


CSE113: Parallel Programming

Feb. 24, 2023

- **Topics:**

- General concurrent sets



Announcements

- Midterm was due last Friday
 - Everyone got it in on time; thanks!
 - We are planning to release grades either today or Monday
 - Let us know ASAP if there are issues; you have 1 week
 - make a private piazza post
- HW 2 has 2 extra free late days
 - get it in by Wednesday
- Because we canceled class, HW 3 will be released Wednesday as well

Previous Quiz

Previous Quiz

Which one of the following is NOT a drawback of a global workstealing parallel schedule

-
- Requires a concurrent data structure

 - Contention on shared cache lines

 - Contention on a single location with RMWs

Previous Quiz

Which of the following is NOT an overhead of the local worklist workstealing parallel schedule (that we studied in class)

-
- Initialization of the queues
 - Checking a global variable to ensure all work is completed
 - Managing concurrent enqueues to the worklists

Previous Quiz

Which of the following solutions can guarantee that a static schedule will not be out of bounds?

-
- The last thread always checks to get the minimum between the end of the array or the value allocated

 - The last thread always get the end of the array

 - The last thread never receives more than N tasks

Previous Quiz

Write a few sentences about the pros and cons of using local workstealing queues over the global implicit worklist

New material

C++ Atomic template

- C++ lets you wrap custom objects/types as an atomic type
- included in `<atomic>`
- use like this:
 - `atomic<int> i;`
 - `atomic<float> f;`

C++ Atomic template

- Lets you:
 - load
 - store
 - exchange
 - `compare_and_swap`
- It may use a lock behind the scenes!
- *Examples*

C++ Atomic template

- If you do this to a class, you will lose access to your methods!
- Pattern:
 - load the class atomically into a non atomic variable
 - operate on it
 - store it back. Be careful (others may have updated it!)

Schedule

- **Concurrent set**
 - Coarse-grained lock
 - fine-grained lock
 - optimistic locking

Thanks to Roberto Palmieri (Lehigh University) and material from the text book for some of the slide content/ideas.

Set Interface

- Unordered collection of items
- No duplicates

- We will implement this as a sorted linked list


Set Interface

- Unordered collection of items
- No duplicates
- Methods
 - **add (x)** put **x** in set
 - **remove (x)** take **x** out of set
 - **contains (x)** tests if **x** in set

List Node

```
class Node {  
    public:  
        Value v;  
        int key;  
        Node *next;  
}
```

The List-Based Set



Sorted with Sentinel nodes
(min & max possible keys)

Sequential List Based Set

add(b)




remove(b)




Sequential List Based Set

add(b)




remove(b)




Schedule

- Concurrent set
 - **Coarse-grained lock**
 - fine-grained lock
 - optimistic locking




Coarse-Grained Locking



Coarse-Grained Locking



Coarse-Grained Locking



Simple but inefficient!



Schedule

- Concurrent set
 - Coarse-grained lock
 - **fine-grained lock**
 - optimistic locking


Fine-grained Locking

- Requires **careful** thought
- Split object into pieces
 - Each piece has own lock
 - Methods that work on disjoint pieces need not exclude each other


