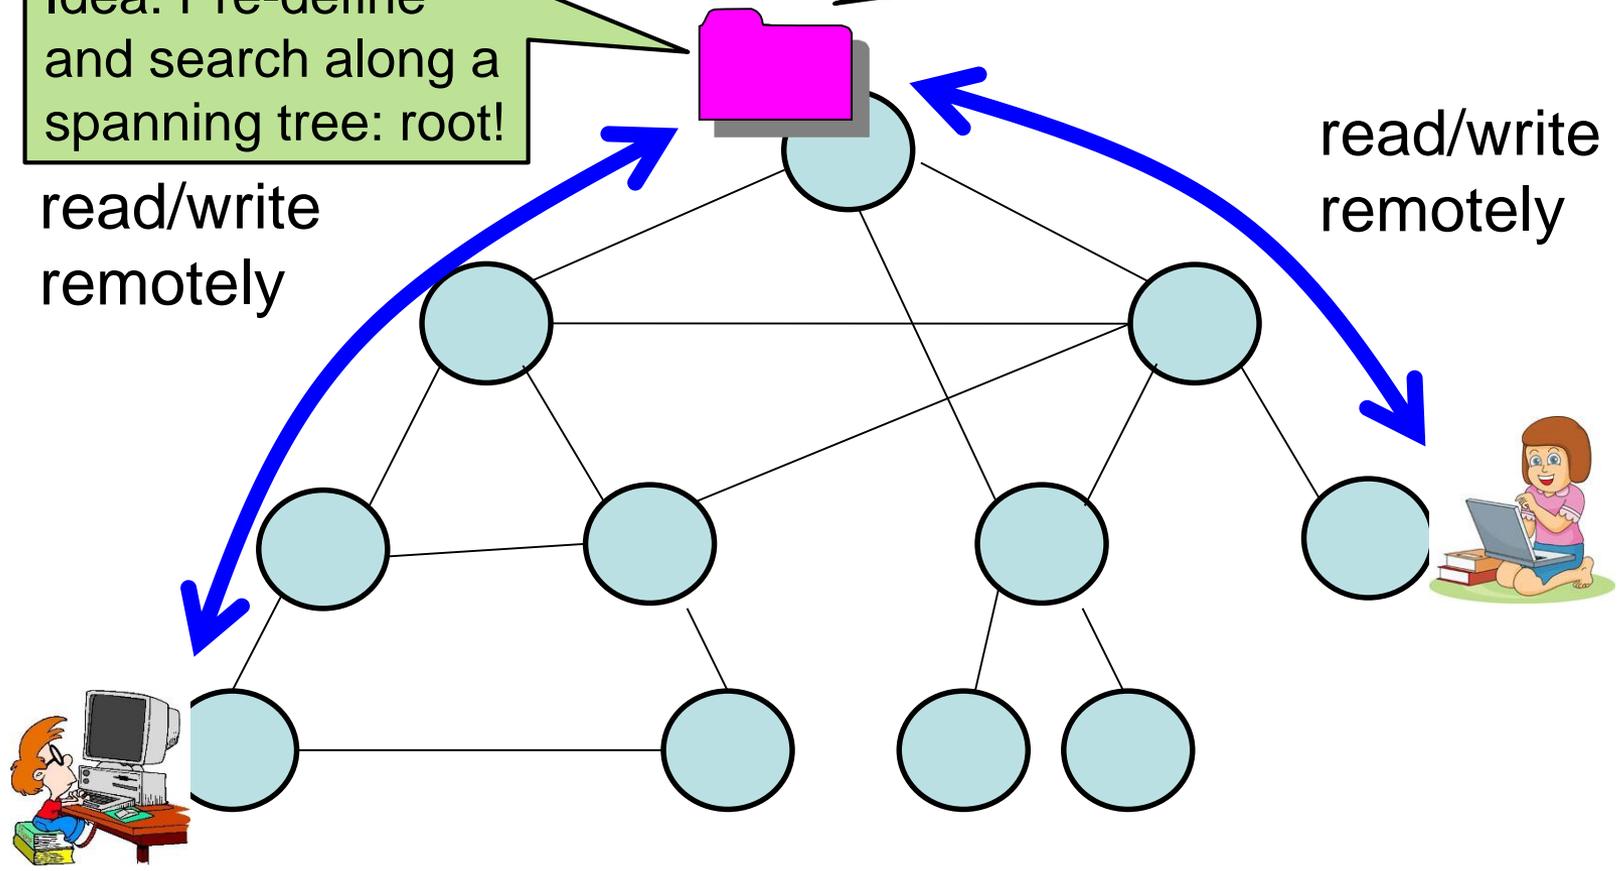
Problem 1: How to find leader?

Idea: just keep a single copy at some *leader node*!

Idea: Pre-define and search along a spanning tree: root!

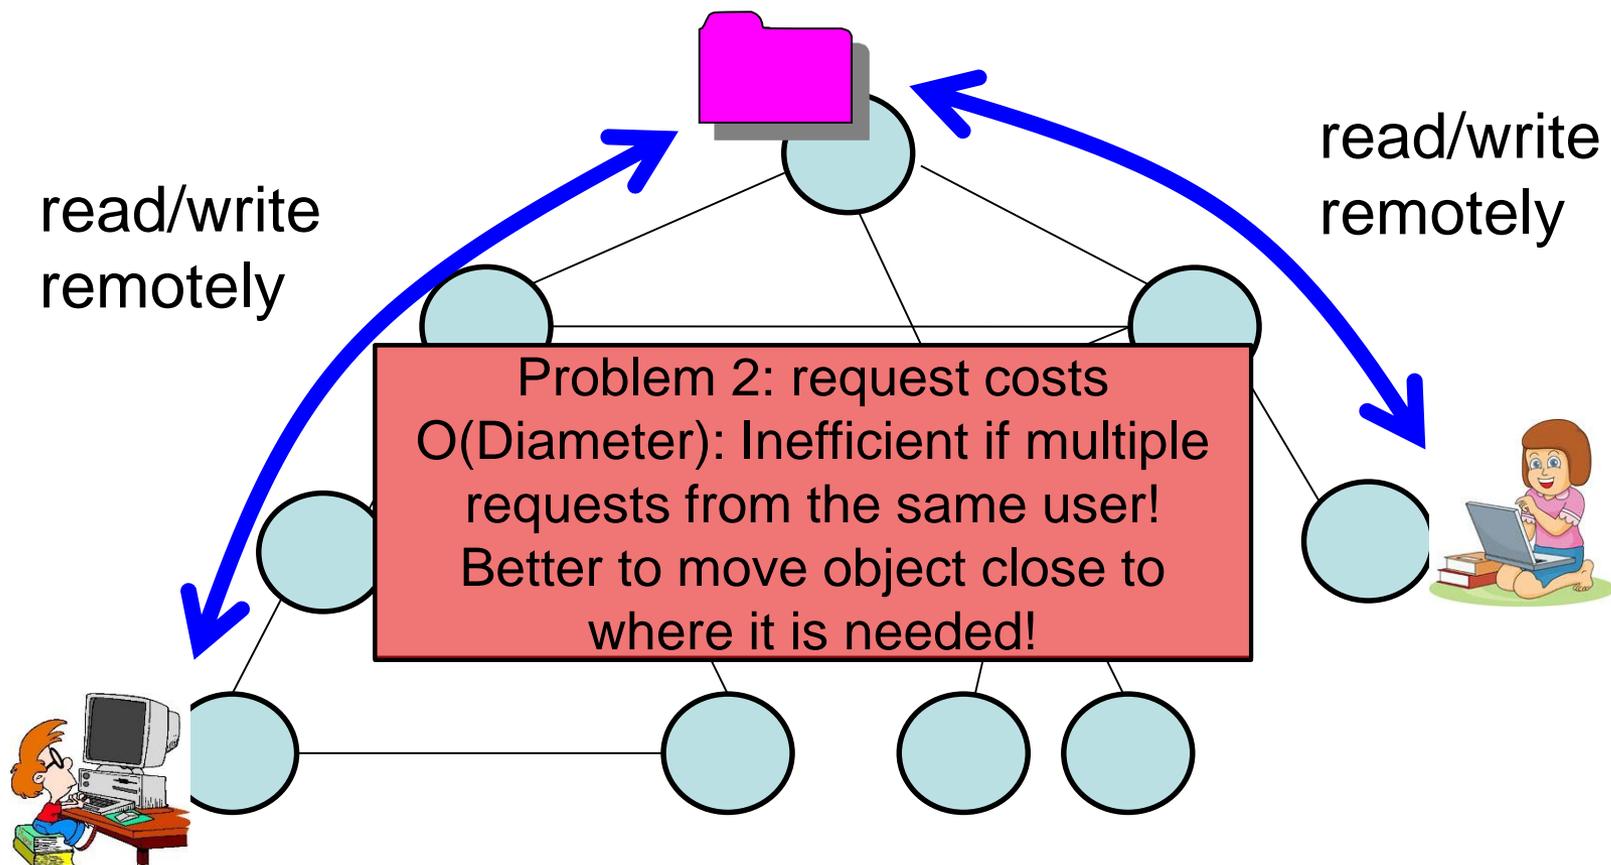
read/write remotely

read/write remotely



Access: 1. send message to leader,
2. leader processes request, 3. result sent back down the tree.

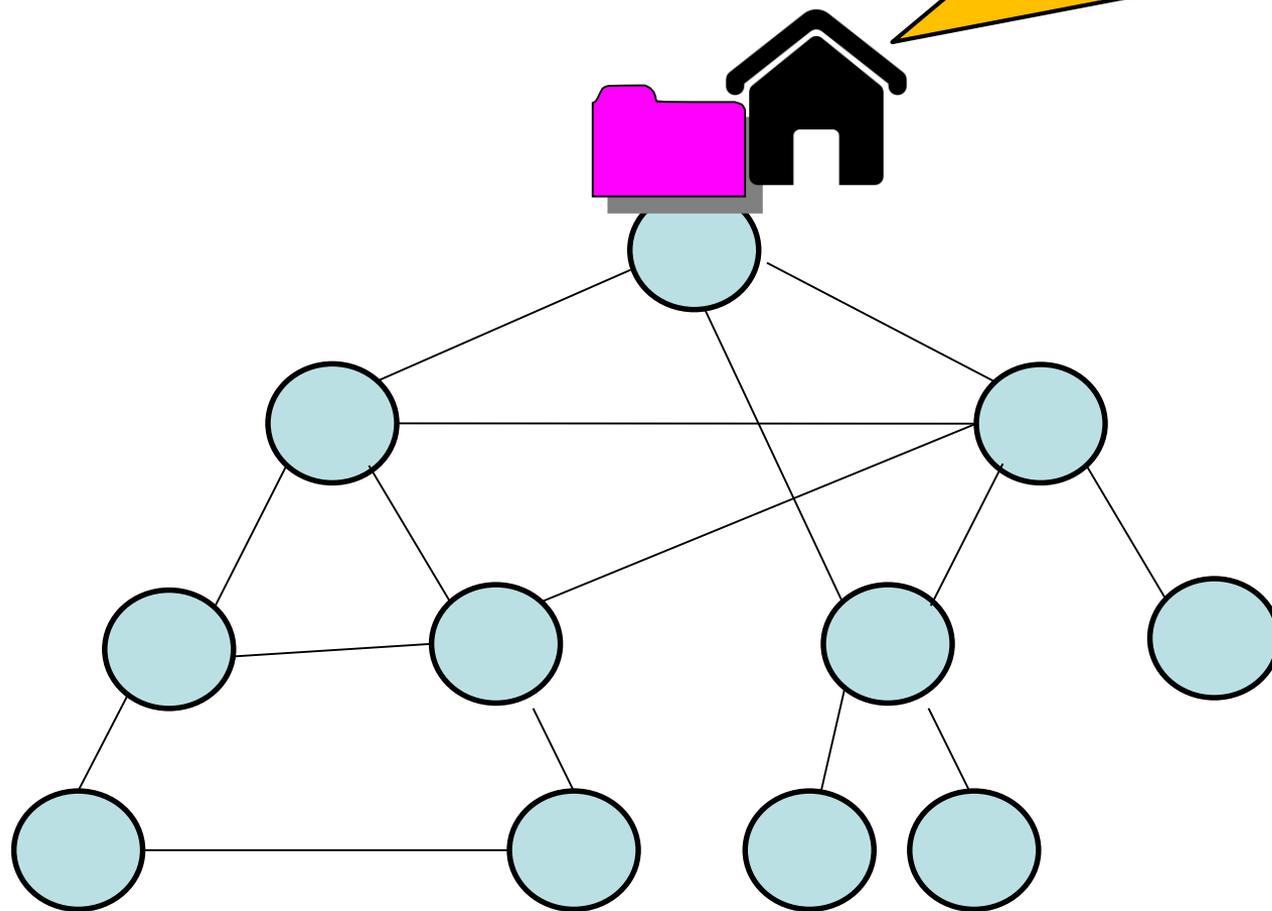
Naive Solution:



- Access:** 1. send message to leader,
2. leader processes request, 3. result sent back down the tree.

