

Verifying a Lazy Concurrent List-Based Set Algorithm in Iris

Daniel Nezamabadi

Supervisor: Isaac van Bakel, Prof. Dr. Ralf Jung

Research in Computer Science, ETH Zurich

February 3rd, 2025



We want a set algorithm that...

...is concurrent

...has good performance

...implementable

...is correct

Our Candidate:

A Lazy Concurrent List-Based Set Algorithm

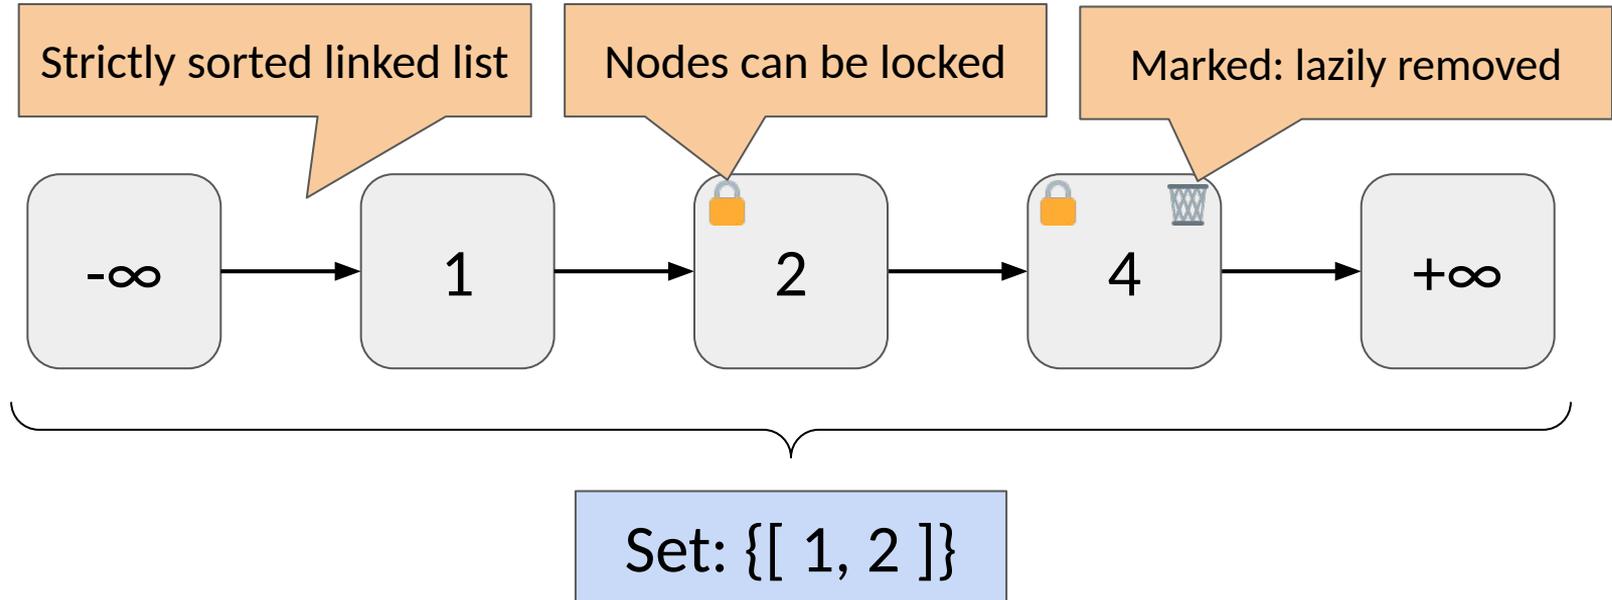
Steve Heller¹, Maurice Herlihy², Victor Luchangco¹, Mark Moir¹, William N. Scherer III³, and Nir Shavit¹

¹ Sun Microsystems Laboratories

² Brown University

³ University of Rochester

Anatomy of a Lazy List-Based Set



High-Level Properties

Less lock contention,
good for performance

Optimistic locking:
Don't lock while walking the list

Common design
technique

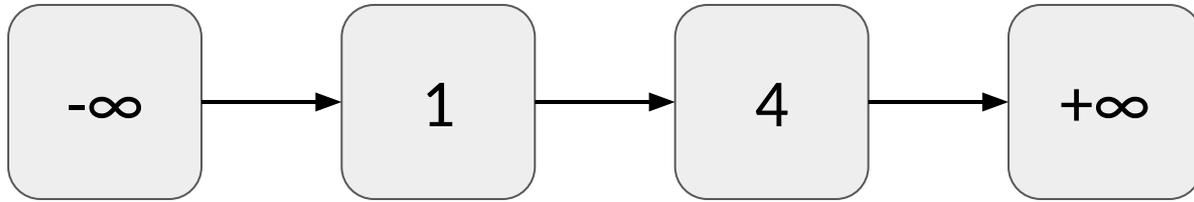
Lazy removal:
First mark a node, then unlink it

No lock (contention)

Wait-free contains:
Always finishes in finite #steps

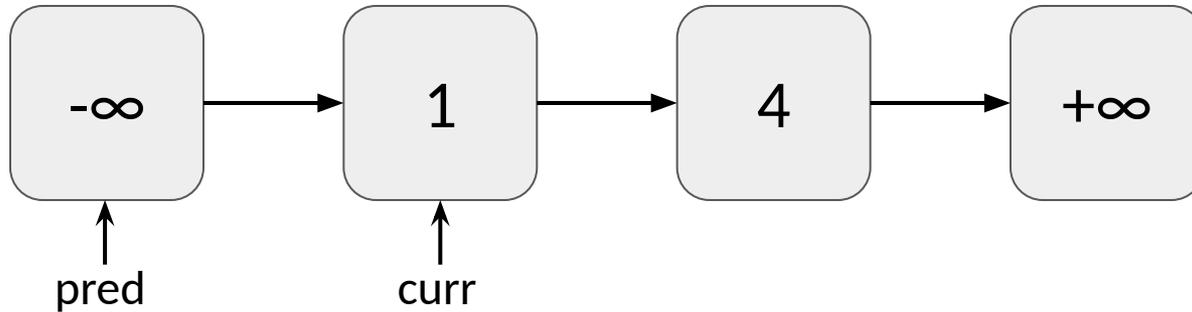
add(2)

Initial State



add(2)

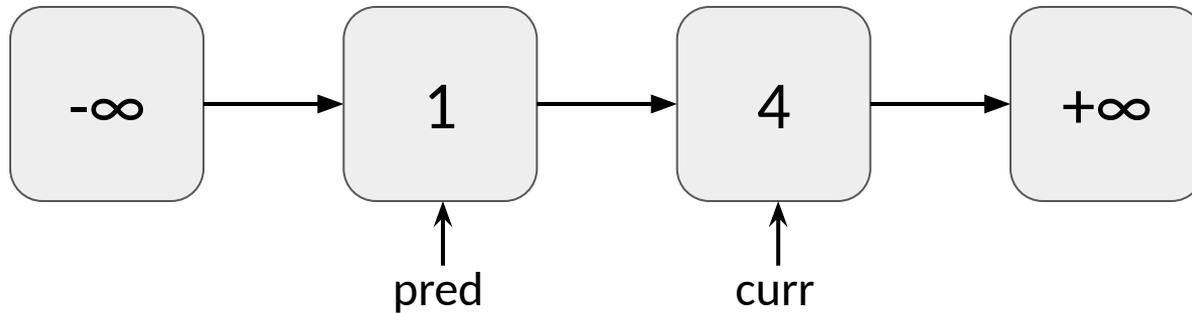
Thread 1 - add(2): walk



Optimistic: No locking yet

add(2)