Hand-over-Hand locking




Hand-over-Hand locking




Hand-over-Hand locking





Hand-over-Hand locking



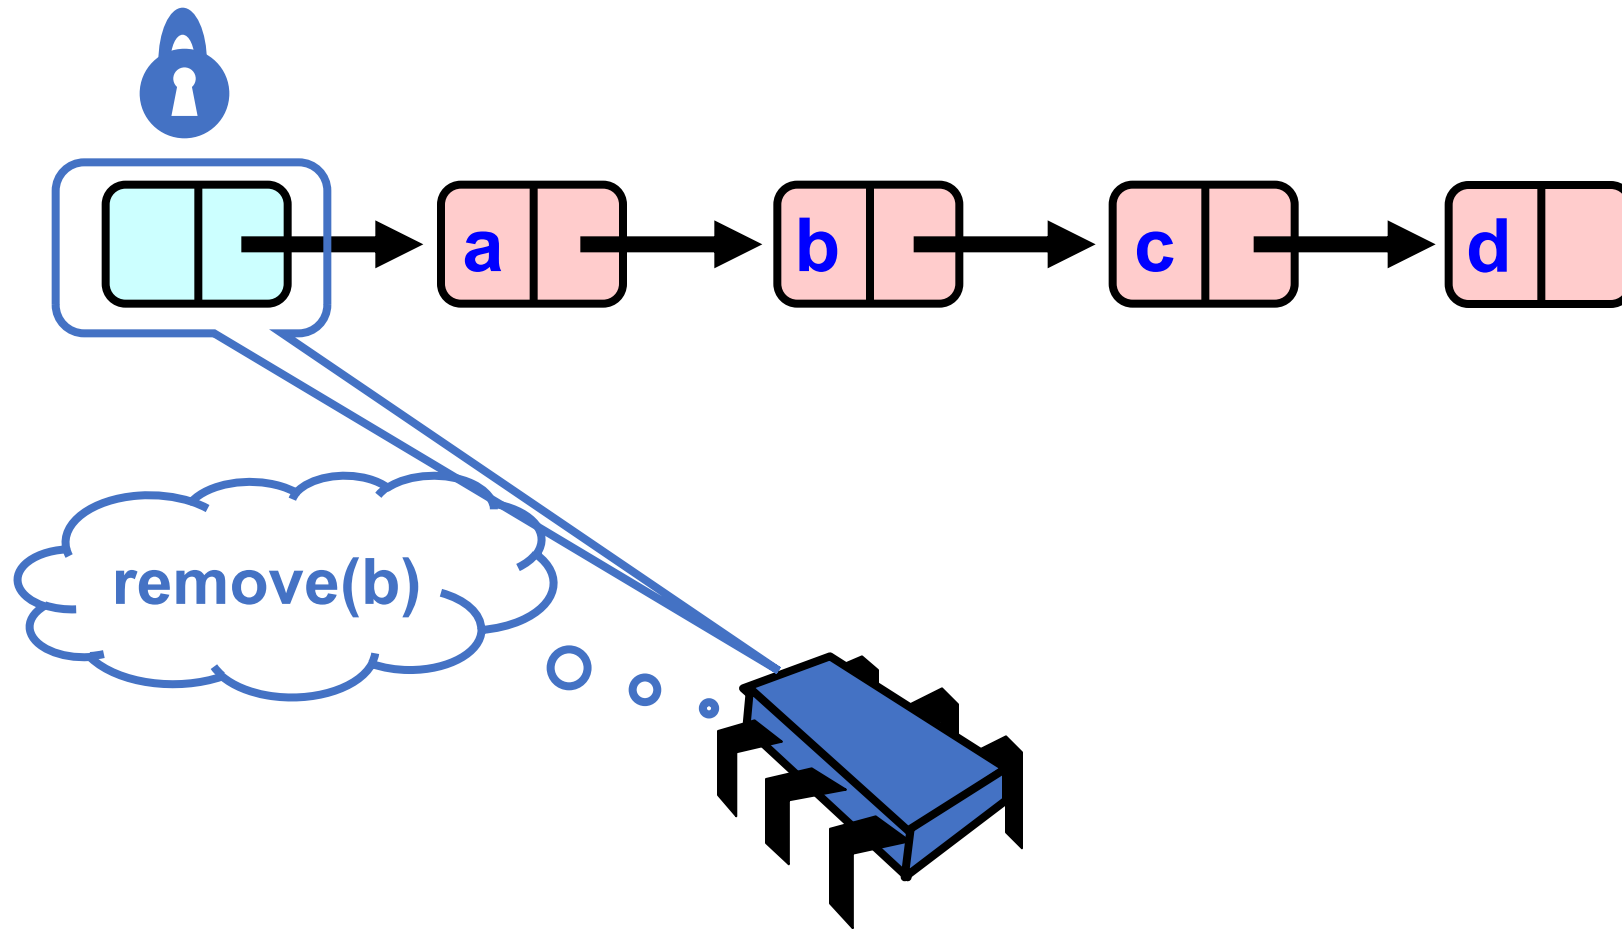
Hand-over-Hand locking




Removing a Node




Removing a Node




Removing a Node




Removing a Node




Removing a Node




Removing a Node




Removing a Node




Concurrent Removes




Concurrent Removes




Concurrent Removes




Concurrent Removes




Concurrent Removes




Concurrent Removes




Concurrent Removes



Concurrent Removes




Uh, Oh





Uh, Oh

Bad news, c not removed



Problem



- To delete node c
 - Swing node b's next field to d
- Problem is,
 - **Data conflict:**
 - Someone deleting b concurrently could direct a pointer to C