Better Solution:

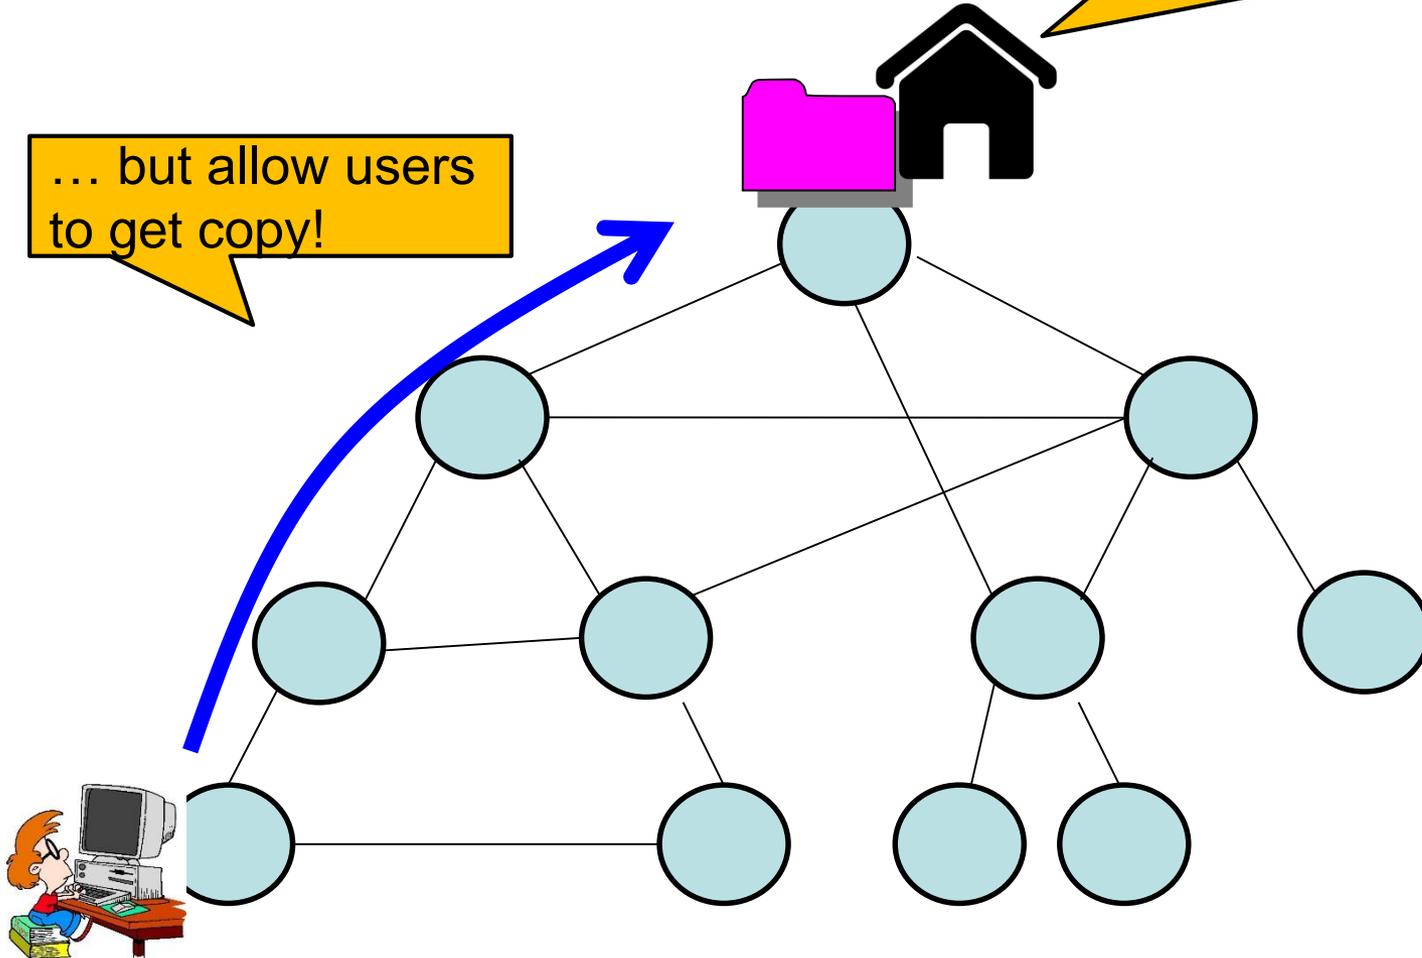
Idea: define home location of object:
e.g., at root of a spanning tree.



Better Solution:

Idea: define home location of object:
e.g., at root of a spanning tree...

... but allow users
to get copy!



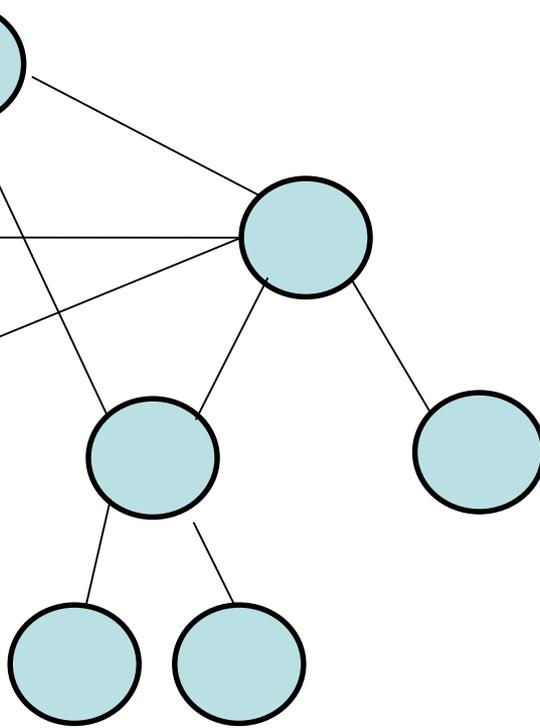
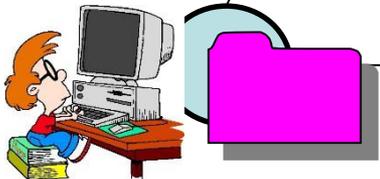
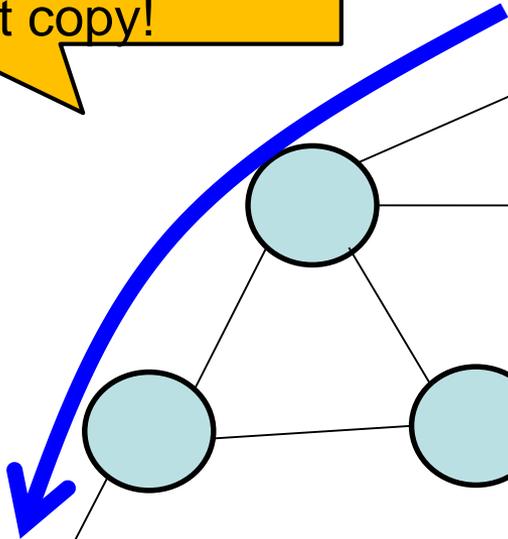
Assume: Access by Bob

Better Solution:

Idea: define home location of object:
e.g., at root of a spanning tree...



... but allow users
to get copy!



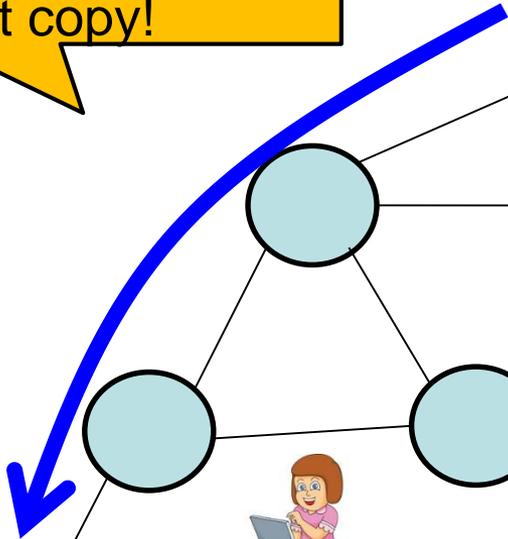
Assume: Access by Bob

Better Solution:

Idea: define home location of object:
e.g., at root of a spanning tree...



... but allow users
to get copy!



What if Alice wants access?

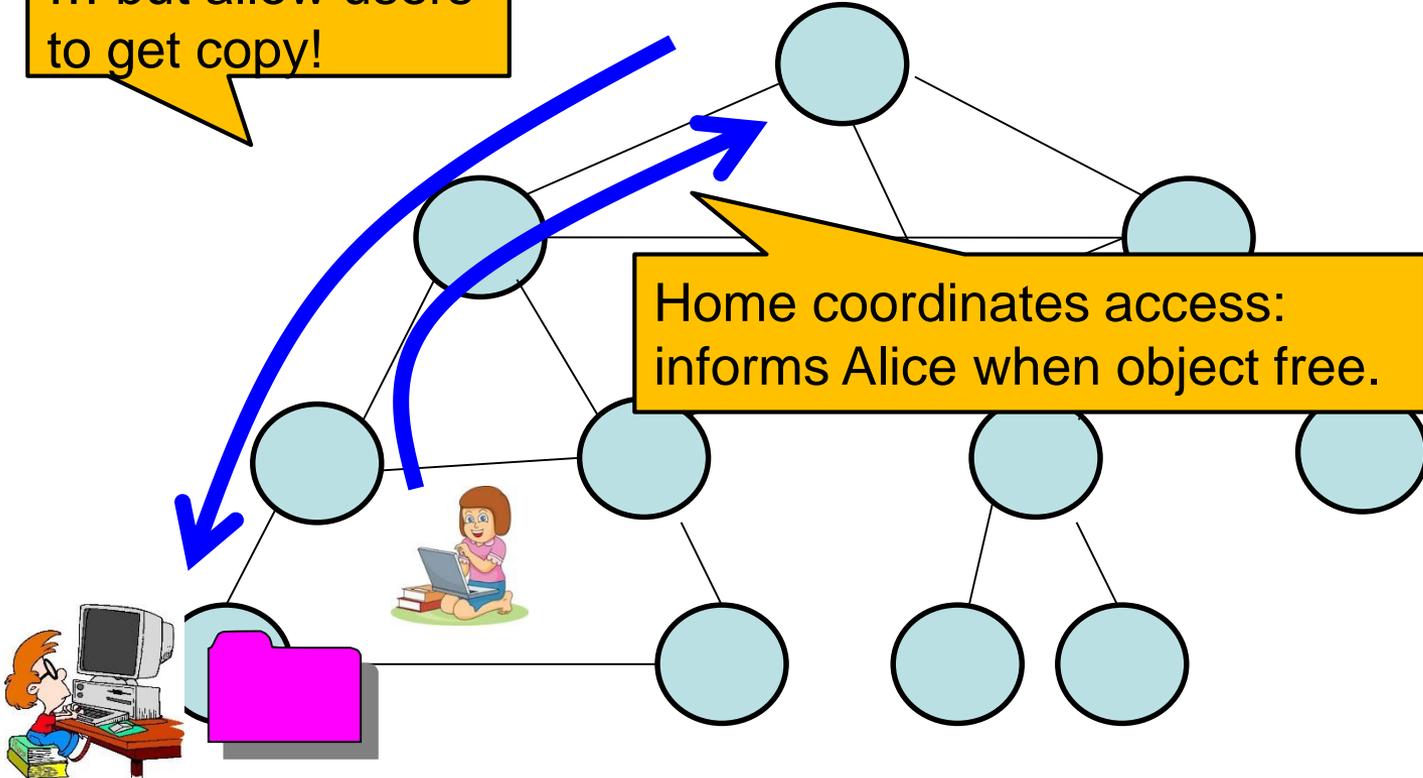
Better Solution:

Idea: define home location of object:
e.g., at root of a spanning tree...



... but allow users
to get copy!

Home coordinates access:
informs Alice when object free.



What if Alice wants access?

Better Solution:

Idea: define home location of object:
e.g., at root of a spanning tree...

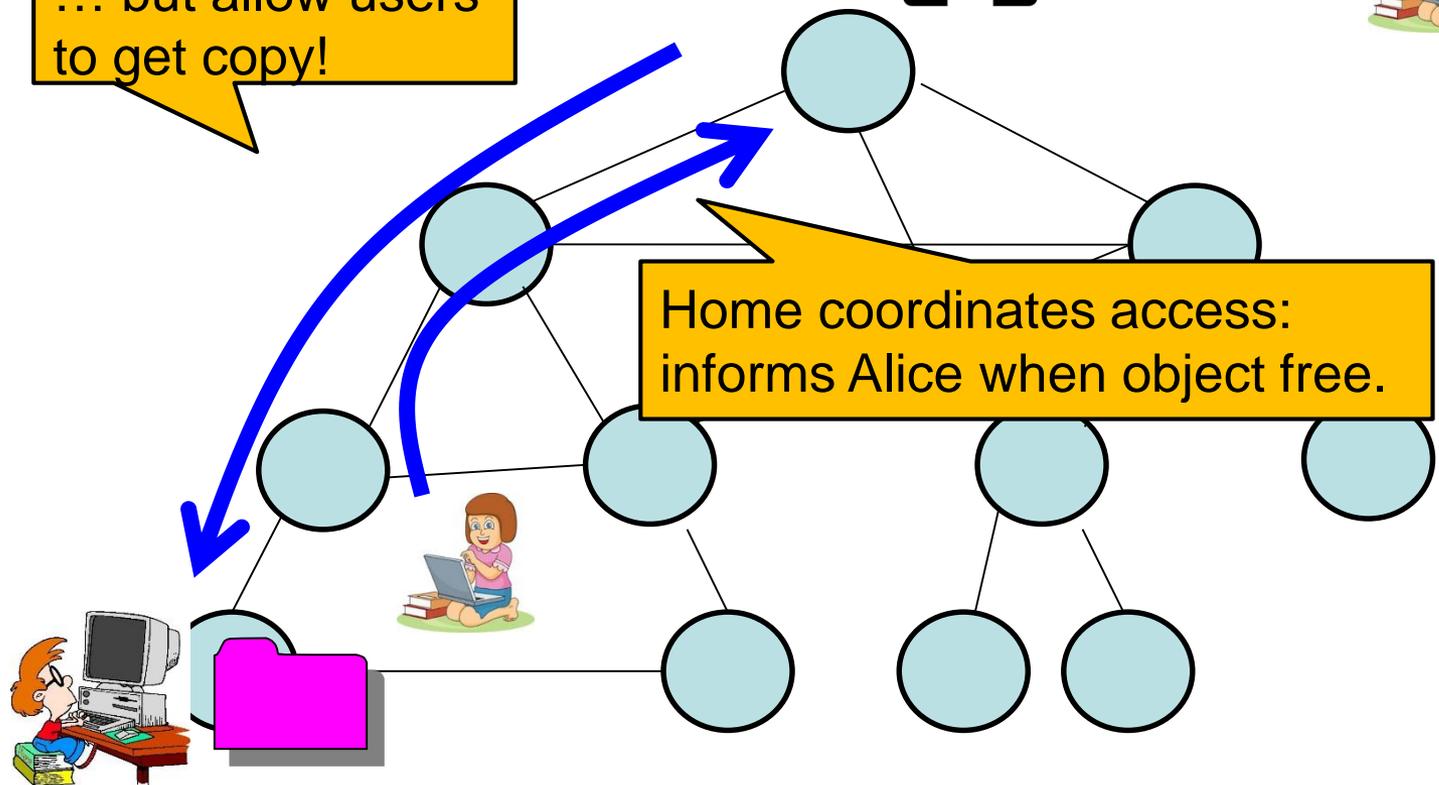


next in line:



... but allow users
to get copy!

Home coordinates access:
informs Alice when object free.

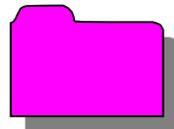


What if Alice wants access?