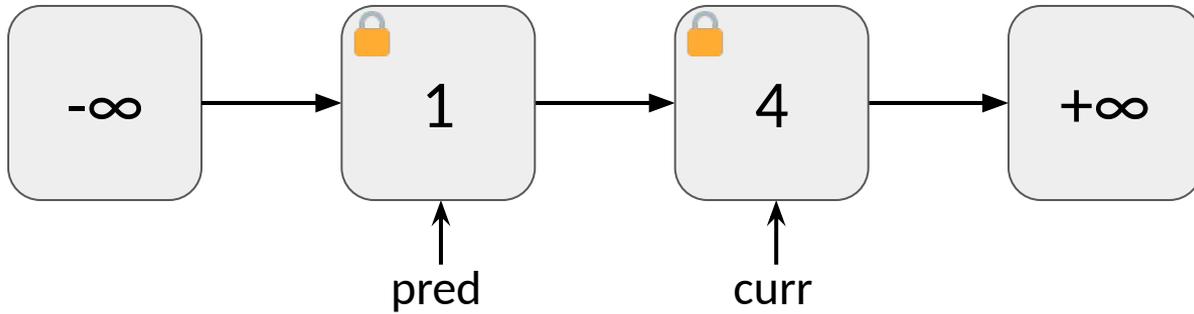
Thread 1 - add(2): walk is done



Optimistic: No locking yet

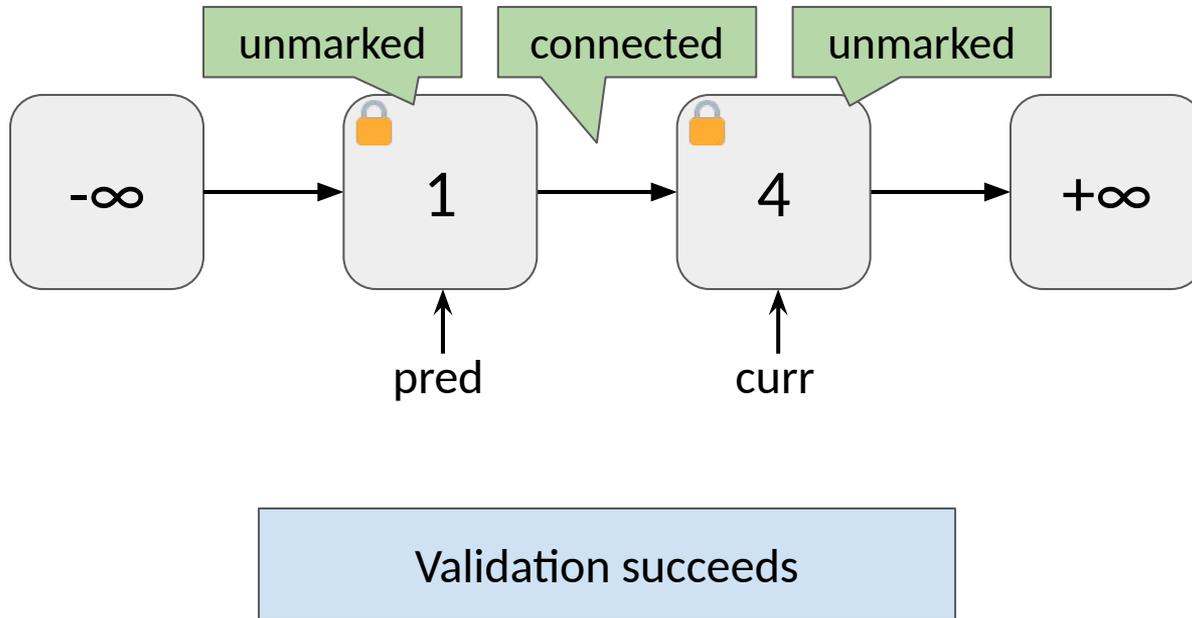
add(2)

Thread 1 - add(2): acquire locks



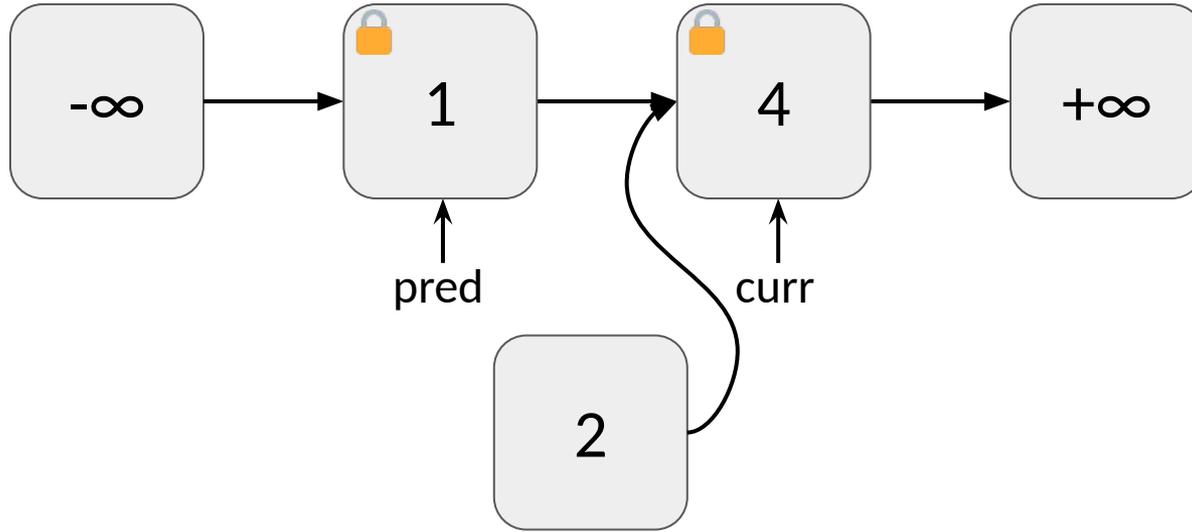
add(2)

Thread 1 - add(2): validate state



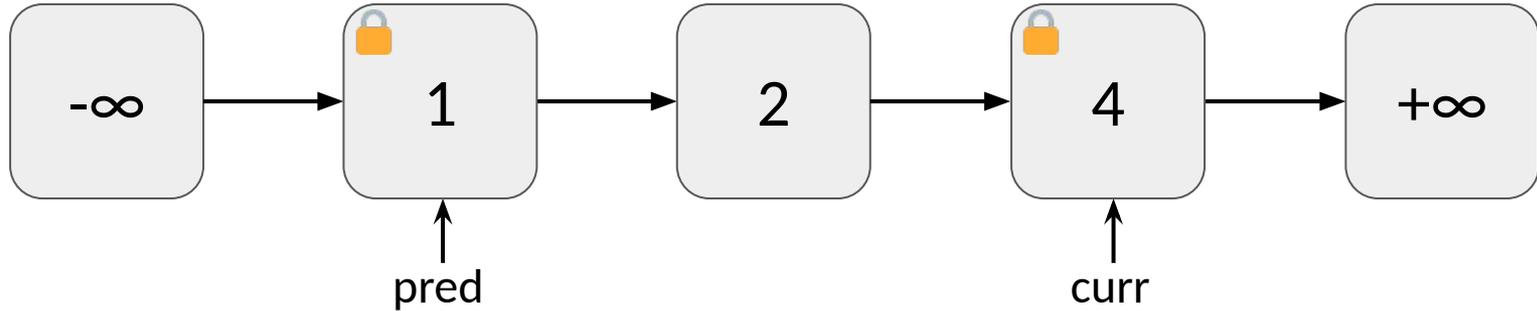
add(2)

Thread 1 - add(2): create new node



add(2)

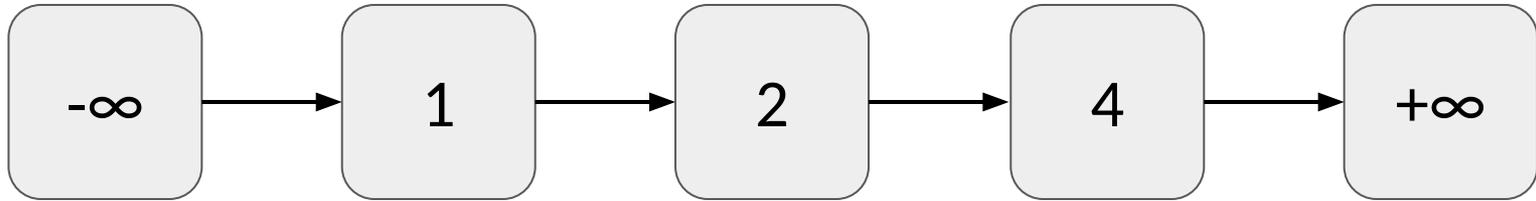
Thread 1: add(2) - **hook in new node**



Linearization Point: $\{[1, 4]\} \rightarrow \{[1, 2, 4]\}$

add(2)

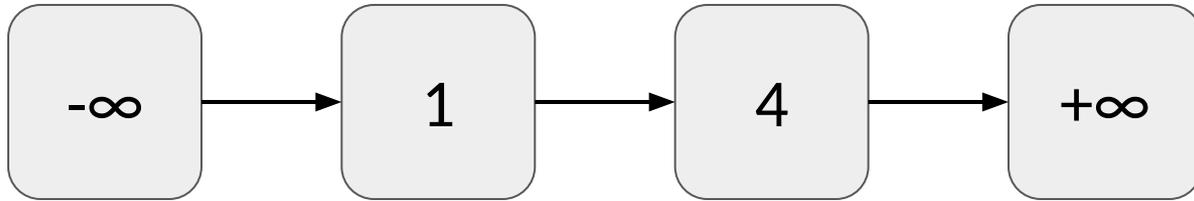
Thread 1: release locks



What if we are too optimistic?

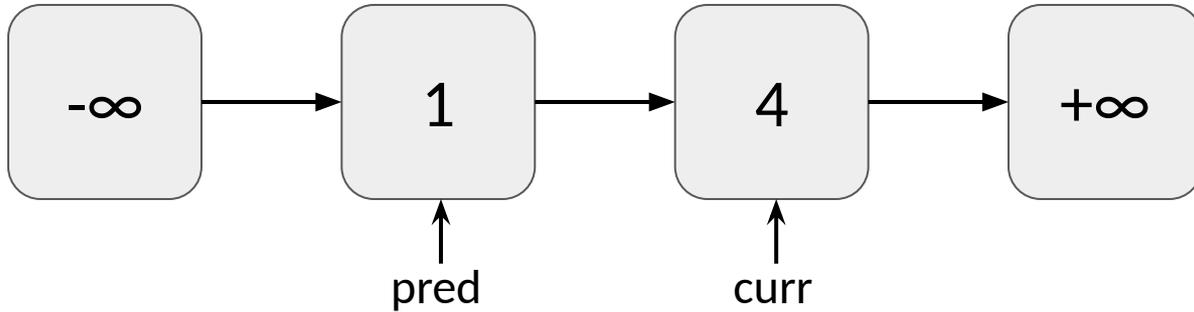
`add(2) || add(2) ;; add(3)`

Initial State



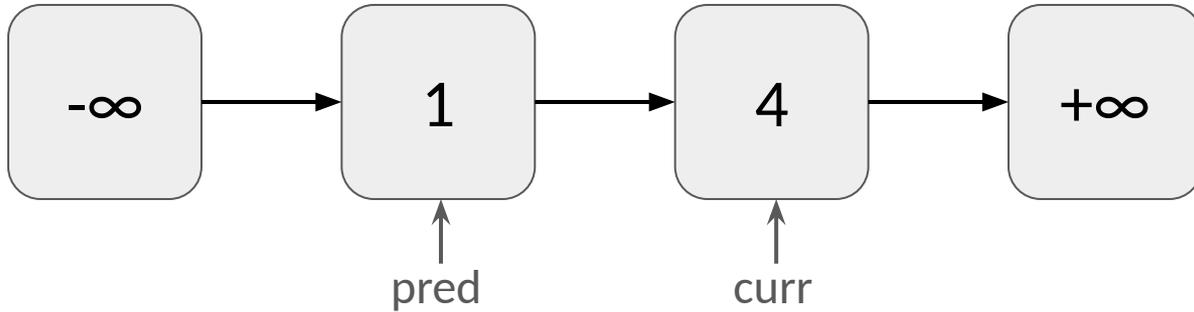
add(2) || add(2) ;; add(3)