Insight

- If a node is locked
 - No one can delete node's *successor*
- If a thread locks
 - Node to be deleted
 - And its predecessor
 - Then it works


Hand-Over-Hand Again




Hand-Over-Hand Again




Hand-Over-Hand Again




Hand-Over-Hand Again




Hand-Over-Hand Again




Hand-Over-Hand Again




Removing a Node




Removing a Node




Removing a Node




Removing a Node




Removing a Node




Removing a Node




Removing a Node




Removing a Node




Removing a Node




Removing a Node




Removing a Node




Removing a Node



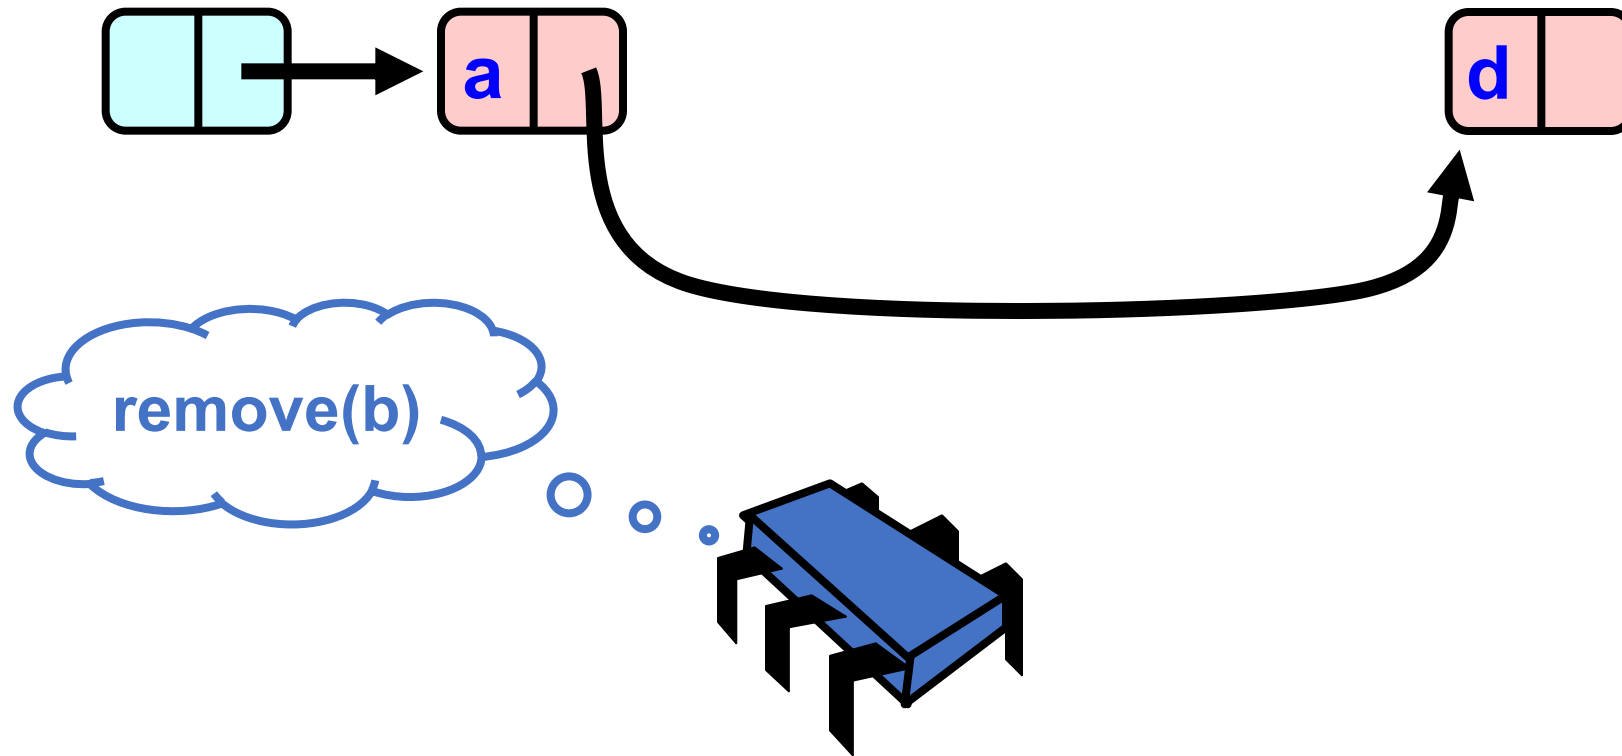
Removing a Node




Removing a Node



Removing a Node



Removing a Node



Adding Nodes


- To add node e
 - Must lock predecessor
 - Must lock successor
- Neither can be deleted