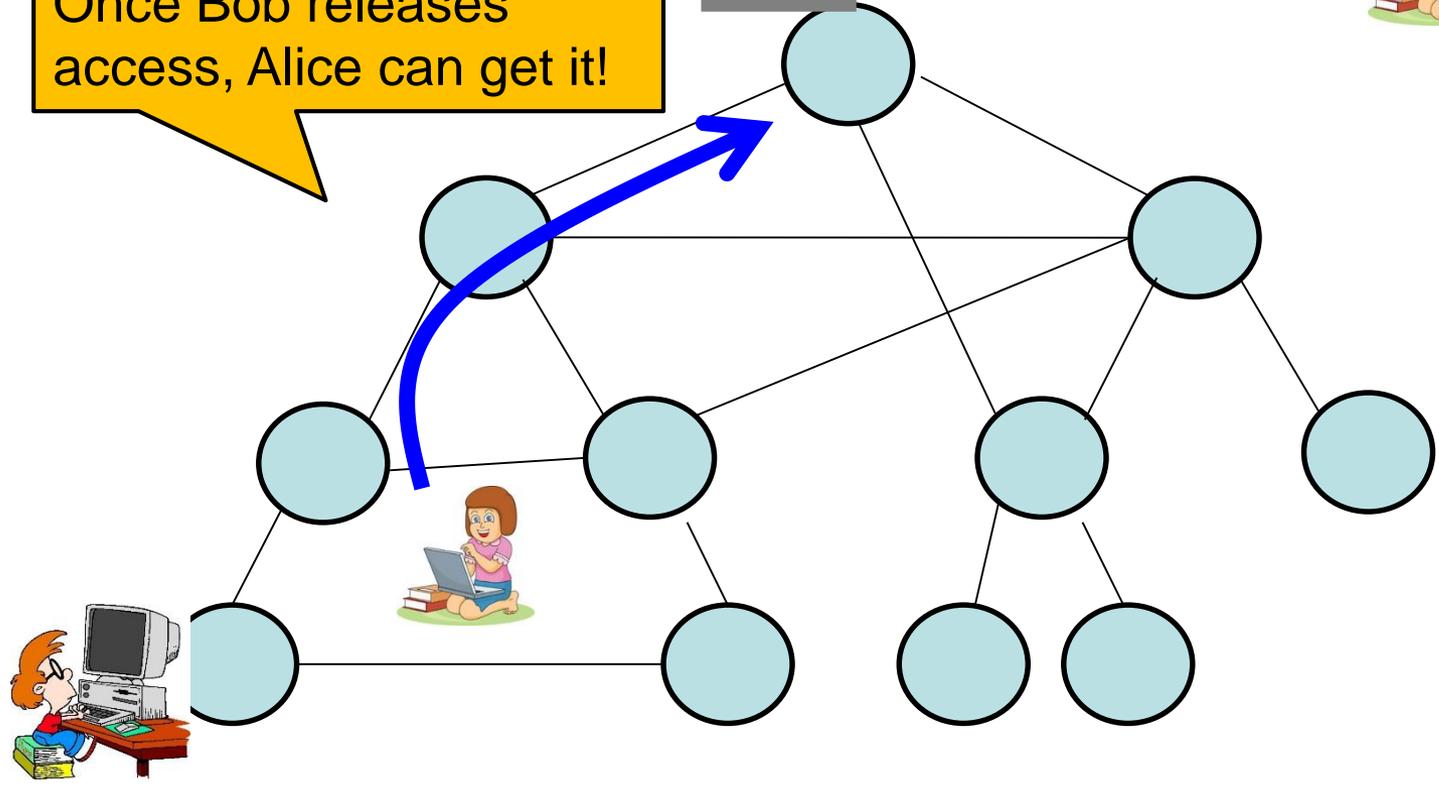
Better Solution:

Idea: define home location of object:
e.g., at root of a spanning tree.

Once Bob releases access, Alice can get it!



next in line:



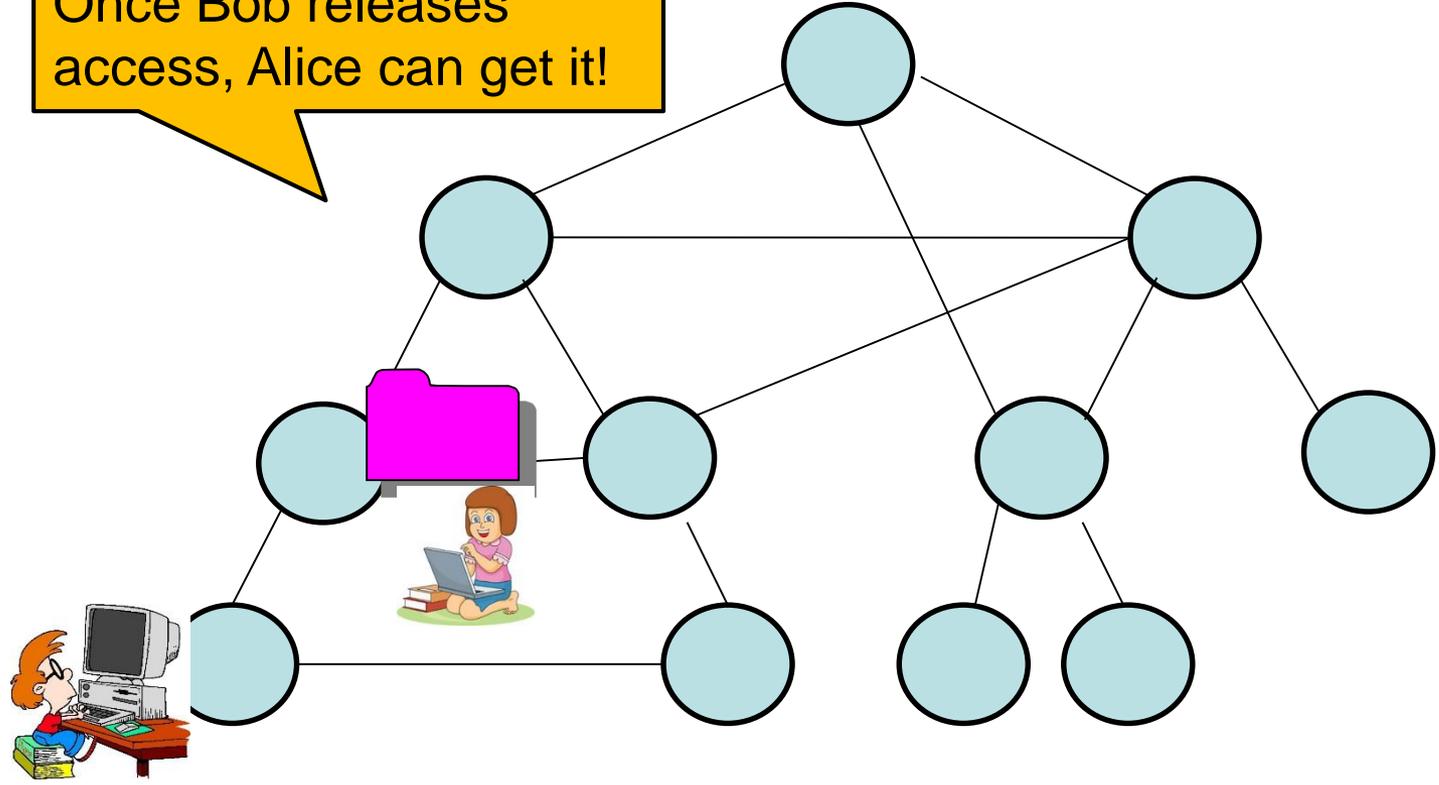
What if Alice wants access?

Better Solution:

Idea: define home location of object: e.g., at root of a spanning tree.



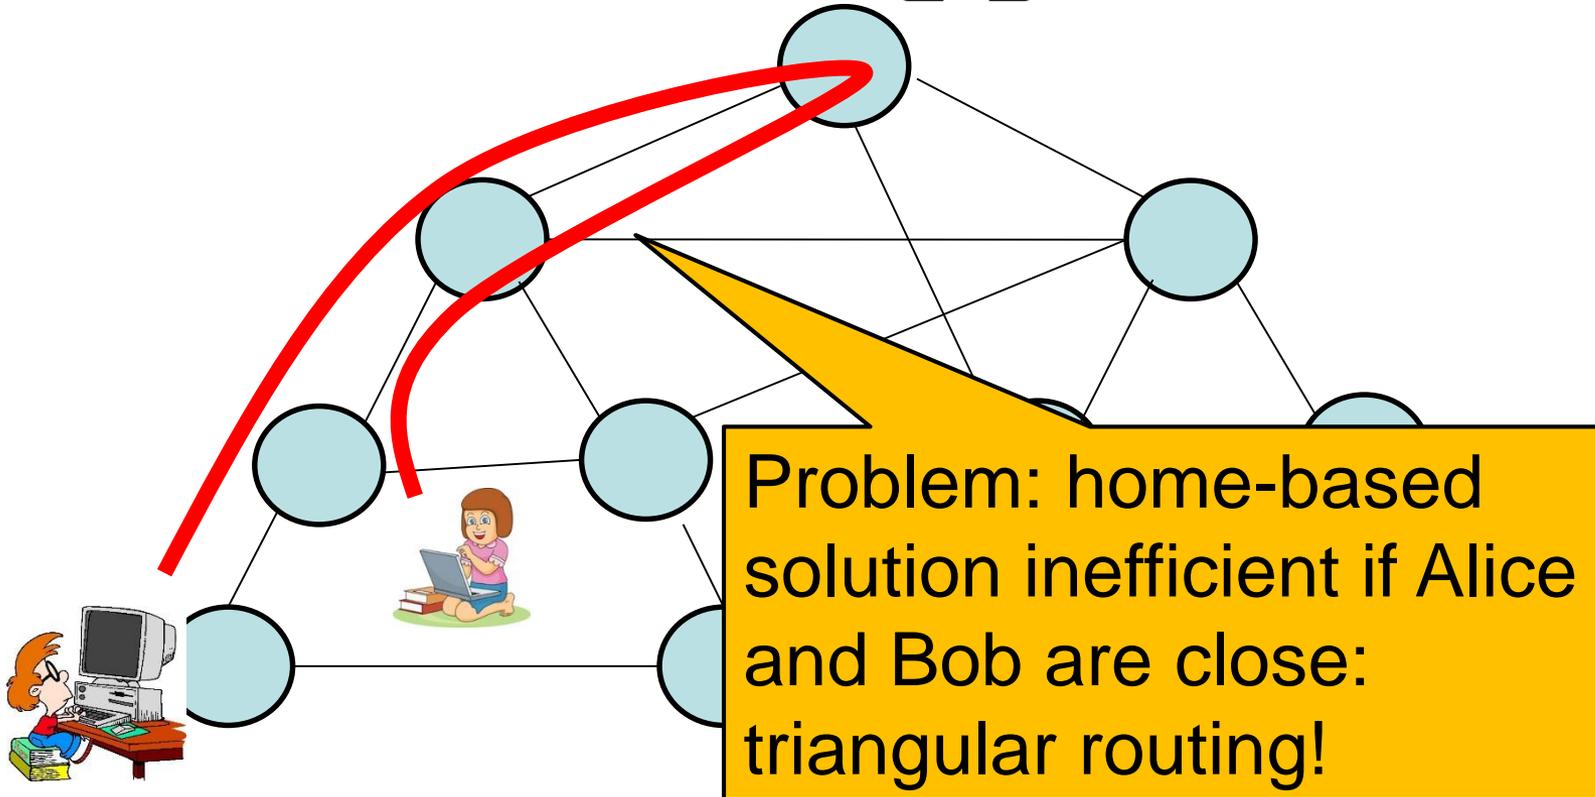
Once Bob releases access, Alice can get it!



What if Alice wants access?

Everything great?

Idea: define home location of object:
e.g., at root of a spanning tree.

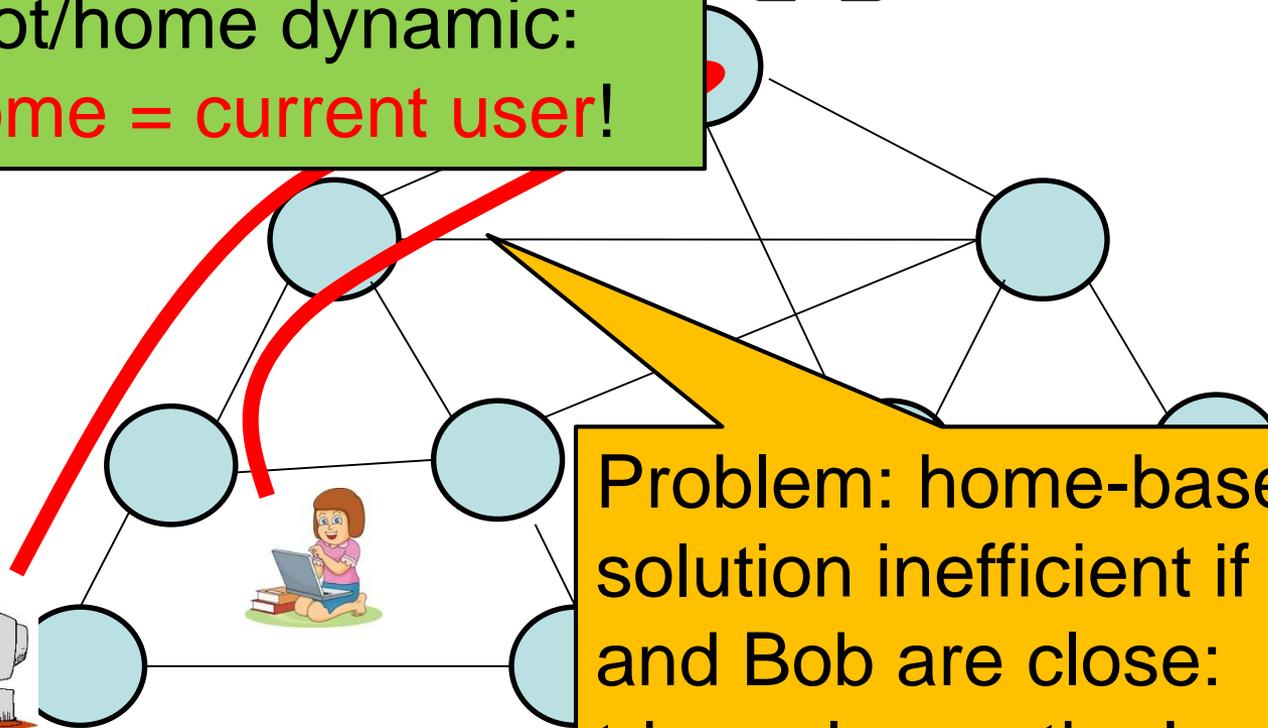


Everything great?

Idea: define home location of object:
e.g., at root of a spanning tree.



Idea: make the
root/home dynamic:
home = current user!