Thread 1 - add(2): walk is done



add(2) || add(2) ;; add(3)

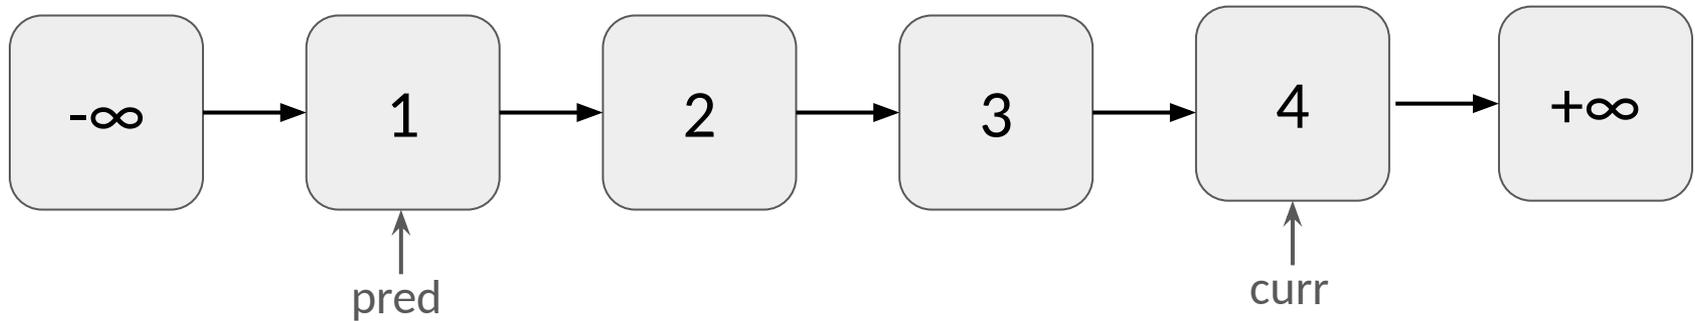
Thread 1 - add(2): preempted



`add(2) || add(2) ;; add(3)`

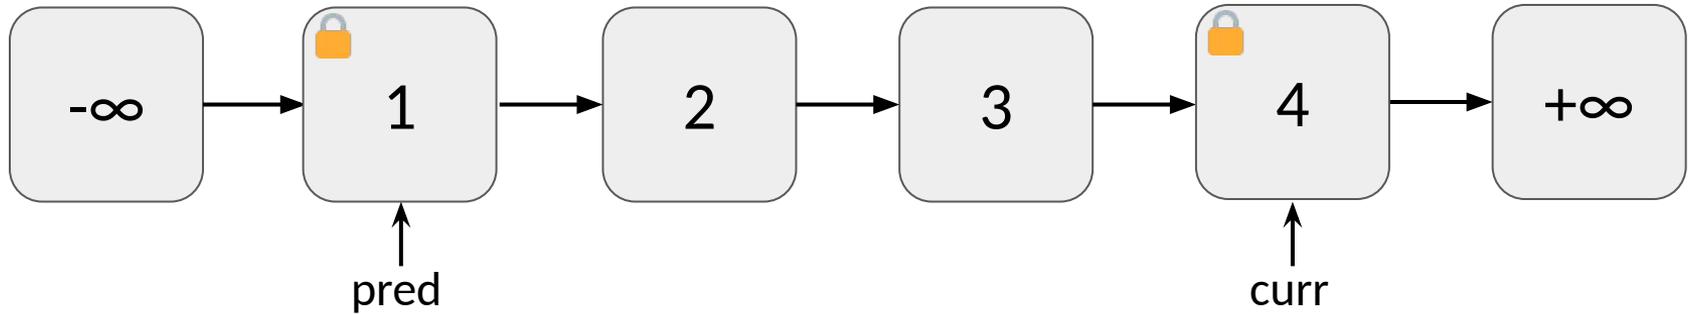
Thread 1 - `add(2)`: preempted

Thread 2 - adds 2 and 3



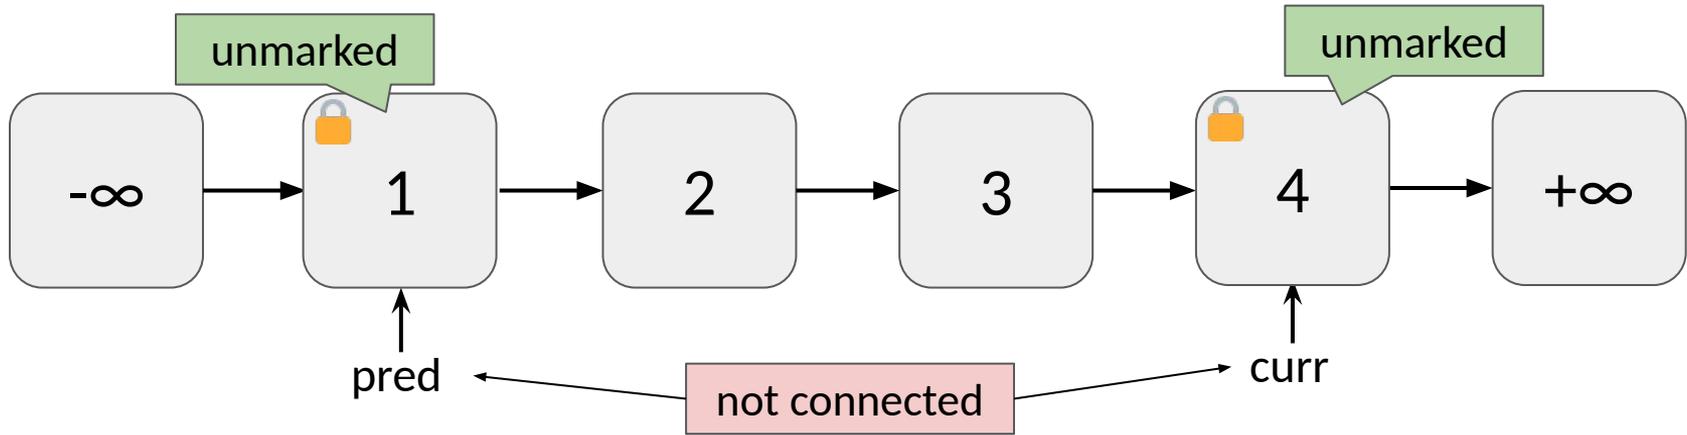
`add(2) || add(2) ;; add(3)`

Thread 1 - `add(2)`: acquire locks



add(2) || add(2) ;; add(3)

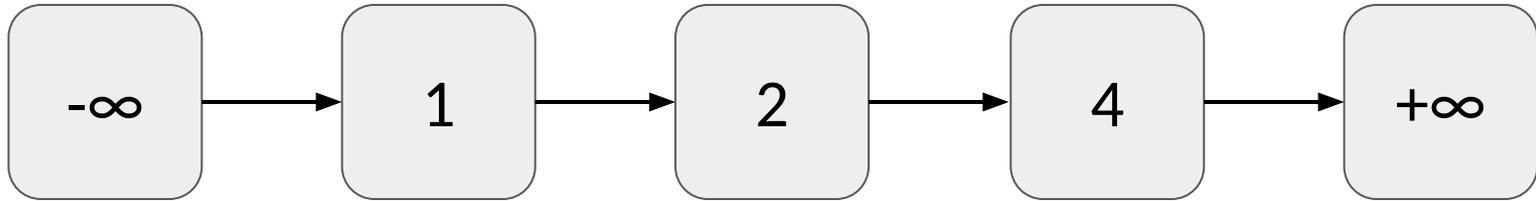
Thread 1 - add(2): validate state



Continuing would remove 3 \Rightarrow Abort & Retry

remove(4)