Drawbacks


- Better than coarse-grained lock
 - Threads can traverse in parallel
- Still not ideal
 - Long chain of acquire/release
 - Inefficient


```
void remove(Value v) {
    Node* pred = NULL, *curr = NULL;
    head.lock();
    pred = head;
    curr = pred.next();
    curr.lock();
    while (curr.value != v) {
        pred.unlock();
        pred = curr;
        curr = curr.next();
        curr.lock();
    }
    pred.next = curr.next;
    curr.unlock();
    pred.unlock();
}
```


```
void remove(Value v) {
    Node* pred = NULL, *curr = NULL;
    head.lock();
    pred = head;
    curr = pred.next();
    curr.lock();
    while (curr.value != v) {
        pred.unlock();
        pred = curr;
        curr = curr.next();
        curr.lock();
    }
    pred.next = curr.next;
    curr.unlock();
    pred.unlock();
}
```




```
void remove(Value v) {
    Node* pred = NULL, *curr = NULL;
    head.lock();
    pred = head;
    curr = pred.next();
    curr.lock();
    while (curr.value != v) {
        pred.unlock();
        pred = curr;
        curr = curr.next();
        curr.lock();
    }
    pred.next = curr.next;
    curr.unlock();
    pred.unlock();
}
```




```
void remove(Value v) {  
    Node* pred = NULL, *curr = NULL;  
    head.lock();  
    pred = head;  
    curr = pred.next();  
    curr.lock();  
    while (curr.value != v) {  
        pred.unlock();  
        pred = curr;  
        curr = curr.next();  
        curr.lock();  
    }  
    pred.next = curr.next;  
    curr.unlock();  
    pred.unlock();  
}
```




```
void remove(Value v) {
    Node* pred = NULL, *curr = NULL;
    head.lock();
    pred = head;
    curr = pred.next();
    curr.lock();
    while (curr.value != v) {
        pred.unlock();
        pred = curr;
        curr = curr.next();
        curr.lock();
    }
    pred.next = curr.next;
    curr.unlock();
    pred.unlock();
}
```




```
void remove(Value v) {
    Node* pred = NULL, *curr = NULL;
    head.lock();
    pred = head;
    curr = pred.next();
    curr.lock();
    while (curr.value != v) {
        pred.unlock();
        pred = curr;
        curr = curr.next();
        curr.lock();
    }
    pred.next = curr.next;
    curr.unlock();
    pred.unlock();
}
```




```
void remove(Value v) {
    Node* pred = NULL, *curr = NULL;
    head.lock();
    pred = head;
    curr = pred.next();
    curr.lock();
    while (curr.value != v) {
        pred.unlock();
        pred = curr;
        curr = curr.next();
        curr.lock();
    }
    pred.next = curr.next;
    curr.unlock();
    pred.unlock();
}
```





```
void remove(Value v) {  
    Node* pred = NULL, *curr = NULL;  
    head.lock();  
    pred = head;  
    curr = pred.next();  
    curr.lock();  
    while (curr.value != v) {  
        pred.unlock();  
        pred = curr;  
        curr = curr.next();  
        curr.lock();  
    }  
    pred.next = curr.next;  
    curr.unlock();  
    pred.unlock();  
}
```




```
void remove(Value v) {  
    Node* pred = NULL, *curr = NULL;  
    head.lock();  
    pred = head;  
    curr = pred.next();  
    curr.lock();  
    while (curr.value != v) {  
        pred.unlock();  
        pred = curr;  
        curr = curr.next();  
        curr.lock();  
    }  
    pred.next = curr.next;  
    curr.unlock();  
    pred.unlock();  
}
```



```
void remove(Value v) {
    Node* pred = NULL, *curr = NULL;
    head.lock();
    pred = head;
    curr = pred.next();
    curr.lock();
    while (curr.value != v) {
        pred.unlock();
        pred = curr;
        curr = curr.next();
        curr.lock();
    }
    pred.next = curr.next;
    curr.unlock();
    pred.unlock();
}
```



Schedule

- Concurrent set
 - Coarse-grained lock
 - fine-grained lock
 - **optimistic locking**

How can we improve

- Acquires and releases lock for every node traversed
 - If we have a long list to search, it can be bad!
 - reduces concurrency (traffic jams)

Optimistic Synchronization

Assume there will be no conflicts. Check before committing. If there was a conflict, try again.