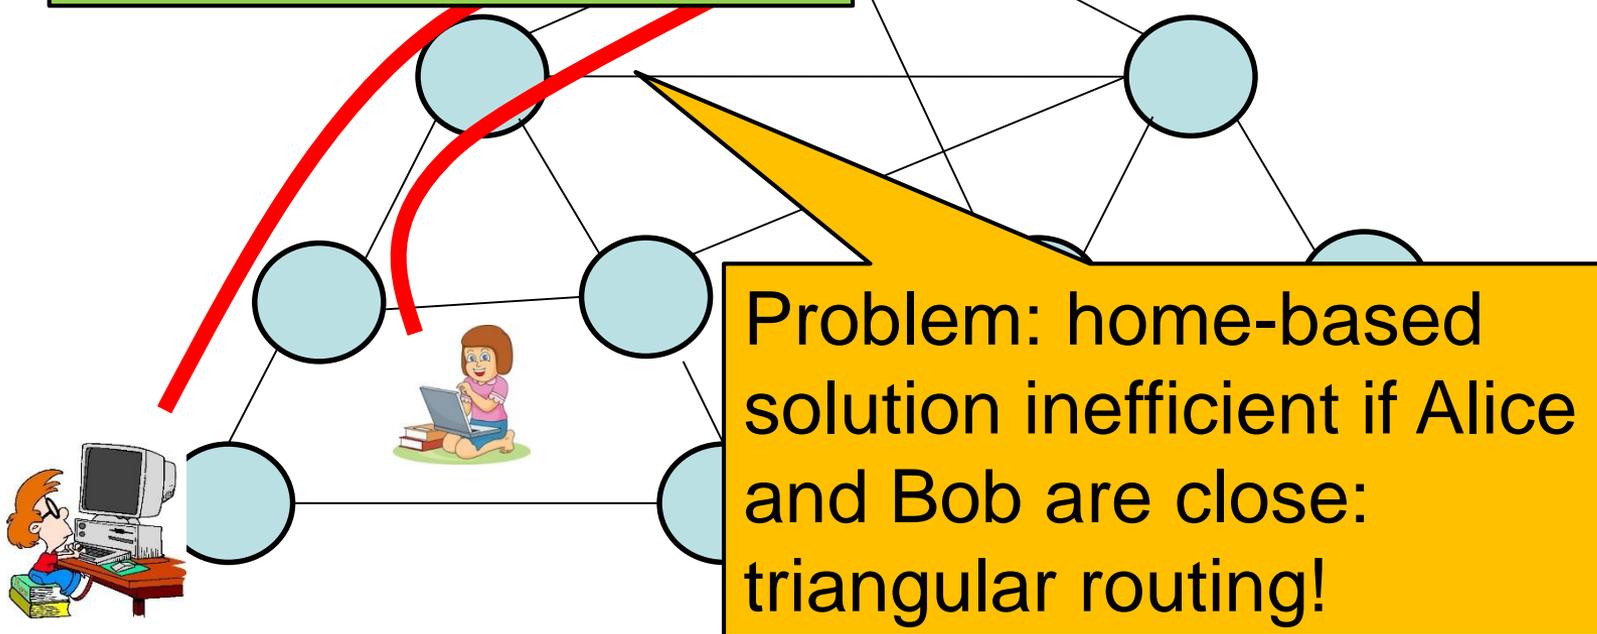
Problem: home-based
solution inefficient if Alice
and Bob are close:
triangular routing!

Everything great?

Idea: define home location of object:
e.g., at root of a spanning tree.



Idea: make the
root/home dynamic:
home = current user!



Problem: home-based
solution inefficient if Alice
and Bob are close:
triangular routing!

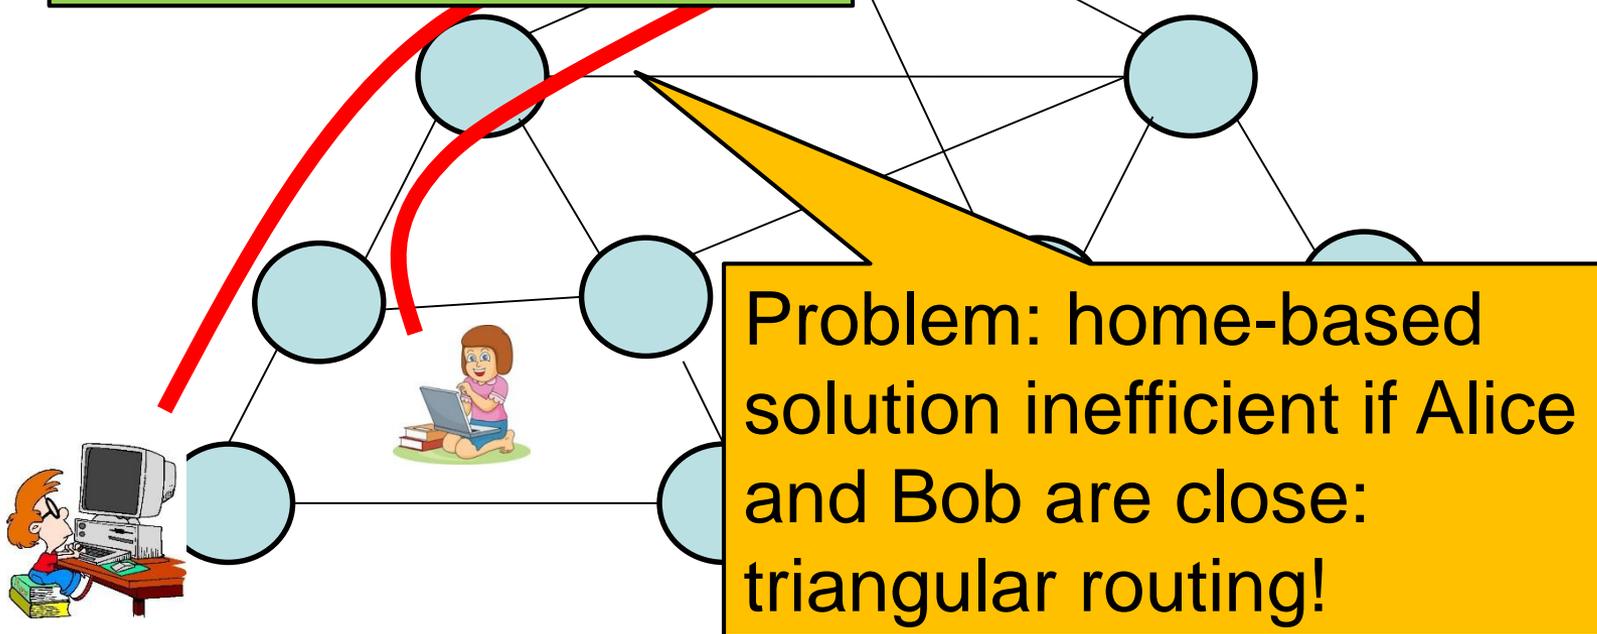
The Arrow Protocol!

Everything great?

Idea: define home location of object:
e.g., at root of a spanning tree.



Idea: make the
root/home dynamic:
home = current user!



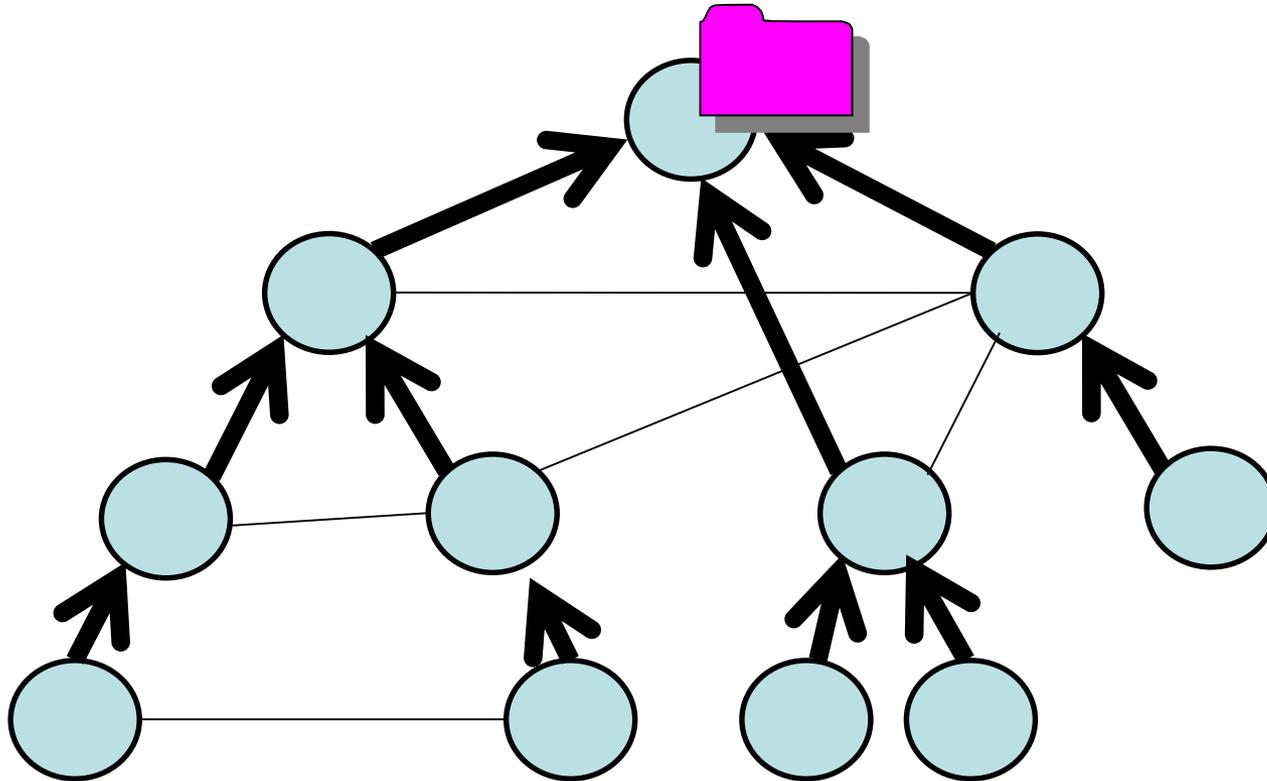
Problem: home-based
solution inefficient if Alice
and Bob are close:
triangular routing!

The Arrow Protocol!

Based on link reversals:
remember MST!

The Arrow Protocol

Idea maintain invariant: directed spanning tree pointing to the object.



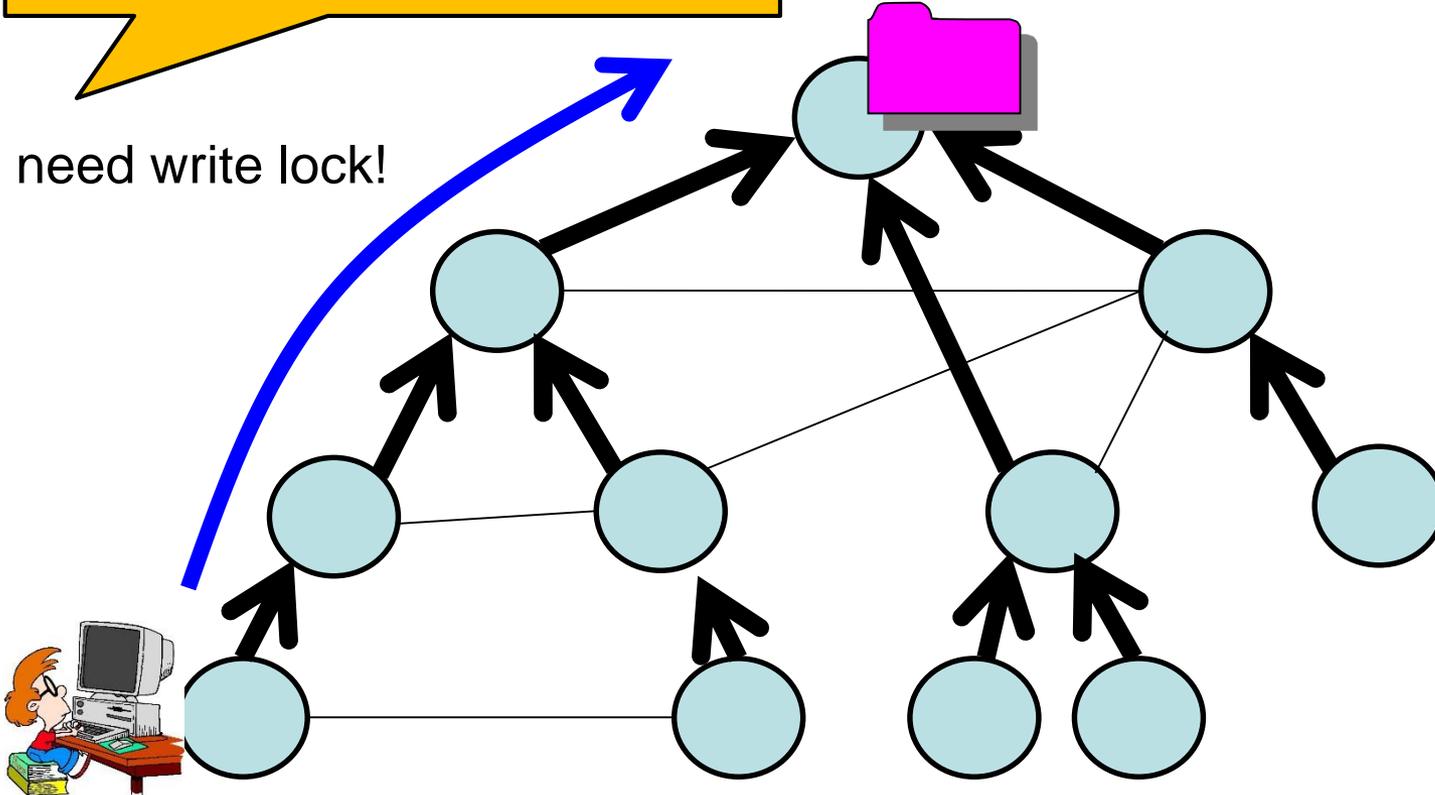
In the Beginning

The Arrow Protocol

Find object is easy due to invariant: simply route along directed spanning tree edges!

Idea maintain invariant: directed spanning tree pointing to the object.

I need write lock!