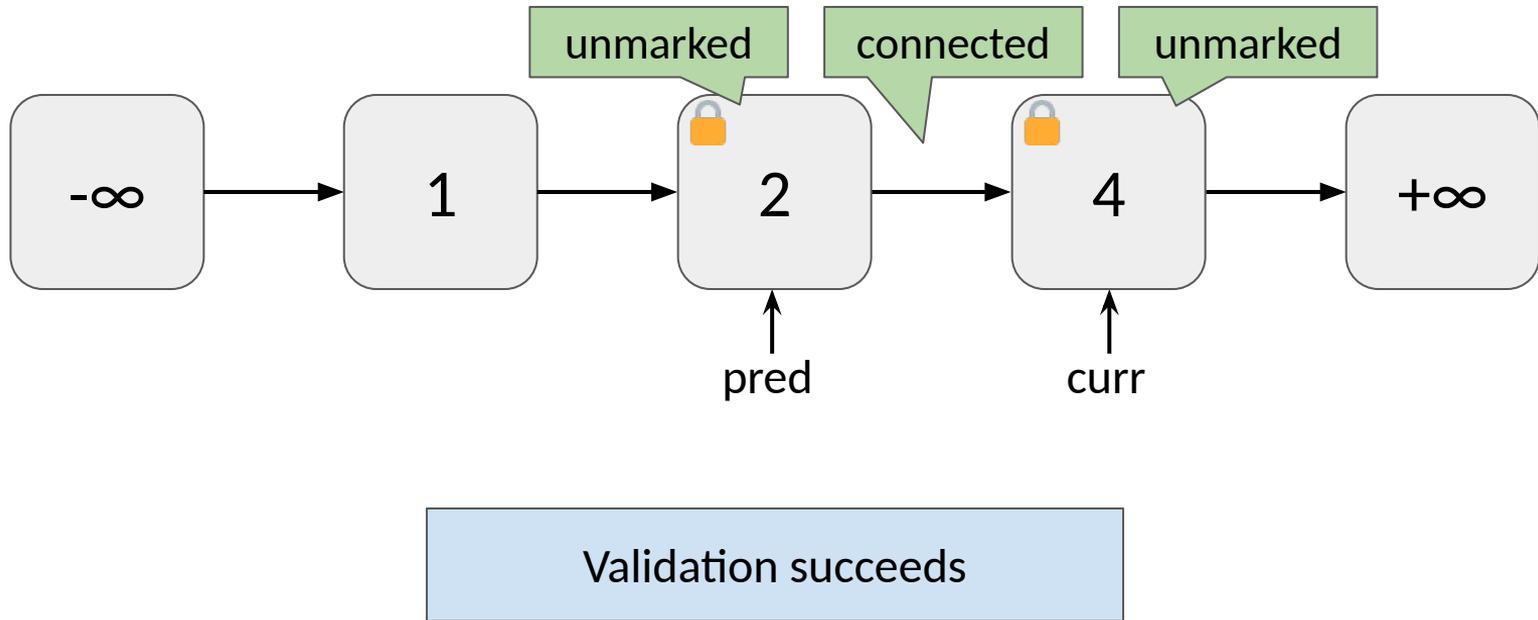
Initial State



walk, lock, and validate are the same as add

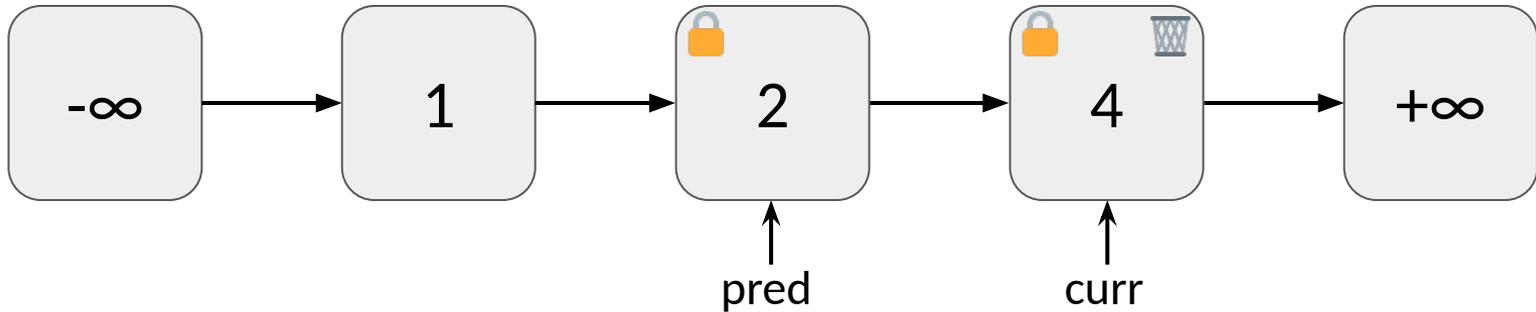
remove(4)

Thread 1 - remove(4): validate state



remove(4)

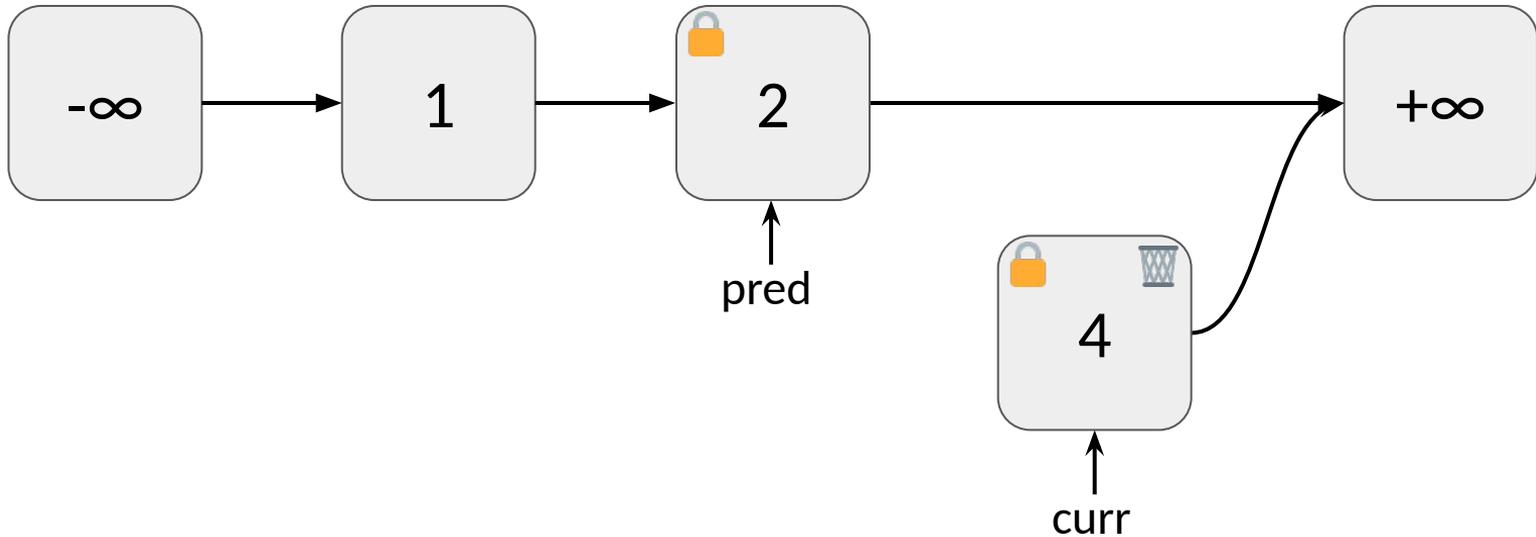
Thread 1 - remove(4): **lazily remove 4**



Linearization Point: $\{[1, 2, 4]\} \rightarrow \{[1, 2]\}$

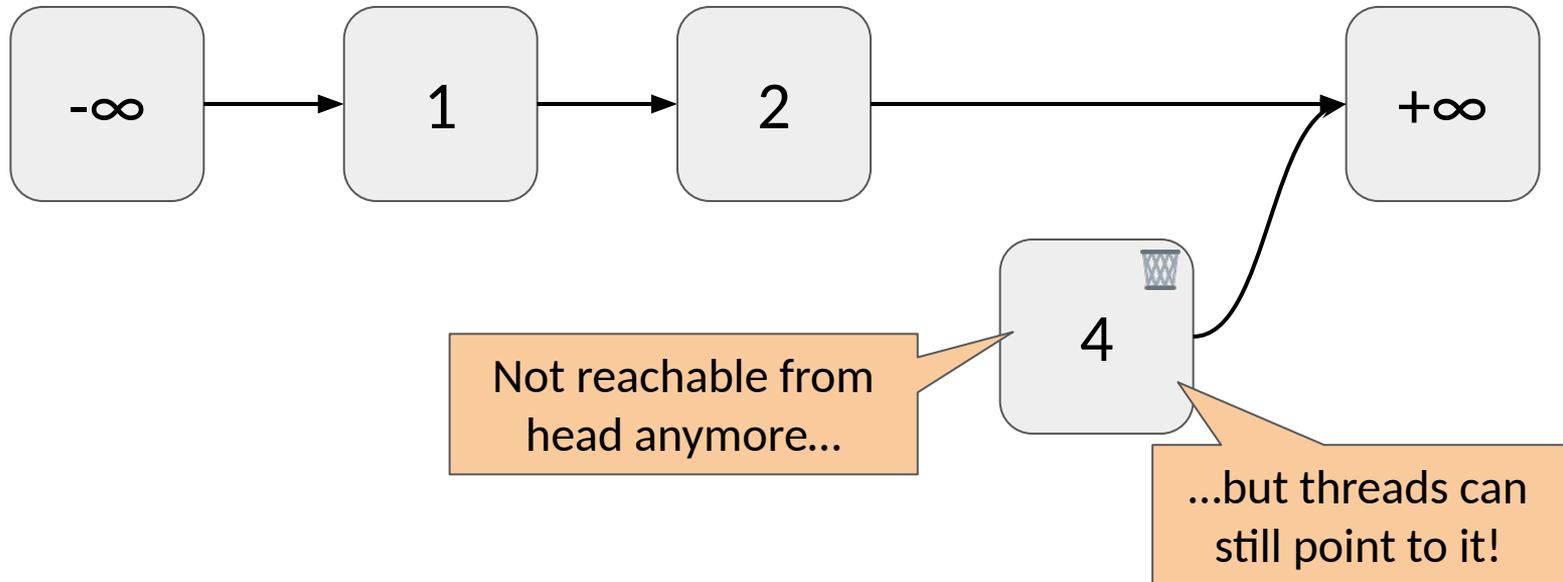
remove(4)