Optimistic Synchronization

- Find nodes without locking


Optimistic Synchronization

- Find nodes without locking
- Lock nodes


Optimistic Synchronization

- Find nodes without locking
- Lock nodes
- Check that everything is OK


Optimistic: Traverse without Locking




Optimistic: Lock and Load




Optimistic: Lock and Load




What could go wrong?



What could go wrong?



What could go wrong?



Data conflict!

- Red thread has the lock on a node (so it can modify the node)
- Blue thread is traversing without locks
- What do we do?

Data conflict!

- Red thread has the lock on a node (so it can modify the node)
- Blue thread is traversing without locks
- What do we do? We decided that locking when traversing is too expensive.

Lock-free reasoning

- We can use atomic variables

Lock-free reasoning

- Default atomic accesses are documented to be sequentially consistent.

```
class Node {  
    public:  
        Value v;  
        int key;  
        Node *next;  
}
```

Lock-free reasoning

- Default atomic accesses are documented to be sequentially consistent.

```
class Node {  
    public:  
        Value v;  
        int key;  
        atomic<Node*> next;  
}
```

Create an atomic pointer type using C++ templates

Lock-free reasoning

- Default atomic accesses are documented to be sequentially consistent.


```
void traverse(node *n) {  
    while (n->next != NULL) {  
        n = n->next;  
    }  
}
```

Lock-free reasoning


- Default atomic accesses are documented to be sequentially consistent.

```
void traverse(node *n) {  
    while (n->next.load() != NULL) {  
        n = n->next.load();  
    }  
}
```


What could go wrong?



What could go wrong?




What could go wrong?




No more data conflict, but we do need to reason about interleavings and threads concurrent threads contending for values.


What could go wrong?




What could go wrong?



What could go wrong?



Validate – Part 1



What happens if failure?

- Ideas?

What happens if failure?


- Could try to recover? Back up a node?
 - Very tricky!
 - Just start over!

What happens if failure?


- Could try to recover? Back up a node?
 - Very tricky!
 - Just start over!
- Private method:
 - `try_remove`
 - remove loops on `try_remove` until it succeeds

What about deletion?


Can threads that remove a node delete it?




Can threads that remove a node delete it?




Can threads that remove a node delete it?



Can threads that remove a node delete it?



Our own garbage collector



Java's garbage collection will remove b

We are using a better™ language though...

maintain a list to delete:


Our own garbage collector



Java's garbage collection will remove b

We are using a better™ language though...

maintain a list to delete:





Our own garbage collector




Java's garbage collection will remove b

We are using a better™ language though...

maintain a list to delete:




Our own garbage collector




Java's garbage collection will remove b



We are using a better™ language though...




maintain a list to delete:




Our own garbage collector



maintain a list to delete:





Our own garbage collector



Similar to a reader/writer lock:
Allows an arbitrary number of threads that operate on the list
Only 1 garbage collector thread
Erases the list of nodes

maintain a list to delete:




Garbage collector lock

- Many strategies!
 - A big research area ~10 years ago
- **Strat 1:** Threads always try once to take the garbage collector lock:
 - if failed, no worries, the next operation will get a chance
 - if succeeded, then there was no contention
 - can starve garbage collection
- **Strat 2:** Wait until size grows to a threshold:
 - Wait on the lock (hope for a fair implementation!)
 - Can cause performance spikes


Back to the linked list

What if 2 threads try to add a node in the same position?


What Else Could Go Wrong?




What Else Could Go Wrong?



What Else Could Go Wrong?




What Else Could Go Wrong?