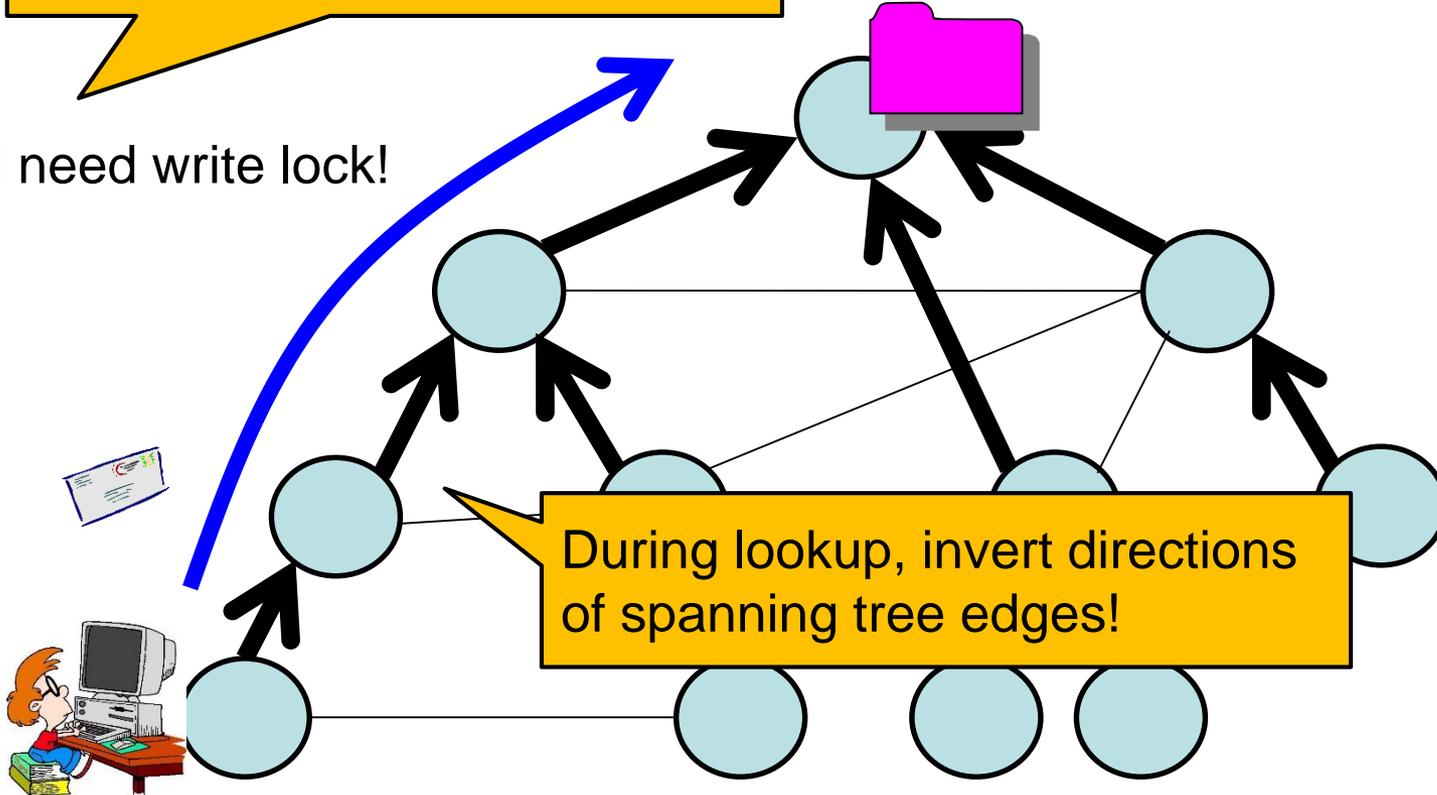
At Runtime

The Arrow Protocol

Find object is easy due to invariant: simply route along directed spanning tree edges!

Idea maintain invariant: directed spanning tree pointing to the object.

I need write lock!



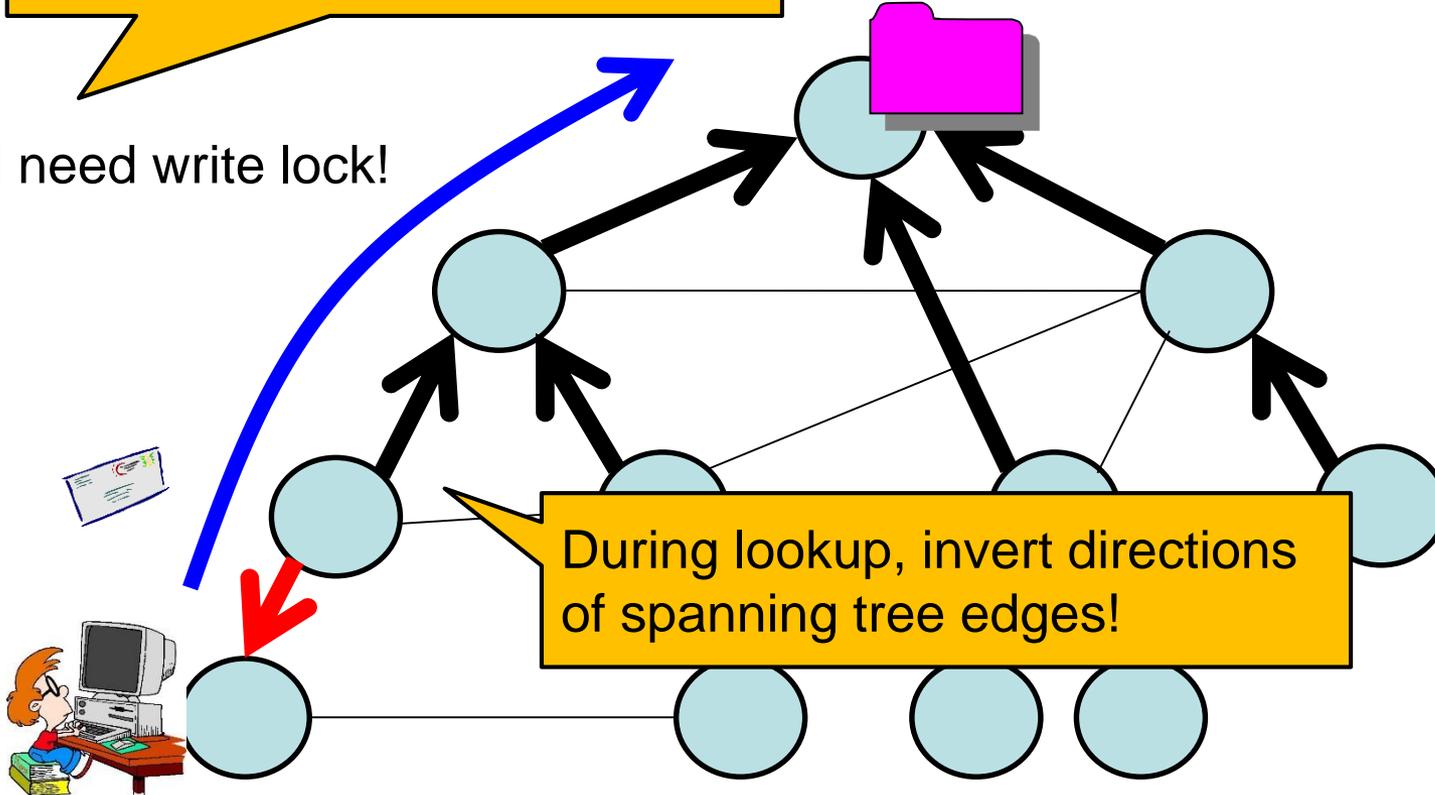
At Runtime

The Arrow Protocol

Find object is easy due to invariant: simply route along directed spanning tree edges!

Idea maintain invariant: directed spanning tree pointing to the object.

I need write lock!



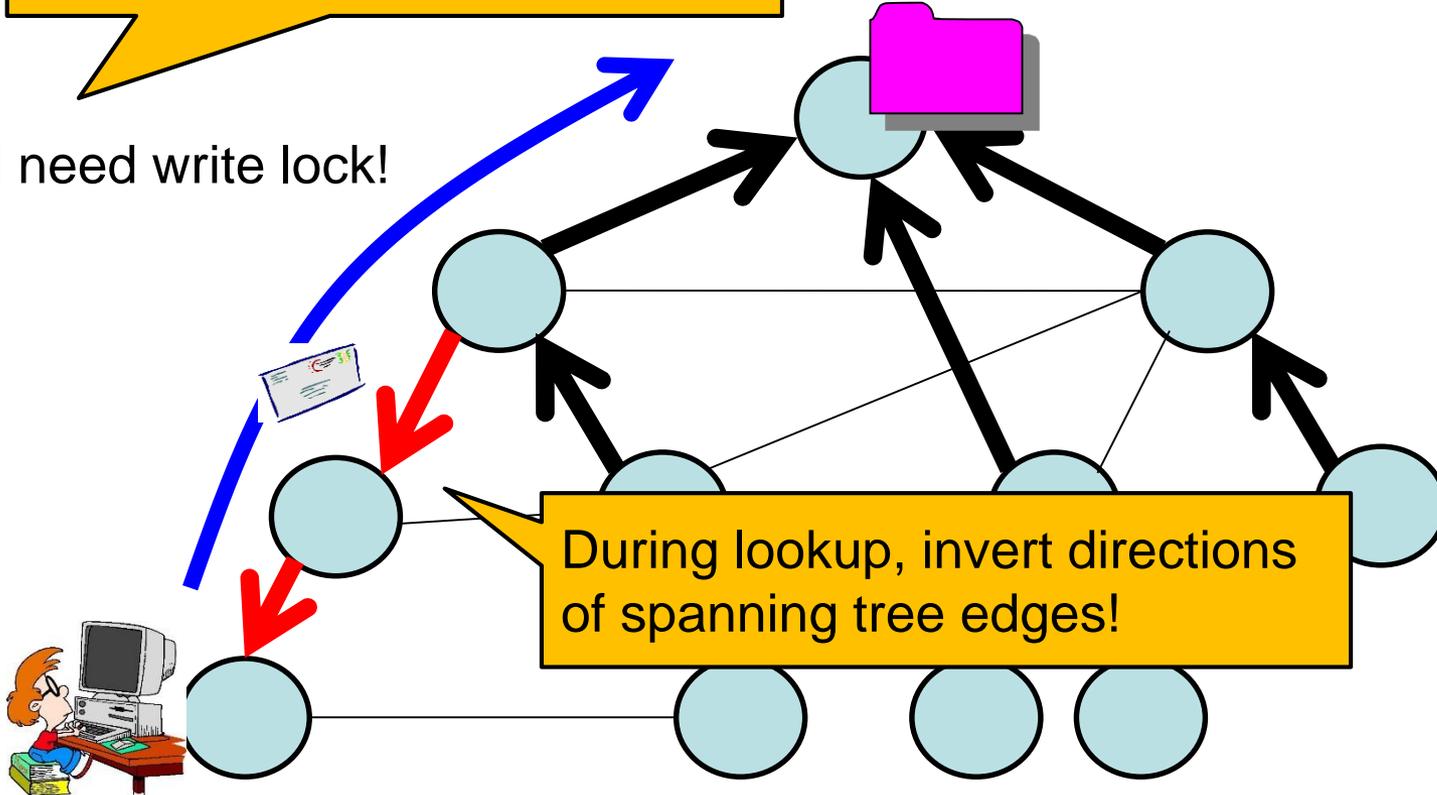
At Runtime

The Arrow Protocol

Find object is easy due to invariant: simply route along directed spanning tree edges!

Idea maintain invariant: directed spanning tree pointing to the object.

I need write lock!



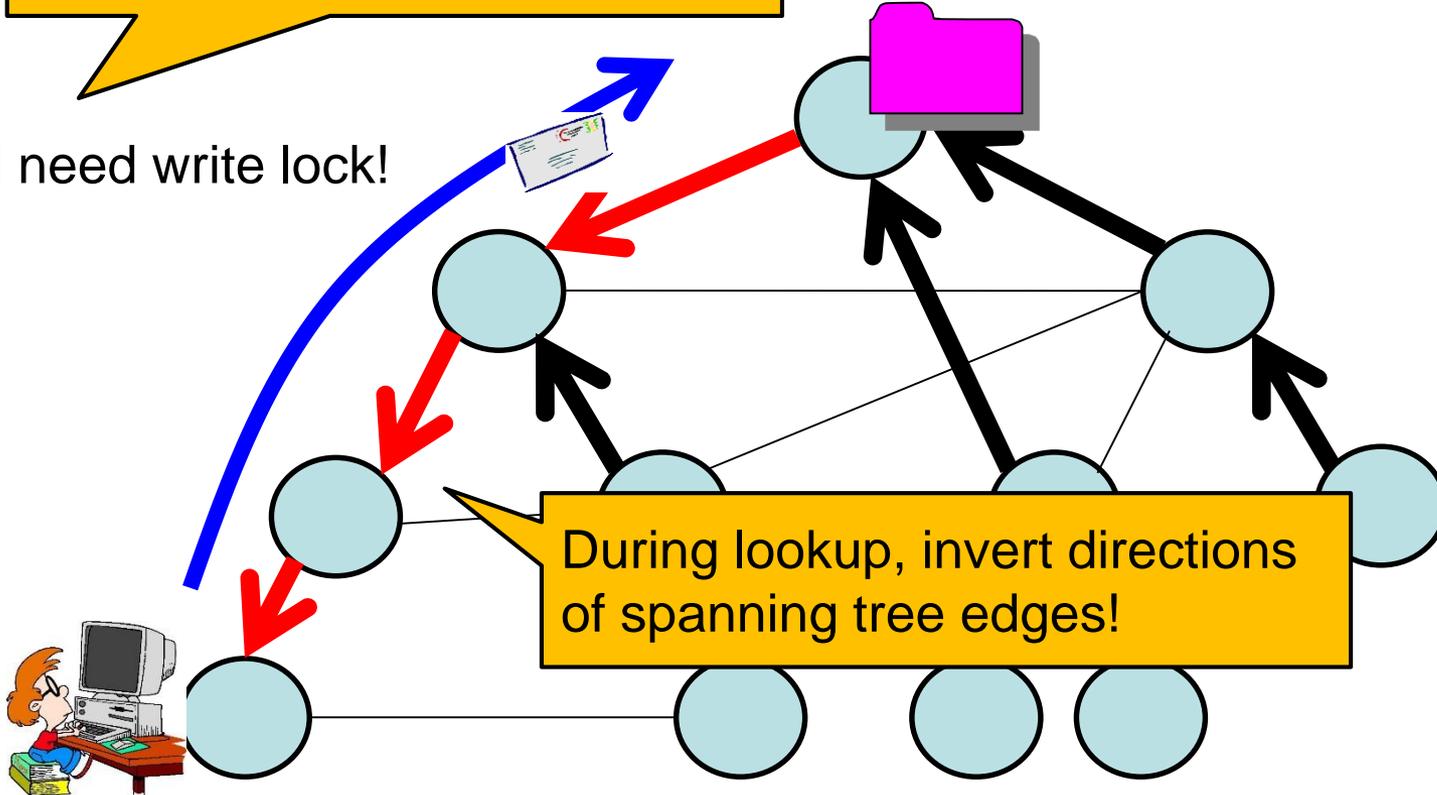
At Runtime

The Arrow Protocol

Find object is easy due to invariant: simply route along directed spanning tree edges!

Idea maintain invariant: directed spanning tree pointing to the object.

I need write lock!