Thread 1 - remove(4): physically remove 4



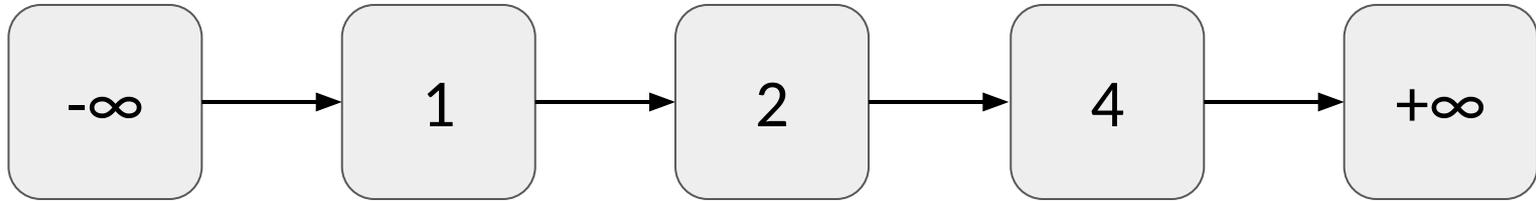
remove(4)

Thread 1 - remove(4): release locks



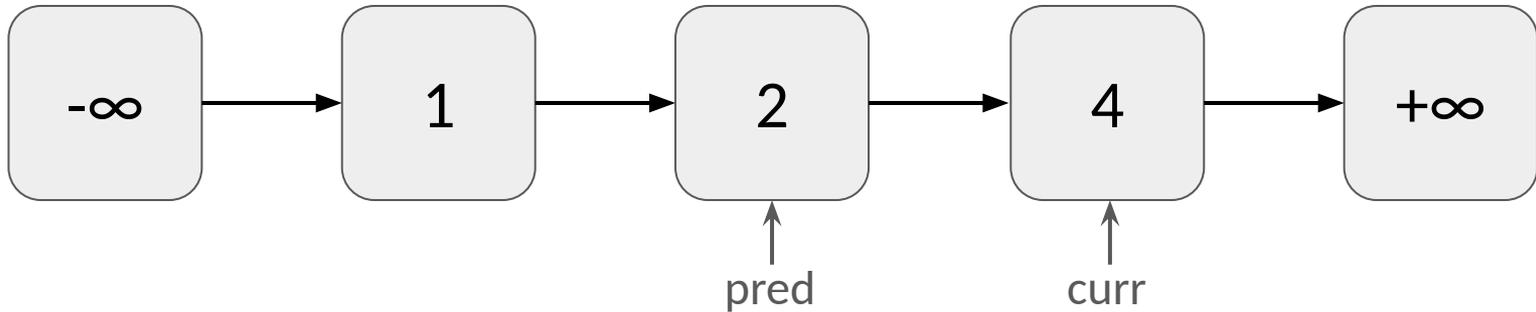
add(5) || remove(4)

Initial State



add(5) || remove(4)

Thread 1 - add(5): preempted during walk

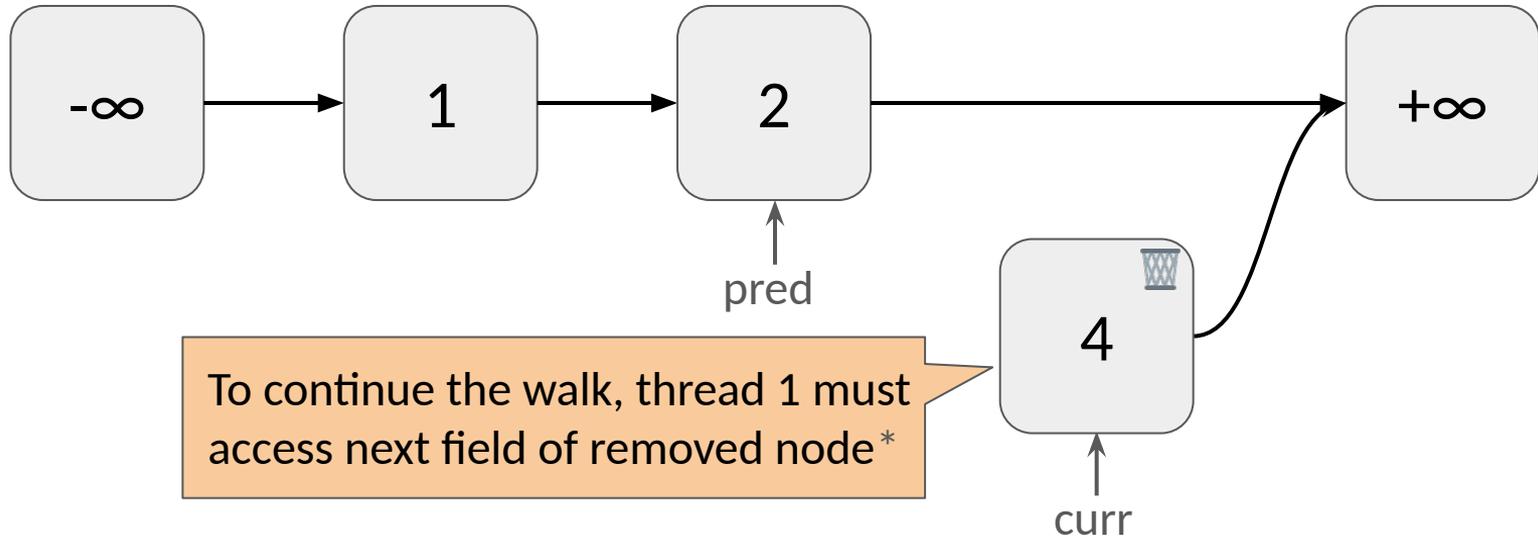


add(5) || remove(4)

Thread 1 - add(5): preempted during walk

Thread 2 - removes 4

Possible, because no locking during walk



* Figured this out the hard way.

What about contains?

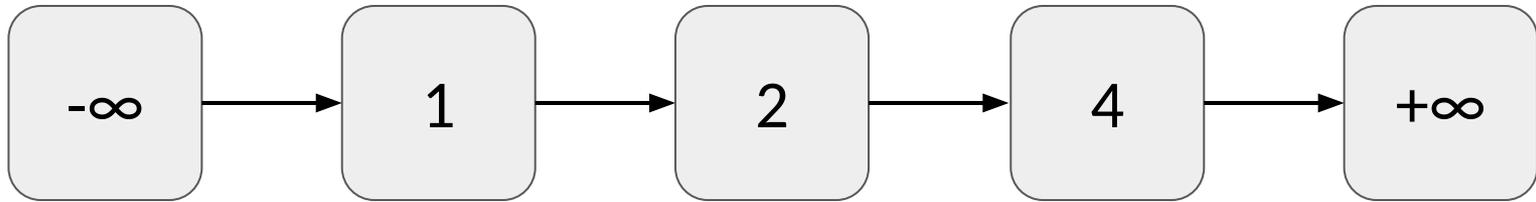
```
public boolean contains(int key) {  
    Entry curr = this.head;  
    while (curr.key < key)  
        curr = curr.next;  
    return curr.key == key && !curr.marked;  
}
```

wait-free \Rightarrow
lock-free

Simple implementation, tricky linearization point.

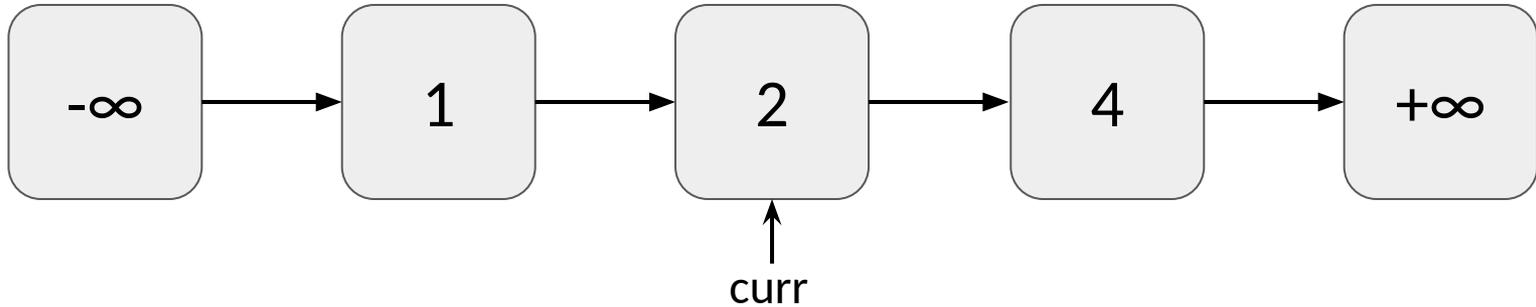
contains(2) || remove(2)

Initial State



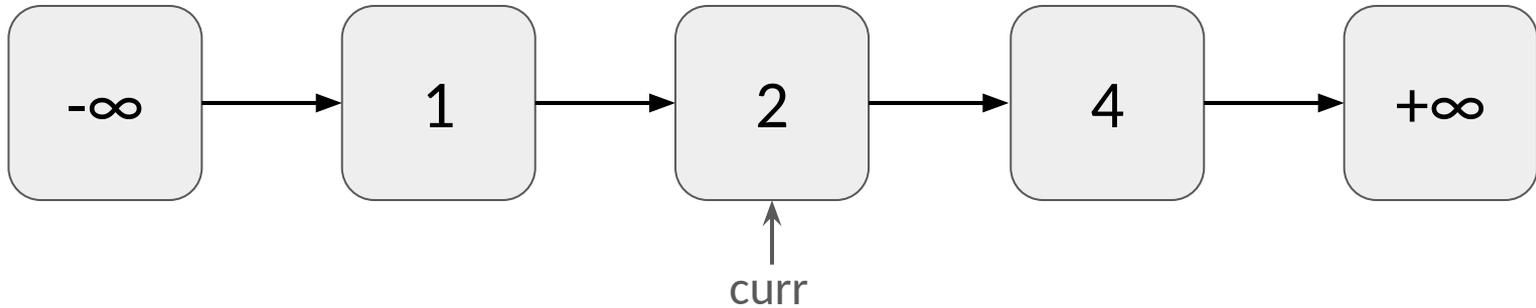
contains(2) || remove(2)