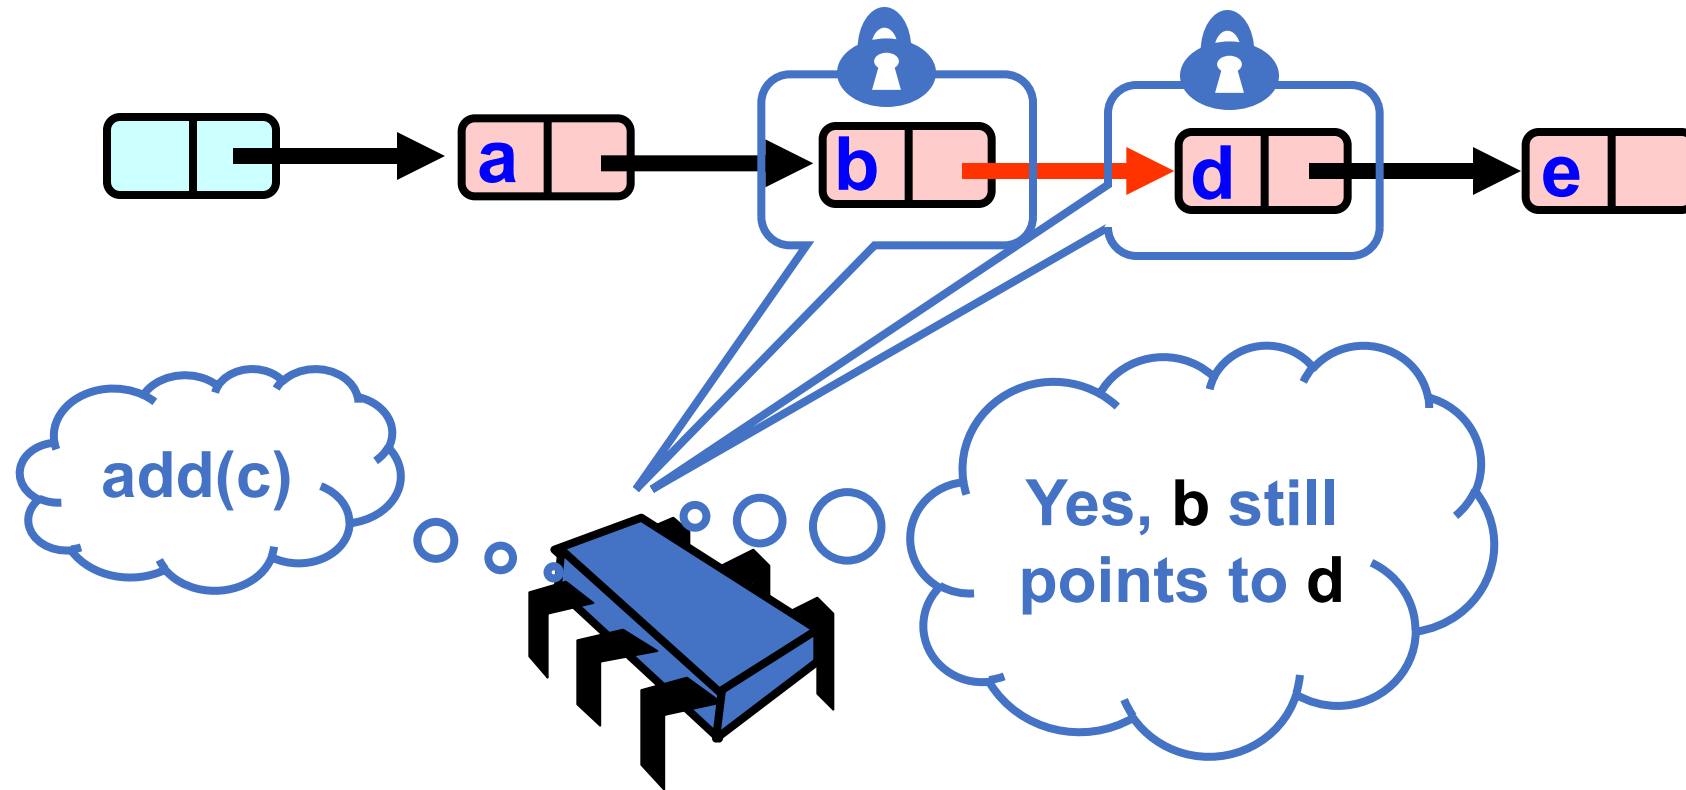


What Else Could Go Wrong?

Validation passes!



Validate Part 2 (while holding locks)



Summary

- We traverse without lock
 - Traversal may access nodes that are locked
 - Its okay because we have atomic pointers!
- We might traverse deleted nodes
 - Its okay because we validate after we obtain locks
 - Two validations:
 - our node is still reachable (it was not deleted)
 - Our insertion point is still valid (no thread has inserted in the meantime)
- We don't actually free node memory, but we put them in a list to be freed later

Enjoy your weekend!

- On Monday: making the list lock-free!
- Get HW 2 in and look for midterm grades