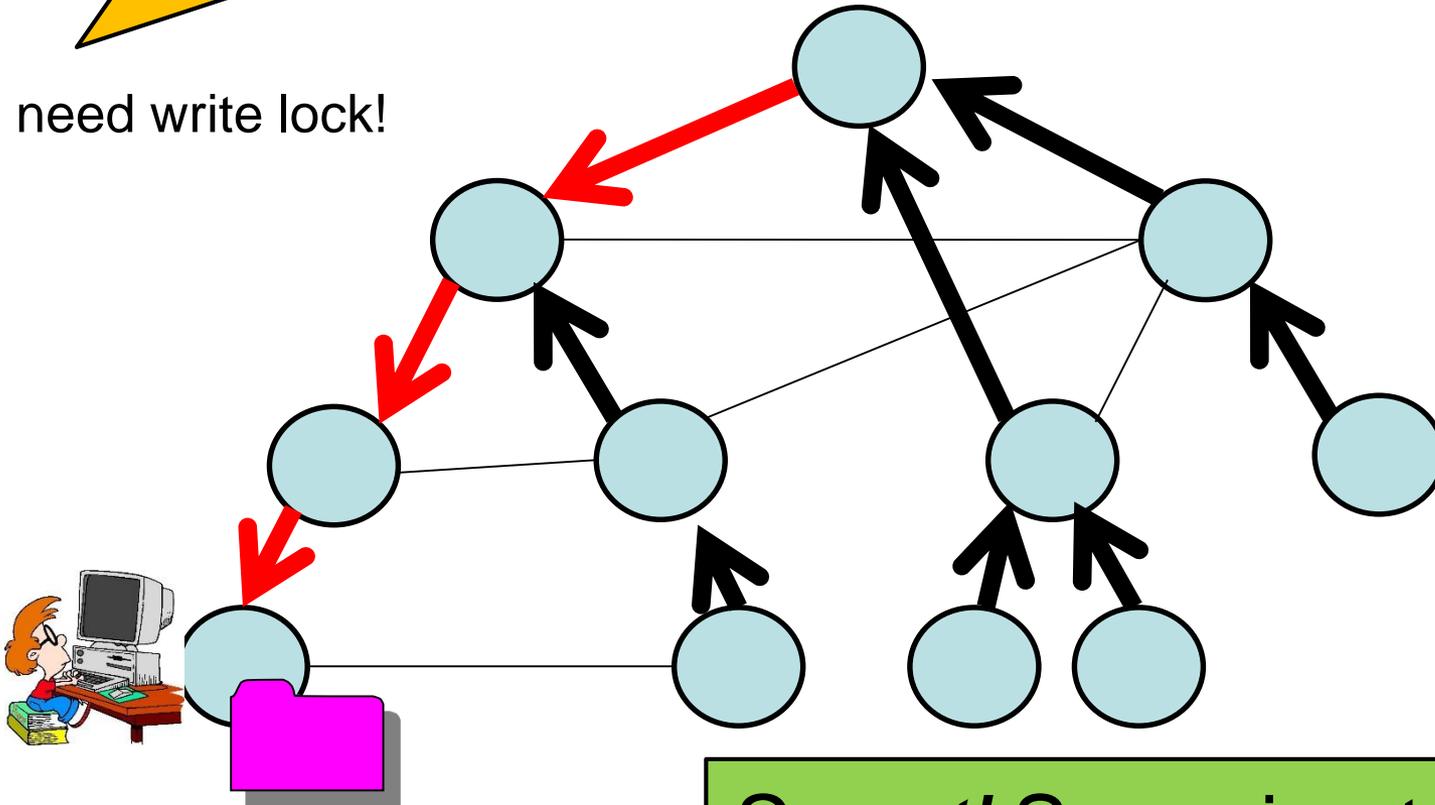
At Runtime

The Arrow Protocol

Find object is easy due to invariant: simply route along directed spanning tree edges!

Idea maintain invariant: directed spanning tree pointing to the object.

I need write lock!



At Runtime

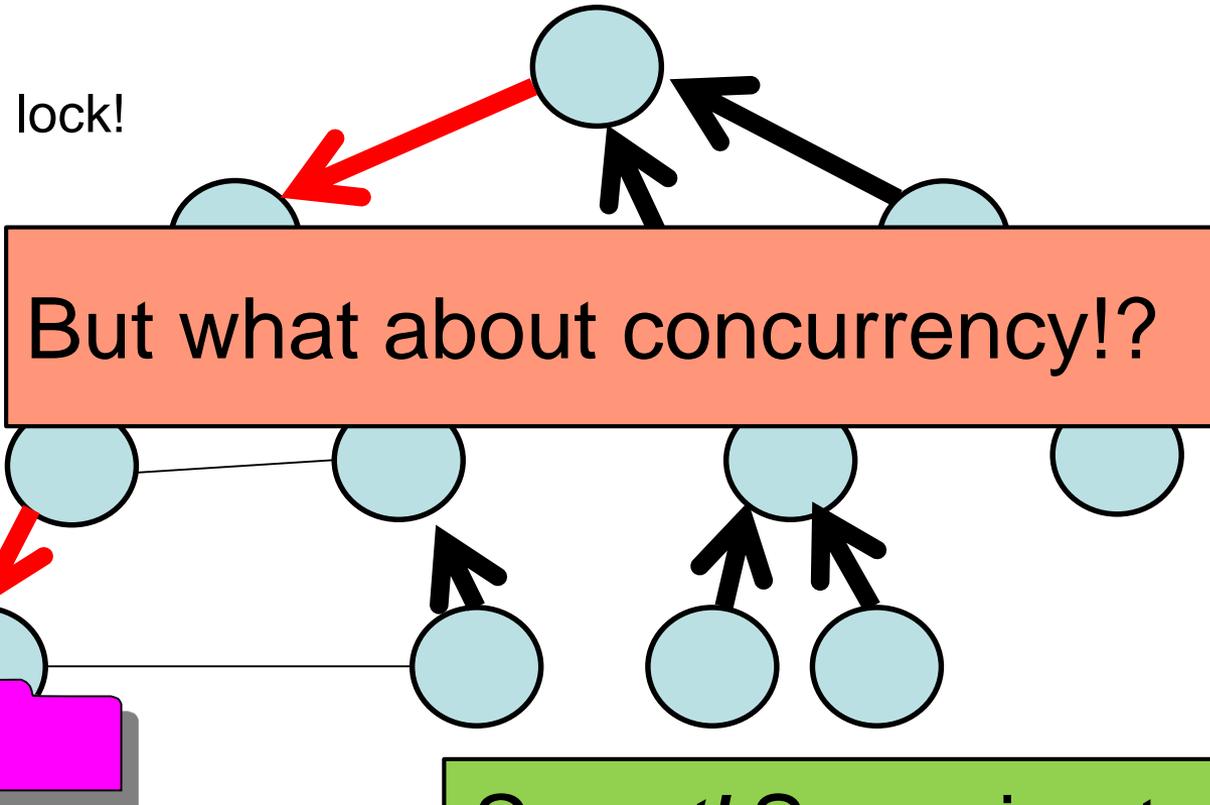
Sweet! Spanning tree rooted at object location again. 😊

The Arrow Protocol

Find object is easy due to invariant: simply route along directed spanning tree edges!

Idea maintain invariant: directed spanning tree pointing to the object.

I need write lock!



At Runtime

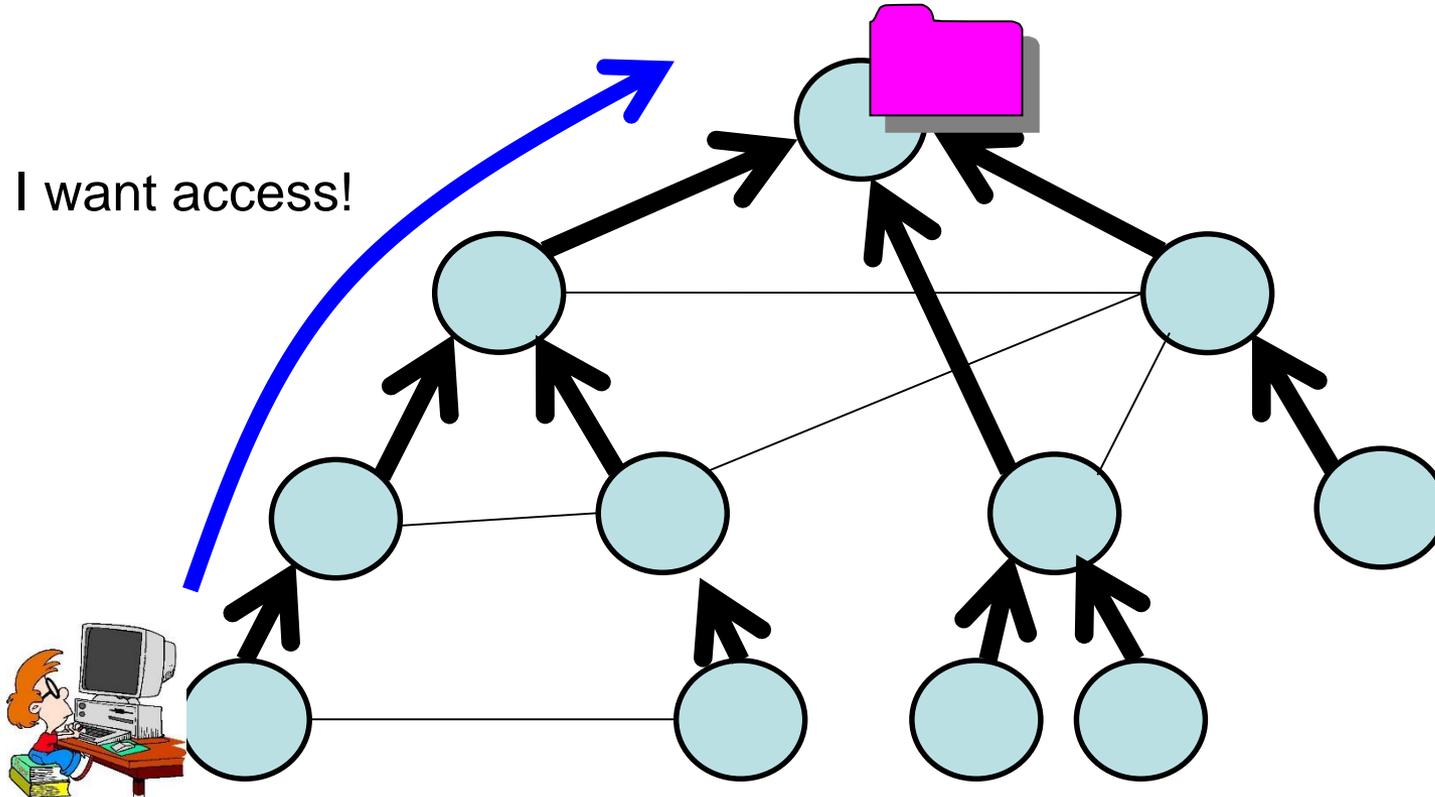
Sweet! Spanning tree rooted at object location again. 😊

The Arrow Protocol: Concurrency

Still in use by
someone else!

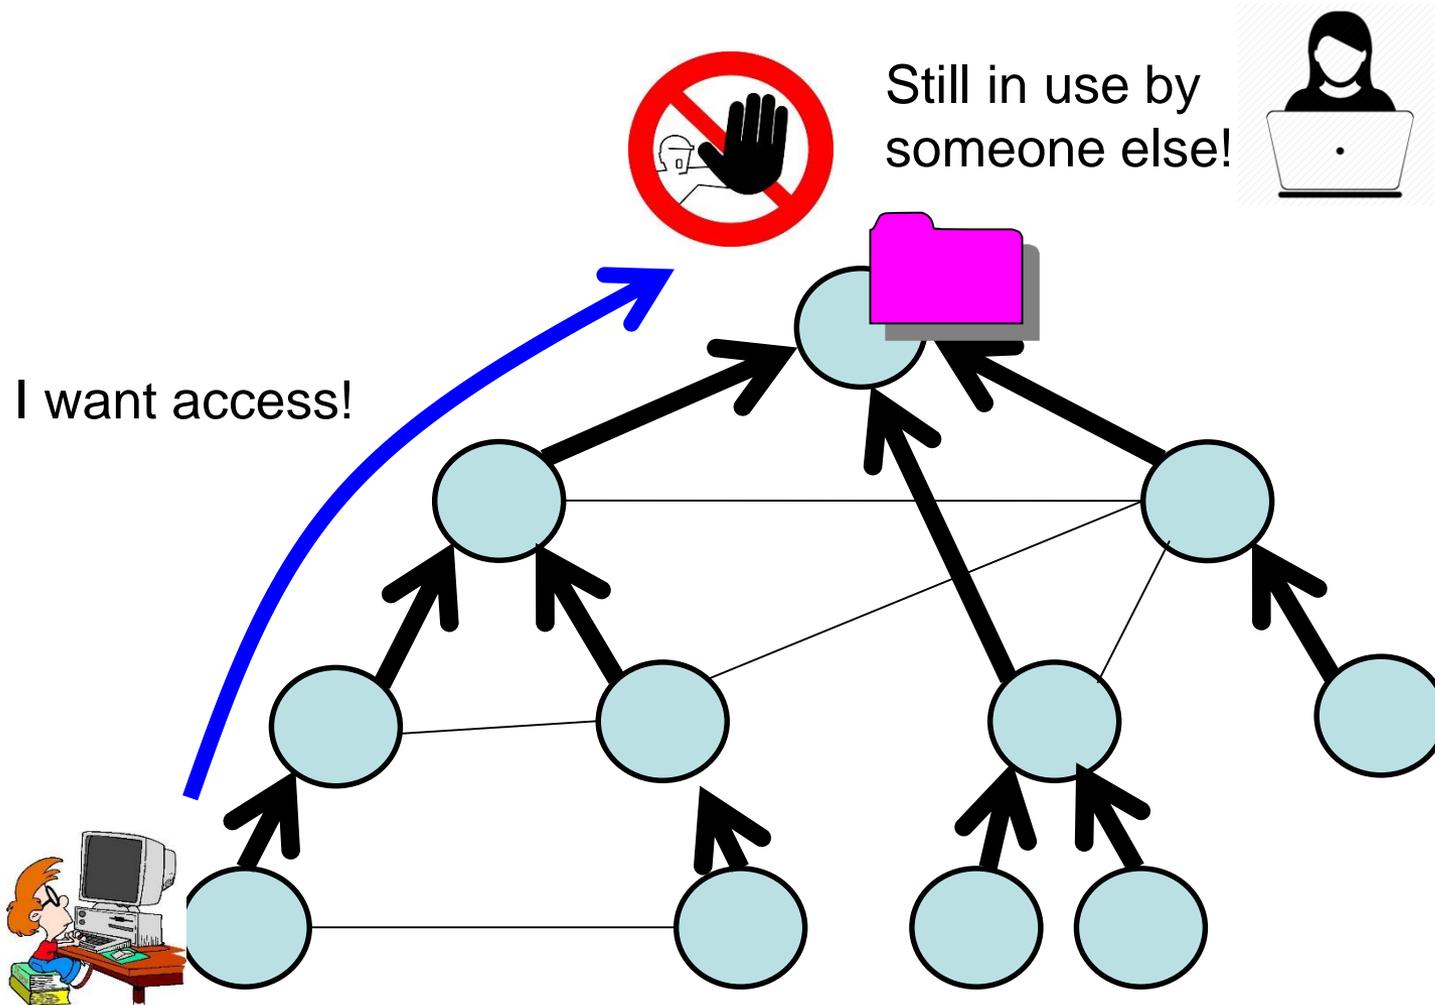


I want access!