Thread 1 - contains(2): walk is done



contains(2) || remove(2)

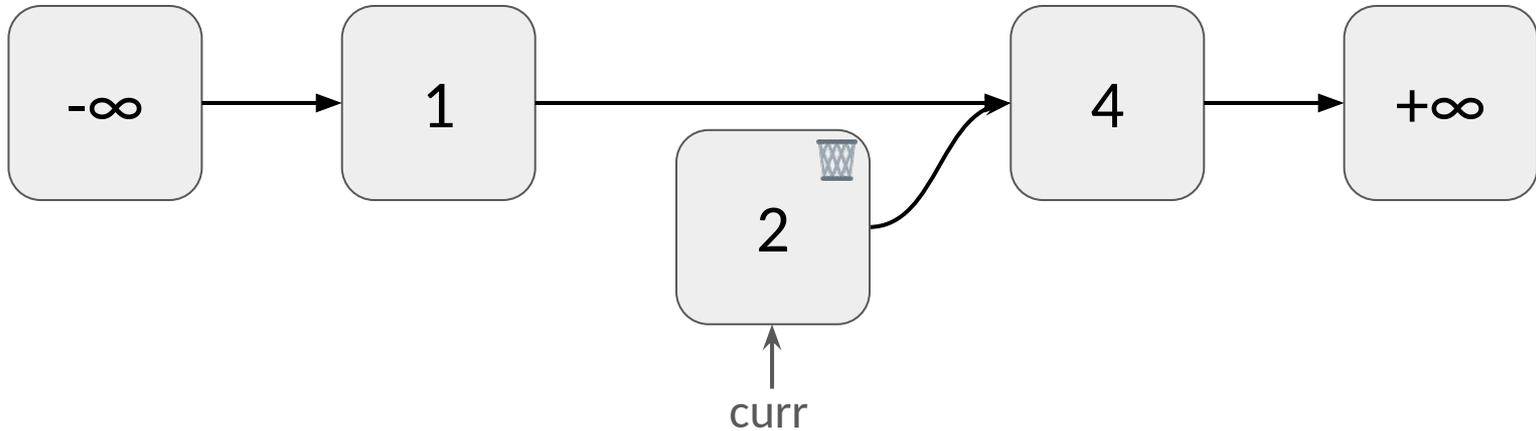
Thread 1 - contains(2): preempted



contains(2) || remove(2)

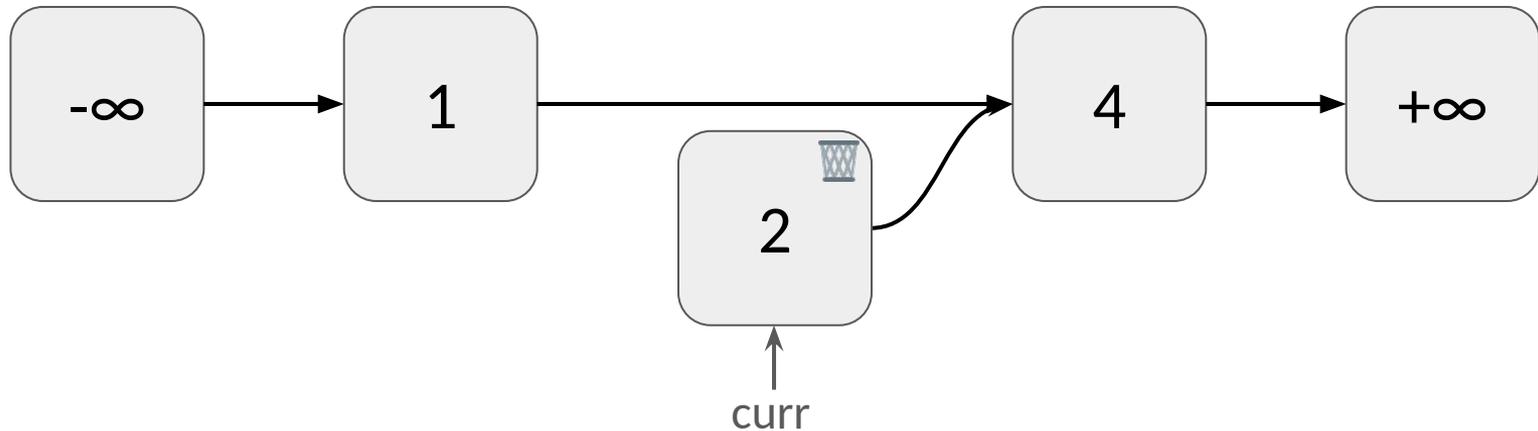
Thread 1 - contains(2): preempted

Thread 2 removes 2



contains(2) || remove(2)

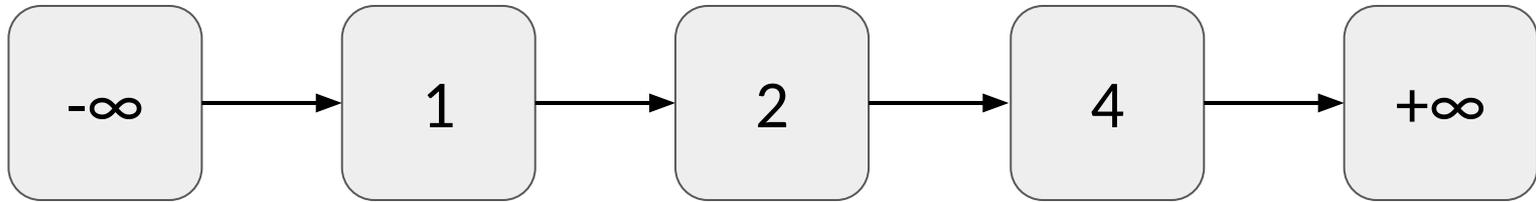
Thread 1 - contains(2): reads marked is true



Linearization Point: 2 is not in the set

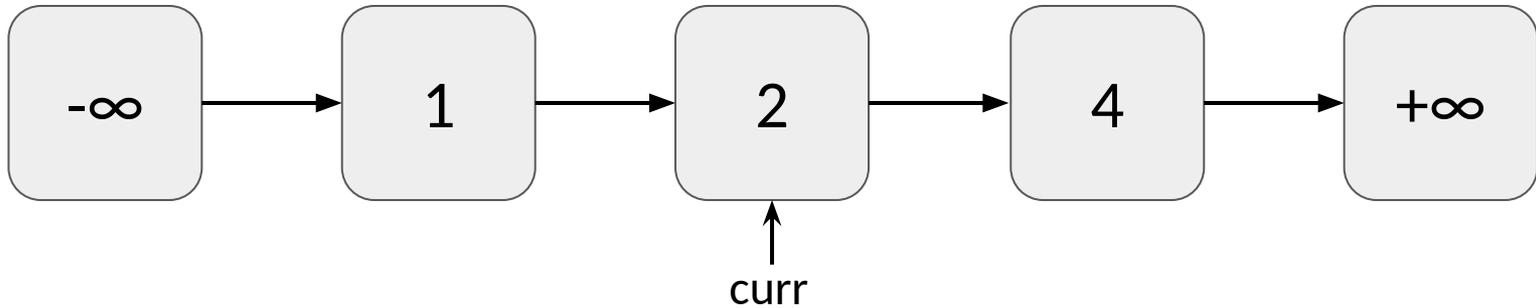
contains(2) || remove(2) ;; add(2)