In the Beginning

The Arrow Protocol: Concurrency

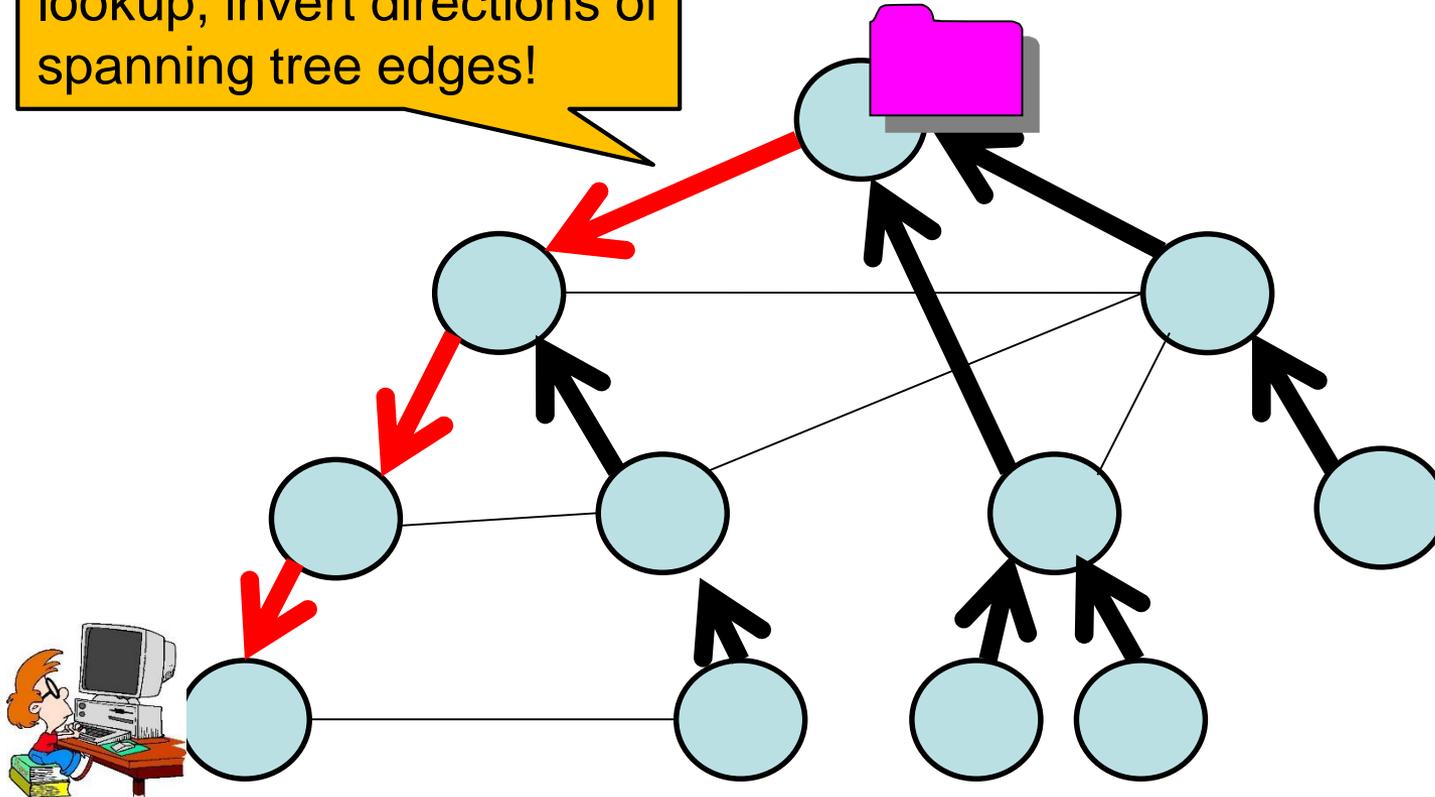


In the Beginning

The Arrow Protocol: Concurrency

However, like before:
lookup, invert directions of
spanning tree edges!

Still in use by
someone else!



Phase 1

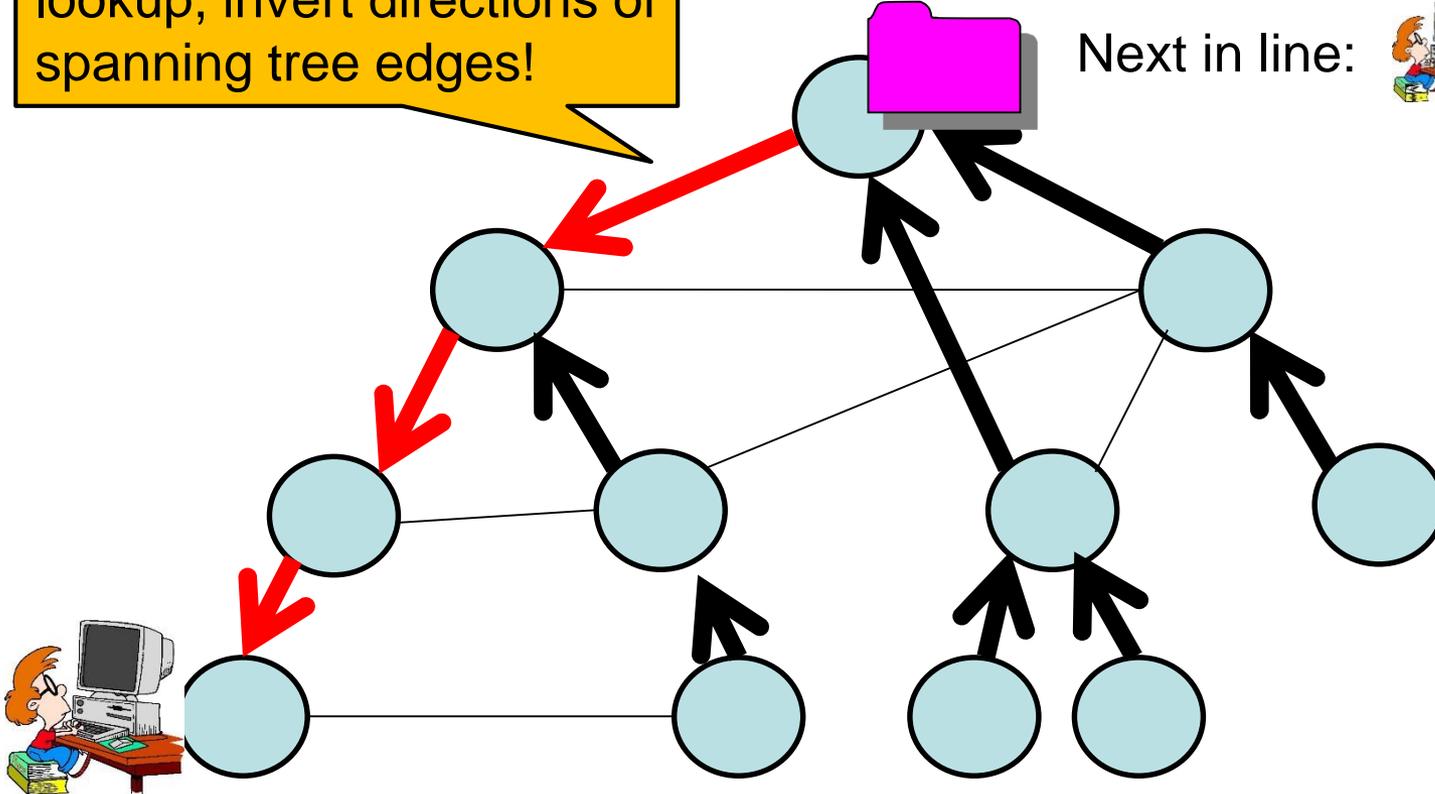
The Arrow Protocol: Concurrency

However, like before:
lookup, invert directions of
spanning tree edges!

Still in use by
someone else!

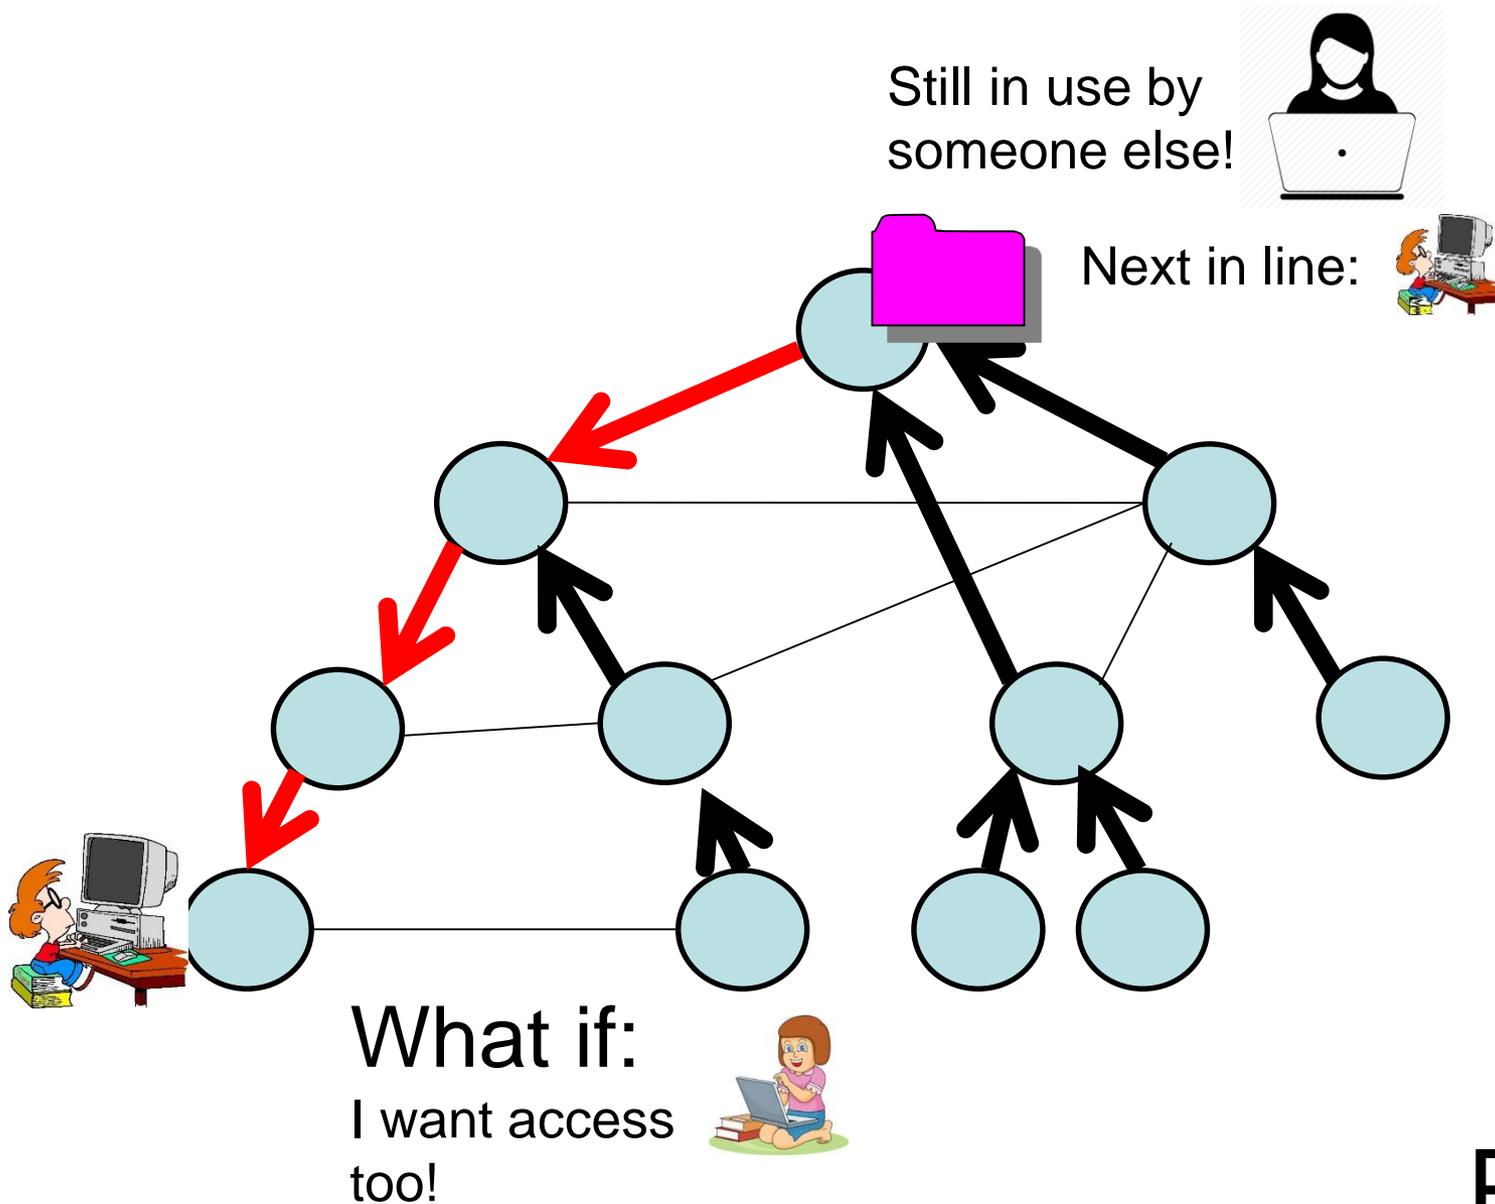


Next in line:



Phase 1

The Arrow Protocol: Concurrency



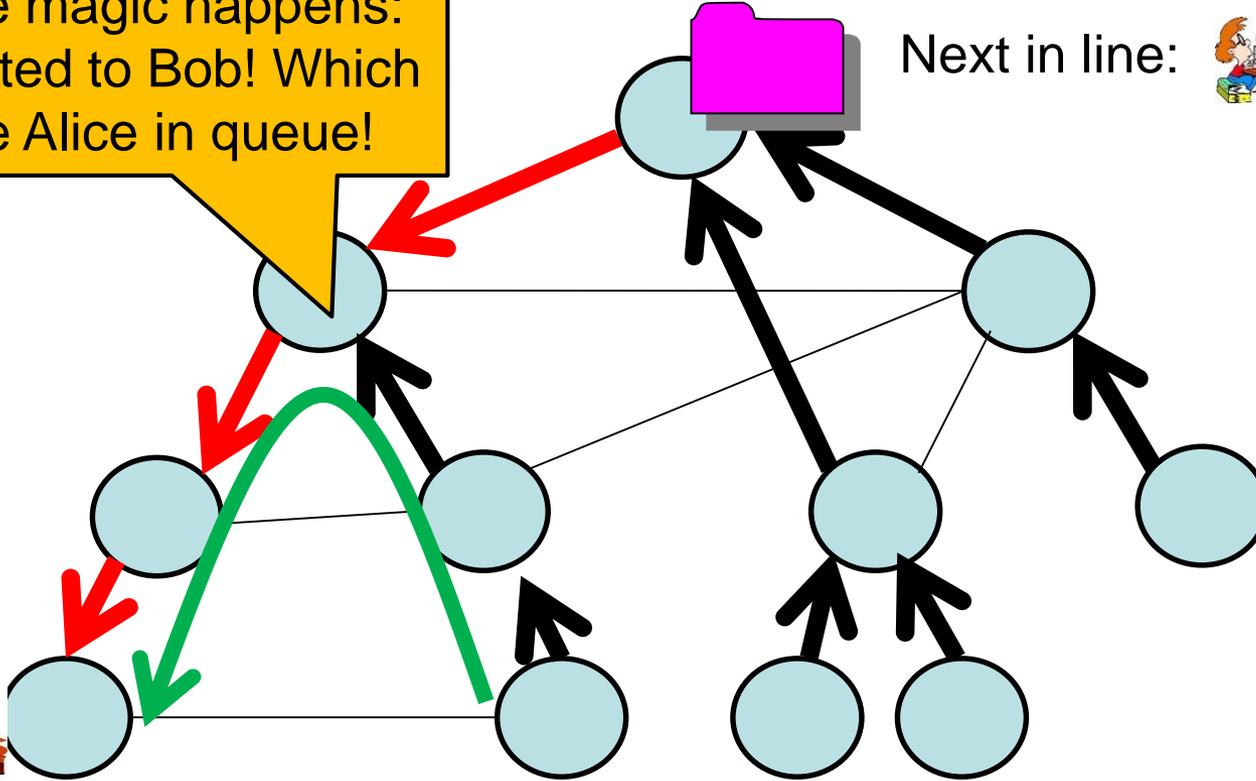
The Arrow Protocol: Concurrency

Here the magic happens:
gets routed to Bob! Which
is before Alice in queue!

Still in use by
someone else!



Next in line:



What if:
I want access
too!



Phase 2

The Arrow Protocol: Concurrency

Here the magic happens:
gets routed to Bob! Which
is before Alice in queue!

Still in use by
someone else!



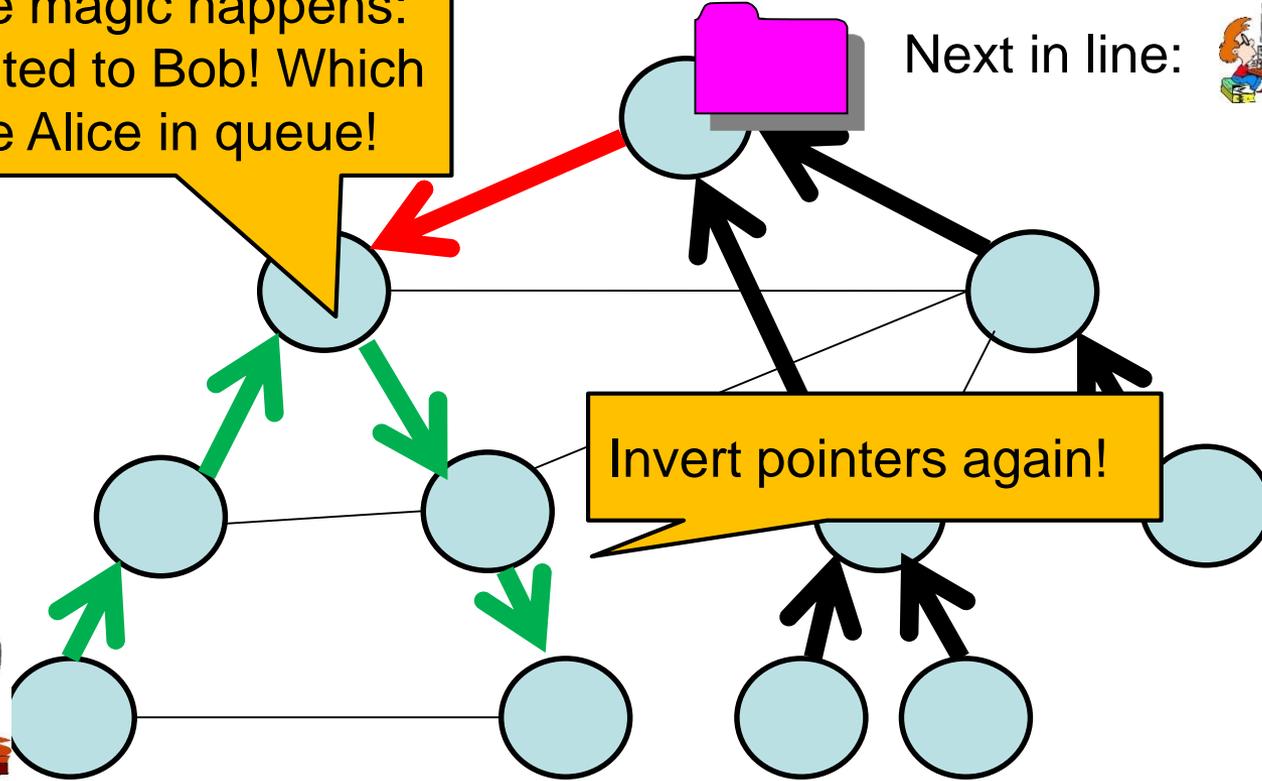
Next in line:



Invert pointers again!



Phase 2



The Arrow Protocol: Concurrency

Here the magic happens:
gets routed to Bob! Which
is before Alice in queue!

Still in use by
someone else!



Next in line:



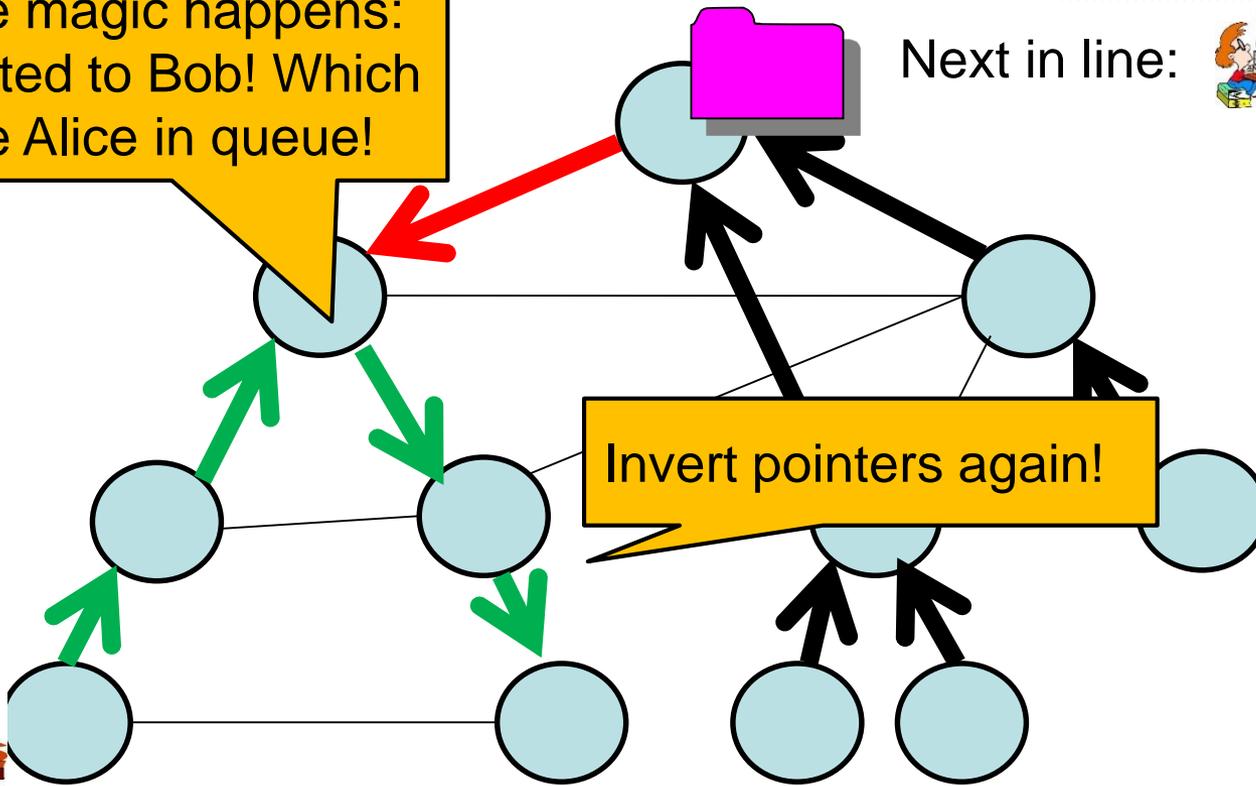
Invert pointers again!



Next in line:



Phase 2



The Arrow Protocol: Concurrency

Still in use by
someone else!



Next in line:



Here the magic happens:
gets routed to Bob! Which
is before Alice in queue!

And so on! Alice will be the
next root, users will be directed
to her: A distributed queue
implementation!



Next in line



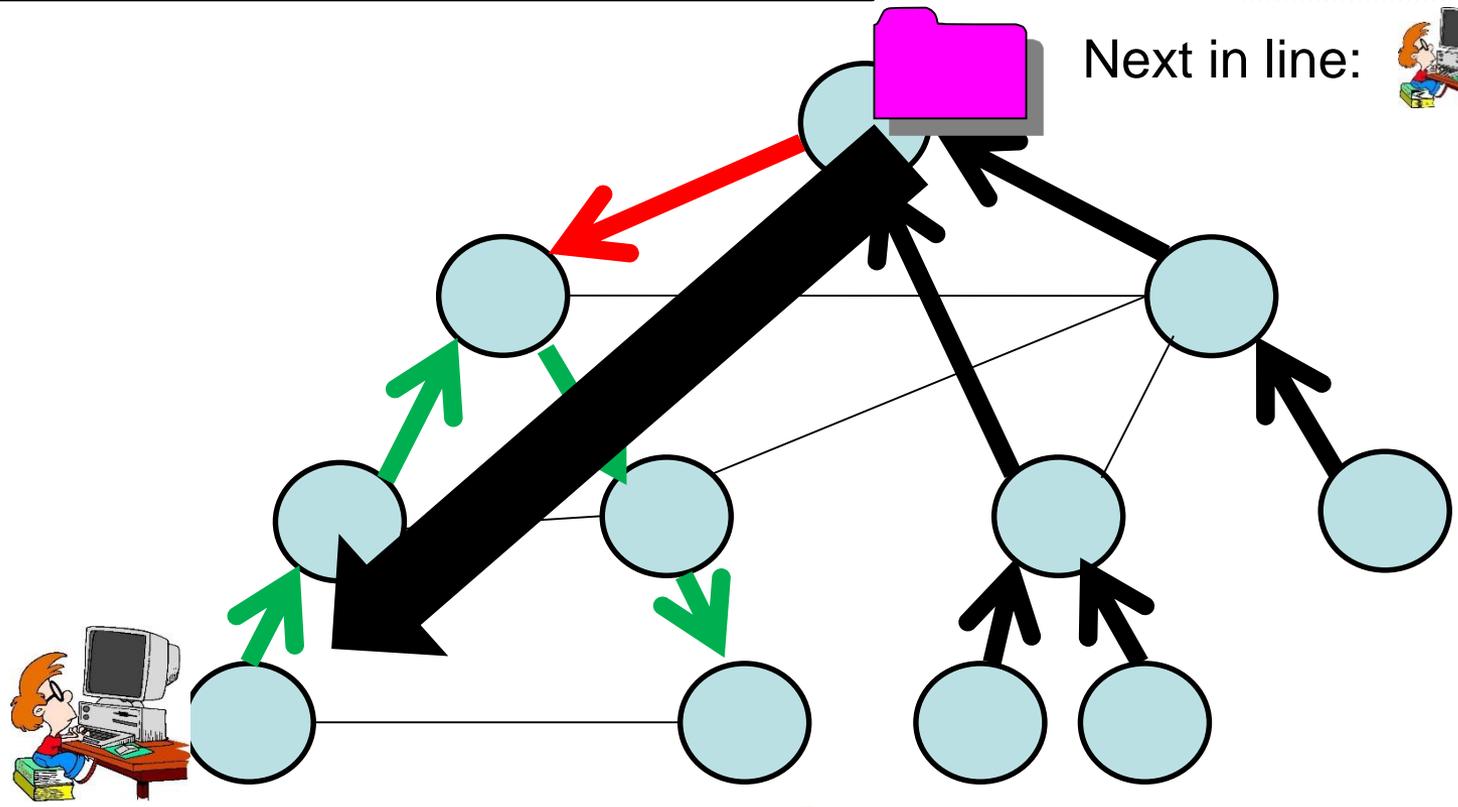
Remark: Of course, once object becomes free, it can be sent to Bob directly. (No link reversals, no spanning tree.)

ency

Still in use by someone else!



Next in line:



Next in line:



Arrow

Start Find Request at Node u :

- 1: **do atomically**
- 2: u sends “find by u ” message to parent node
- 3: $u.parent := u$
- 4: $u.wait := \mathbf{true}$
- 5: **end do**

Upon w Receiving “Find by u ” Message from Node v :

- 6: **do atomically**
- 7: **if** $w.parent \neq w$ **then**
- 8: w sends “find by u ” message to parent
- 9: $w.parent := v$
- 10: **else**
- 11: $w.parent := v$ **invert edge!**
- 12: **if not** $w.wait$ **then**
- 13: send variable to u // w holds var. but does not need it any more
- 14: **else**
- 15: $w.successor := u$ // w will send variable to u a.s.a.p.
- 16: **end if**
- 17: **end if** **wait myself?**
- 18: **end do**

Upon w Receiving Shared Object:

- 19: perform operation on shared object
- 20: **do atomically**
- 21: $w.wait := \mathbf{false}$
- 22: **if** $w.successor \neq \mathbf{null}$ **then**
- 23: send variable to $w.successor$
- 24: $w.successor := \mathbf{null}$
- 25: **end if**
- 26: **end do**

Analysis

- Message will only travel on static tree!
- And can never traverse an edge twice (in opposite direction).

Arrow

Arrow is correct: find() terminates with message and time complexity D , where D is the diameter of the spanning tree. Completely **asynchronous and concurrent** environments!

Proof.

- Each edge $\{u,v\}$ in the spanning tree is in one of four states:
 - (A) u points to v , no message on the edge, v does not point to u
 - (B) Message on the move from u to v (no pointer along edge)
 - (C) v points to u , no message on edge, u does not point to v
 - (D) Message on the move from v to u (no pointer along edge)

QED

End of Lecture