Initial State



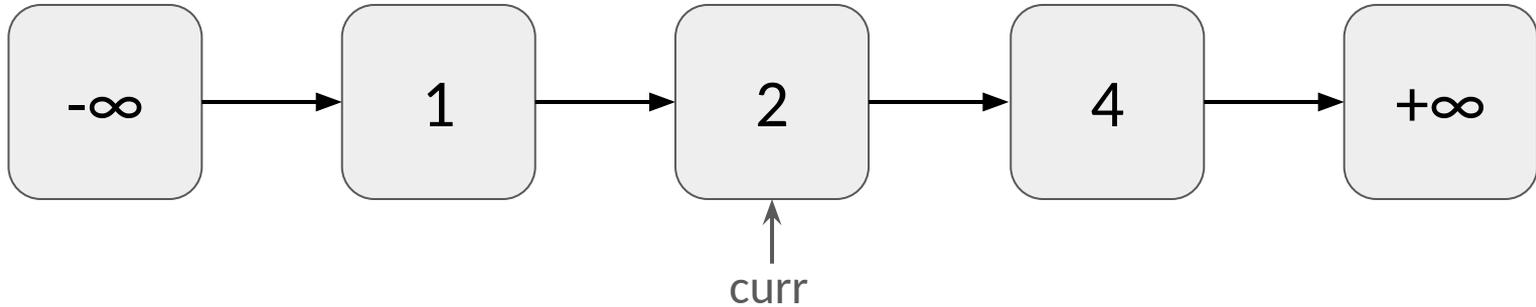
`contains(2) || remove(2) ;; add(2)`

Thread 1 - `contains(2)`: walk is done



`contains(2) || remove(2) ;; add(2)`

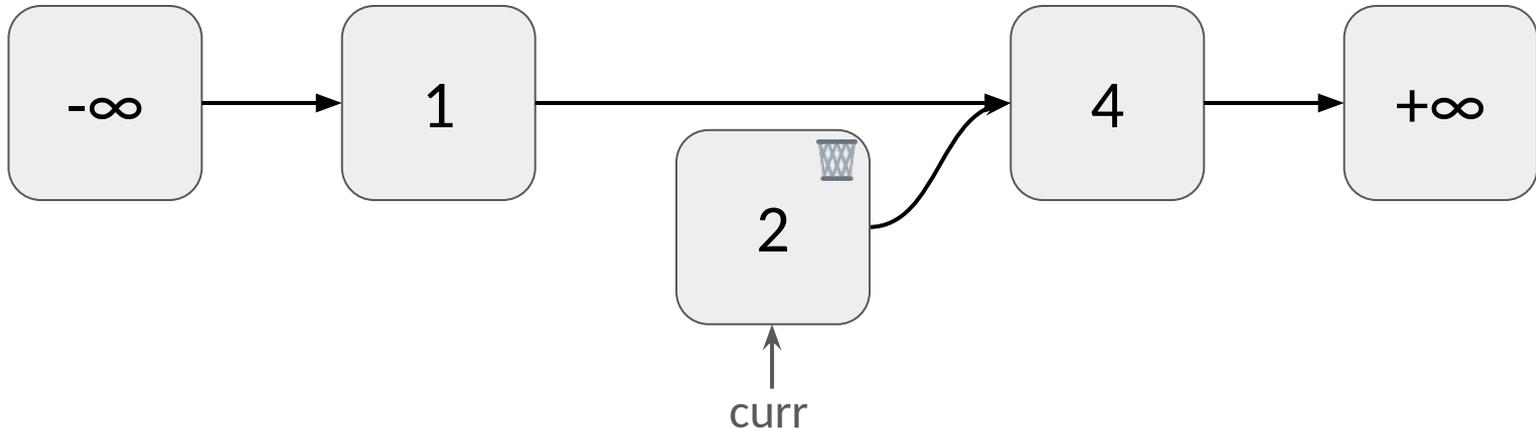
Thread 1 - `contains(2)`: preempted



`contains(2) || remove(2) ;; add(2)`

Thread 1 - `contains(2)`: preempted

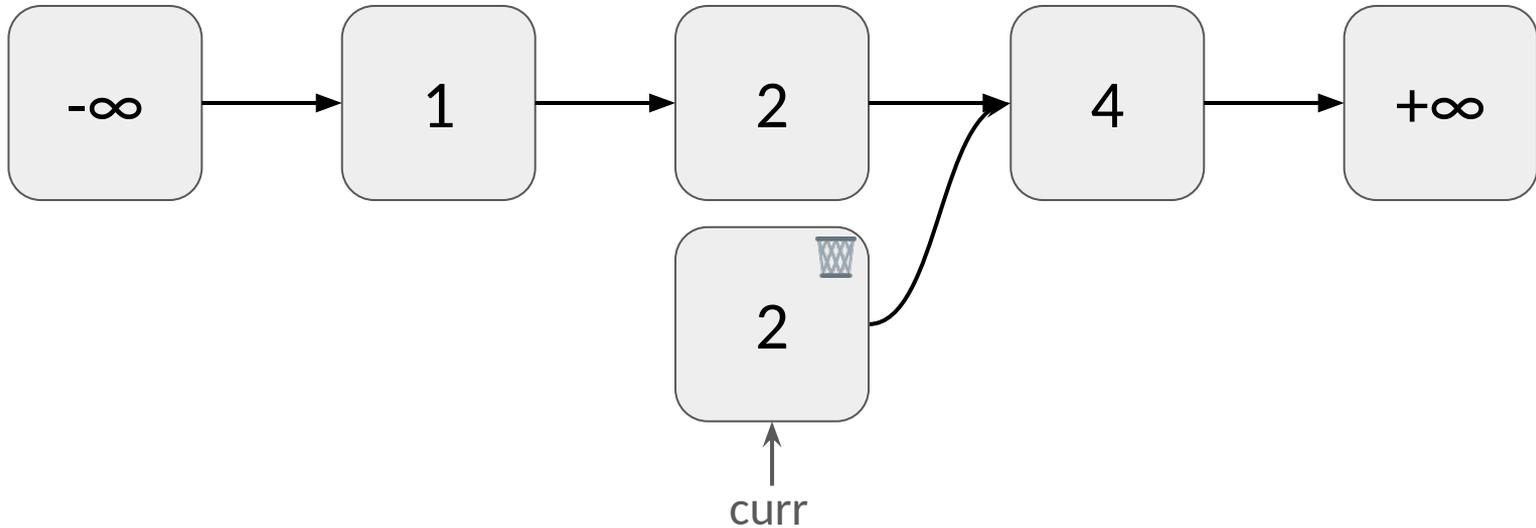
Thread 2 removes 2



contains(2) || remove(2) ;; add(2)

Thread 1 - contains(2): preempted

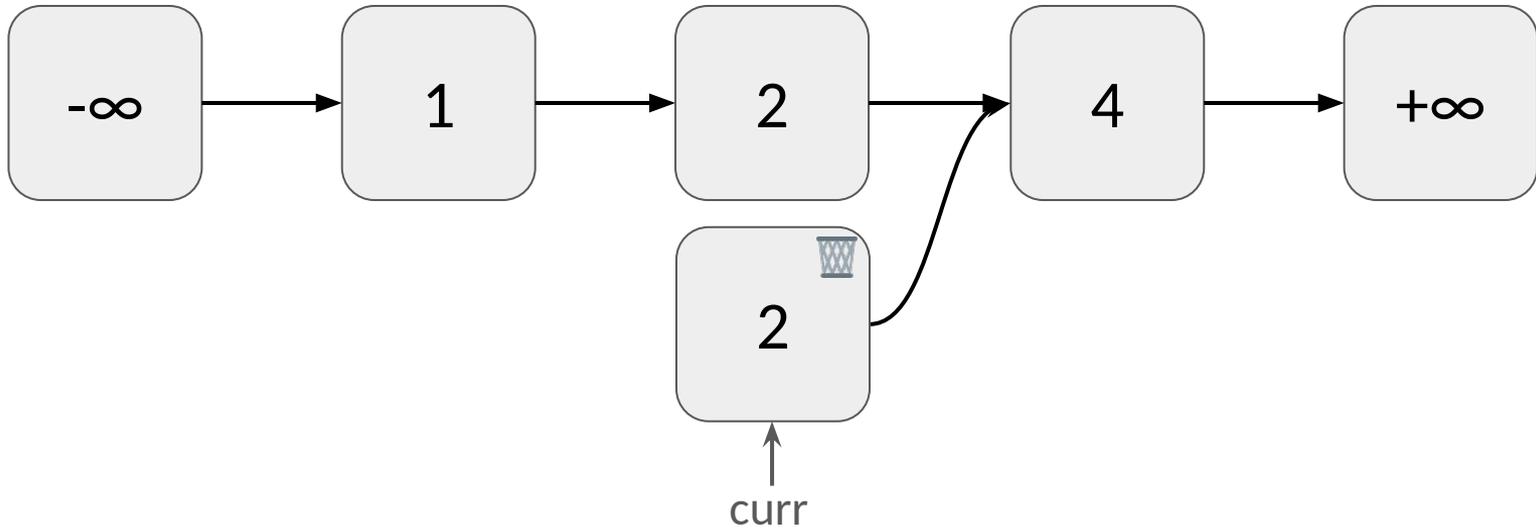
Thread 2 adds 2



contains(2) || remove(2) ;; add(2)

Thread 1 - contains(2): reads marked is true

This cannot be the linearization point, as 2 is in the list!



Linearization Point for contains if Element Is Not in the Set

Earlier of the following points:

Point at which a removed matching entry is found

Point immediately before a new matching entry is added to the list

Proving Correctness in Iris

Implement methods
in HeapLang

Define specifications
in Iris

Prove specifications
in Iris

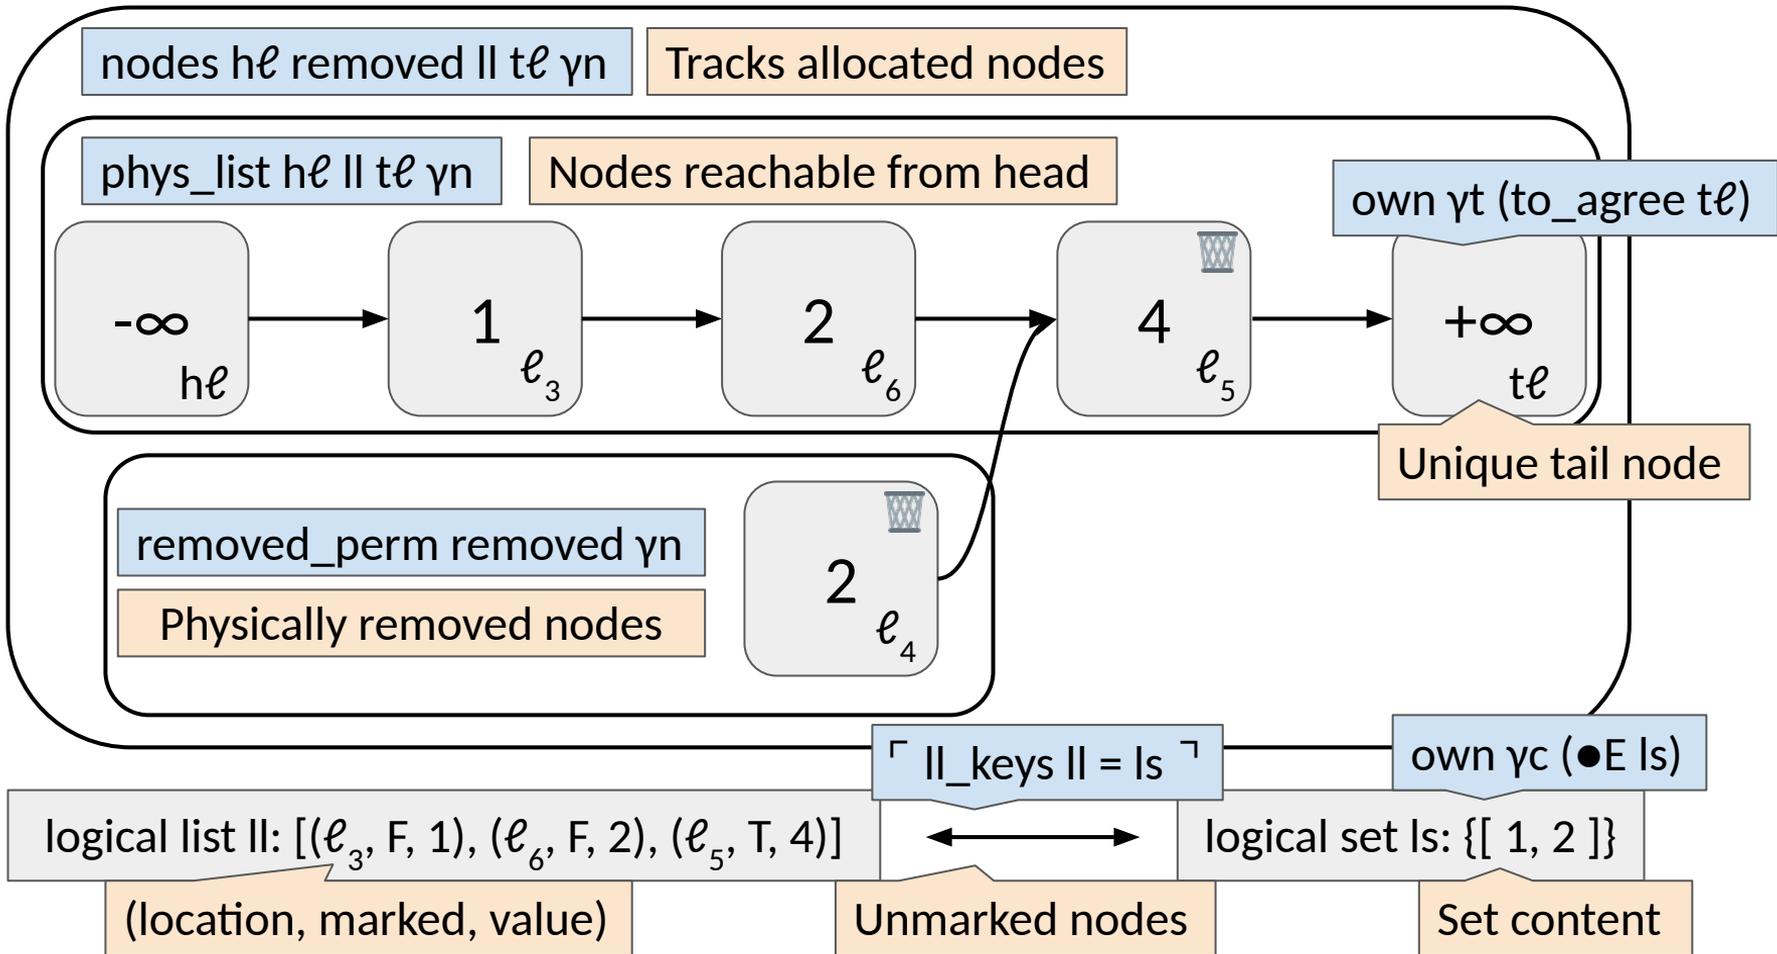
Need to define a set invariant!

Set Invariant

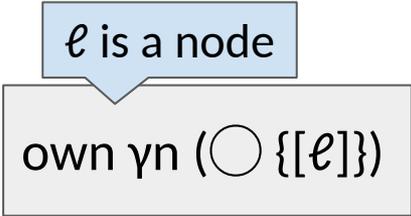
Pointer to set

Ghost State

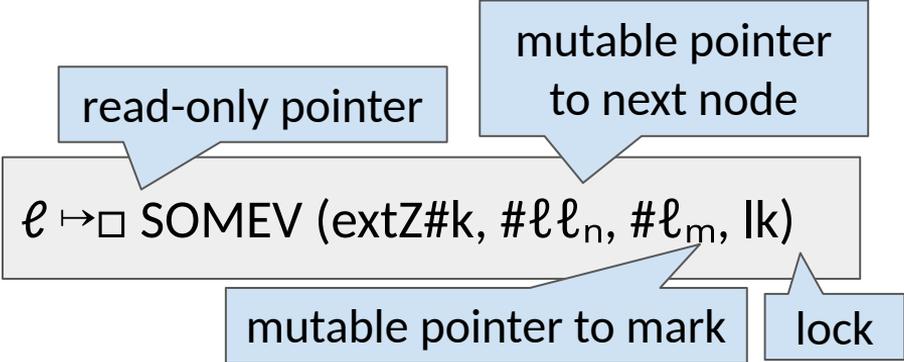
Definition $\text{set_inv} (h\ell : \text{loc}) (\gamma_n : \text{gname}) (\gamma_c : \text{gname}) (\gamma_t : \text{gname}) : \text{iProp} :=$
 $\exists (\text{removed} : \text{gset } \text{loc}) (\text{ls} : \text{gset } \mathbb{Z}) (\text{ll} : \text{llist}) (\text{t}\ell : \text{loc}),$
 $\text{nodes } h\ell \text{ removed } \text{ll } \text{t}\ell \ \gamma_n * \text{own } \gamma_c (\bullet \text{E } \text{ls}) * \text{own } \gamma_t (\text{to_agree } \text{t}\ell) *$
 $\lceil \text{ll_keys } \text{ll} = \text{ls} \rceil * \text{phys_list } h\ell \ \text{ll } \text{t}\ell \ \gamma_n * \text{removed_perm } \text{removed } \gamma_n.$



Modeling Nodes



implies \longrightarrow



+

Often used in proofs

- ℓ is the head node
- $\forall \ell$ is the tail node
- $\forall \ell$ is reachable from the head
- $\forall \ell$ has been physically removed

Specification for add

Key to be added

Lemma add_spec (hℓ : loc) (γn : gname) (γc : gname) (γt : gname) (key : Z) :

is_set hℓ γn γc γt -* **Set Invariant**

Logical Atomicity

<<{ ∇ ∇ (ls : gset Z), set_content γc ls }>>

add #hℓ extZ#(Fin key) @ ↑N **Key added to set**

<<{ set_content γc (ls ∪ {[key]}) | **RET** #(bool_decide (key ∉ ls)) }>>.

Whether key was already in the set

Conclusion and Future Work

What We Did

Implemented *all* methods
in HeapLang

Defined specifications for
all methods in Iris

Proved specifications for
add, and *remove* in Iris

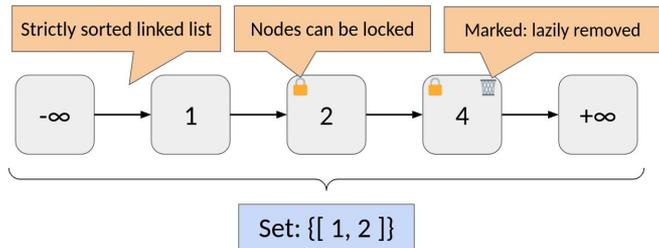
What Could Be Next

Prove *contains* using the
helper pattern

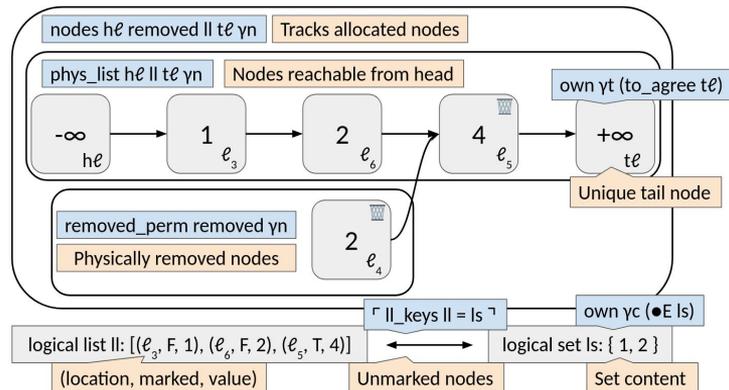
Nested invariants instead of
tracking the set of all nodes?

Use more local instead of
global information

Anatomy of a Lazy List-Based Set

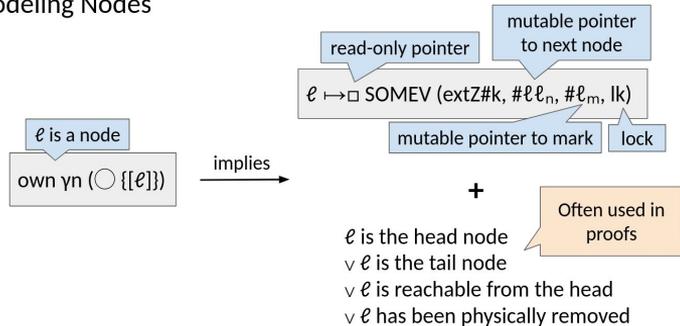


3



42

Modeling Nodes



43

Conclusion and Future Work

What We Did

- Implemented *all* methods in HeapLang
- Defined specifications for *all* methods in Iris
- Proved specifications for *add*, and *remove* in Iris

What Could Be Next

- Prove *contains* using the helper pattern
- Nested invariants instead of tracking the set of all nodes?
- Use more local instead of global information

45